Sampling and Visualizing Creases with Scale-Space Particles

Gordon L. Kindlmann, Raúl San José Estépar, Stephen M. Smith, and Carl-Fredrik Westin
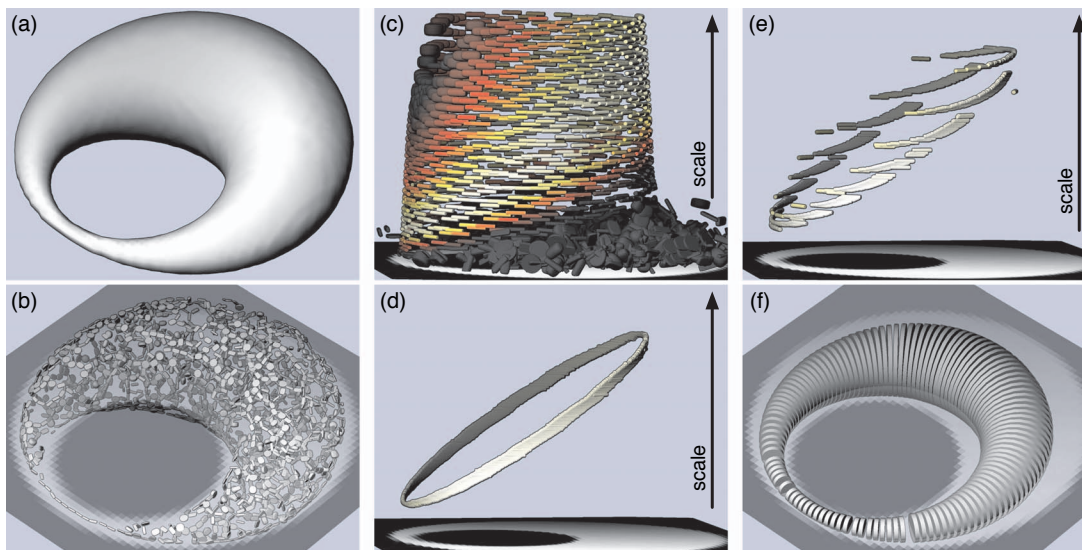

Fig. 1. The shape of a synthetic volume is shown with an isosurface (a). Ridge lines sampled at the original scale (b), shown with a semi-transparent cutting plane, capture only the narrowest portion correctly. Scale-space particles sample the feature manifold through the 4-D scale-space of different blurring levels with spline interpolation across scales (c); the colormap highlights with white the spatially-varying scale of maximal ridge strength. Feature scale is accurately recovered in (d) with an energy term from ridge strength; this is not possible with linear (e) instead of spline interpolation across the same pre-blurred scale samples. Glyphs in (f) show scale with disc radius, indicating a successful sampling of the complete ridge line localized in both space and scale.

Abstract—Particle systems have gained importance as a methodology for sampling implicit surfaces and segmented objects to improve mesh generation and shape analysis. We propose that particle systems have a significantly more general role in sampling structure from *unsegmented* data. We describe a particle system that computes samplings of crease features (*i.e.* ridges and valleys, as lines or surfaces) that effectively represent many anatomical structures in scanned medical data. Because structure naturally exists at a range of sizes relative to the image resolution, computer vision has developed the theory of scale-space, which considers an n-D image as an $(n+1)$-D stack of images at different blurring levels. Our scale-space particles move through continuous four-dimensional scale-space according to spatial constraints imposed by the crease features, a particle-image energy that draws particles towards scales of maximal feature strength, and an inter-particle energy that controls sampling density in space and scale. To make scale-space practical for large three-dimensional data, we present a spline-based interpolation across scale from a small number of pre-computed blurrings at optimally selected scales. The configuration of the particle system is visualized with tensor glyphs that display information about the local Hessian of the image, and the scale of the particle. We use scale-space particles to sample the complex three-dimensional branching structure of airways in lung CT, and the major white matter structures in brain DTI.

Index Terms—Particle Systems, Crease Features, Ridge and Valley Detection, Lung CT, Diffusion Tensor MRI.

◆

1 INTRODUCTION

We present a new approach to sampling features in volumetric data with dynamic particle systems. Particle systems appear most often in computer graphics and visualization for sampling or meshing implicit surfaces from geometric modeling or pre-computed image segmenta-

- Gordon L. Kindlmann is with the Department of Computer Science and the Computation Institute, University of Chicago, E-mail: glk@uchicago.edu.
- Raúl San José Estépar and Carl-Fredrik Westin are with the Department of Radiology, Brigham and Women's Hospital, Harvard Medical School, E-mail: rjosest@bwh.harvard.edu, westin@bwh.harvard.edu.
- Stephen M. Smith is with the Centre for Functional MRI of the Brain, John Radcliffe Hospital, Oxford University, E-mail: steve@fmrib.ox.ac.uk.

tion. One reason for the success of particle systems is that they cleanly separate the task of computing an optimal distribution of vertices on a feature, from the task of connecting the vertices into a polygonal mesh (whereas Marching Cubes [44], for example, does both at the same time). Calculating vertex locations is done with some combination of energy minimization and constraint satisfaction, both of which can be implemented in a straight-forward manner at the level of individual particles, since the required information is limited to a small neighborhood of particles and data values. Particles also enjoy the property of existing in a mathematically continuous representation of the image, so the underlying image resolution and sampling grid have no undue influence. Our work leverages these properties while it extends particle systems into new feature types and image domains.

Effective computational analysis of medical imaging data often requires geometric models of anatomic features. Challenges to feature localization and extraction arise from image noise and limited resolution, as well as from the properties of features themselves, such as branching or variable geometry, or having a wide range of charac-

teristic sizes. *Scale-space* is a computer vision framework for robust feature extraction, in which an n-D image is considered as a $(n+1)$-D stack of images at successive blurring levels, so that large and small features can be detected with equal ease. While theory and methods for scale-space of grayscale, two-dimensional images are well-studied, the promise of scale-space analysis has yet to be realized in practical tools for three-dimensional imaging. Scale-space analysis of non-scalar data, such as diffusion tensor imaging, is largely unexplored.

We propose that particle systems can play a more fundamental role in biomedical visualization and analysis, by sampling complex anatomic features in *unsegmented* data. We focus on ridge and valley features (collectively, *creases*) that can approximate the skeletal cores of objects, rather than the closed isocontours and material boundaries of previous particle systems. Creases are either codimension-one surfaces or codimension-two curves, defined as local extrema with respect to motion along Hessian eigenvectors. Crease surfaces are in general non-orientable, and often appear as multiple disjoint surface patches rather than as closed boundaries. A major challenge associated with using crease features to represent objects in images is that objects naturally occur at a range of sizes relative to the pixel resolution. Scale-space particles locate and sample creases throughout their continuous extent in space and scale. The scale at which a crease appears strongest relates to the characteristic size of the underlying object.

Particle systems are particularly well-suited to crease feature sampling and visualization in scale-space. The mathematical definition of creases directly maps to the rules that individual particles obey to satisfy the constraint of lying within a crease. The implementation of pair-wise particle dynamics, which ultimately drives the system, is the same regardless of feature codimension (surfaces versus lines), and image domain dimension (three or four). In real-world applications, we often don't know where features are, or in how many disjoint pieces they appear; this corresponds to the nature of particles and their lack of explicit topology and connectivity. The point-based representation allows us to simply initialize the system with particles densely sampling the entire image domain, and later reduce the number of particles according to their density on the detected features. At this initial stage of work, however, we are refraining from tackling the task of *meshing* the particle system solution. Subsequent research will investigate the computational geometric considerations for reliably connecting the final particle locations into polygonal feature models.

Our contributions stem from how we design, implement, and apply the combination of particle systems and scale space. At the lowest level, we introduce in Section 3.2 a novel Hermite spline approach for efficiently interpolating through image scales to create a continuous, four-dimensional scale-space. Generalizing the implicit surface constraint previously used for particles, Section 3.3 describes constraints that keep particles within ridges and valleys. We introduce in Section 3.4 inter-particle energy functions that allow particles to either repel or attract along scale, so that the features can either be broadly sampled through scale-space, or be localized at the particular scale that maximizes feature strength. Another novel aspect of our implementation (Sect. 3.5) is that population control (the adding and deleting of particles) is formulated in terms of the same energy minimization that drives the particles towards uniform sampling. We use glyph-based visualizations (Sect. 3.5) to inspect the local properties and over-all configuration of the particle system. Our results (Sect. 4) include visualizations of scale-space particles sampling the branching airways in lung CT, and white matter features in diffusion tensor MRI.

2 RELATED WORK

There are three research areas our work draws upon: scale-space feature extraction (Sect. 2.1), particle systems (Sect. 2.2), and Diffusion Tensor Imaging analysis (Sect. 2.3). Connections to previous work establishing the biomedical utility of crease lines and crease surfaces are drawn in Sections 2.1 and 2.3, respectively.

2.1 Scale-Space Analysis and Crease Lines

The concept of scale and its importance for computer vision led to scale-space theory, which embeds a signal $f(x)$ in a one-parameter

family of functions $L(x;t)$, with diffusion time t inducing Gaussian blurring with standard deviation $\sqrt{t}$ [69, 35, 66]. Florack *et al.* show how principles of linearity, scale-invariance, and well-posed differentiation also imply the Gaussian kernel, independent of a diffusion process [21]. Koenderink notes that significant image features exist at a continuous scale *range* and conceives of image understanding as happening at all scales simultaneously, rather than at a discrete set of blurring levels [35]. A number of scale-space feature-extraction methods have been developed from these ideas. Gauch and Pizer propose multi-resolution analysis for ridge and valley lines by projecting the Hessian at different scales into the level curve tangents, while pointing out the difficulties caused by working on a discrete grid [24]. Eberly presents a general description for ridge detection in n-D signals without focusing on the scale-dependent nature of the problem [17], an extension that was theoretically described by Miller and Furst [49].

A prominent application of ridge and valley analysis is the extraction of tubular structures, like airways and blood vessels, from medical imaging [2, 22, 3]. Lindeberg studies the problem of finding the optimal scale of features (including ridges) by maximizing measurements of feature strength along scale [41]. There is, however, no explicit enforcement of the spatial continuity of feature scale, which we address with a novel scale-attractive inter-particle energy function (Sect. 3.4). An initial application of our work is in airway analysis in lung CT, a modality commonly used to assess various pulmonary conditions [59]. For some diseases, such as chronic obstructive pulmonary disease (COPD), the narrowest airways (near the limits of imaging resolution) are the most clinically significant [28]. This significance creates a need to accurately sample features at the smallest scales, which we accomplish with smooth spatial-domain reconstruction (Sect. 3.2.2), while our scale interpolation (Sect. 3.2.1) captures the connected branching structures continuously across scale.

In the visualization community, there has been some work in exploring scale-space for vector flow-field visualization and feature extraction. Bauer and Peikert use scale-space concepts to recover vortex lines via the parallel vector operator [54] in smoothed volumes [6]. Klein and Ertl track vector field critical points through discretely sampled scales, to better filter out noise effects [34]. In the image processing community, non-linear filtering of tensor fields has been studied [13, 65, 9], which is conceptually related to scale-space by its mathematical connection to earlier work by Perona and Malik on non-linear scale-space [55], but there has been relatively little study of the scale-space features in tensor fields. Florack and Astola argue that a scale-space treatment of diffusion tensor fields should respect commutativity of blurring and tensor inversion, and describe the mathematical components of the resulting non-linear scale-space [19]. We contribute to this previous work by explicitly locating and uniformly sampling the geometry of image features in tensor fields, across their full extent in space and scale, using particle systems.

2.2 Particle Systems

Particle systems have been studied in computer graphics for various applications; the application most closely related to ours is surface representation. Szeliski and Tonnesen model synthetic surfaces using systems of oriented particles with non-rotationally symmetric energy profiles that self-assemble into surfaces [62]. Witkin and Heckbert used a system of particles with rotationally symmetric energy profiles to interactively sample and sculpt implicit surfaces [70, 27]. Continued refinements have enabled further applications to shape modeling and other applications (*e.g.* see [61] and references therein).

Other studies outside computer graphics use particle systems in a more data-driven way for scientific or biomedical applications, including interactive scientific visualization [52], anisotropic mesh generation [8, 72], feature-aware mesh smoothing [30], visualization of Smoothed Particle Hydrodynamics [42], and illustrative volume visualization [10]. For medical image analysis, Cates *et al.* develop entropy-based particle systems that simultaneously sample surfaces across multiple volumes, efficiently determining surface correspondences and modes of shape variation [12, 11]. Isosurface sampling is a prominent application of data-driven particle systems [57, 15]. This

has been studied in detail by Meyer *et al.*, with curvature-dependent sampling [48] and sampling of isosurfaces in high-order finite element grids [46]. Later work connects particles into meshes representing segmented object surfaces [45] and the topology of multi-material volumes [47]. We follow the example of Meyer *et al.* in first determining particle motion, interaction, and population control, while leaving the separate (and significant) computational geometry task of computing vertex connectivity to later work.

A limited amount of previous work shares our approach of using particles to accomplish feature sampling in *unsegmented* data. Szeliski *et al.* direct their oriented particles [62] with an edge detection energy term to form surface models of segmented and unsegmented volume objects [63]. More recently, Jalba *et al.* attach fixed electrostatic charges to the edge pixels or voxels in an image, which then attract mutually repelling and mobile particles that converge in tracing out region boundaries [29]. Other previous work, on active contours, creates long-range attractive forces based on gradient vector flows to ensure insensitivity to initialization [71]. A virtue of particles demonstrated by Jalba *at al.* [29] is that they can be initialized at every image pixel, ensuring that no significant edge features are missed, which side-steps the problem of sensitivity to initial conditions. However, the magnitudes of the data-driven force (for feature localization) and the inter-particle force (for uniform sampling) must be calibrated so that edges attract particles more strongly than particles repel each other.

Relative to this previous work, our contributions are (1) employing spatial *constraints*, rather than forces or potentials, to maintain particles within features, and (2) moving particles smoothly through scale to sample features in continuous four-dimensional scale-space. With its emphasis on feature detection and localization, our use of particles is conceptually related to active contours and deformable models that use energy-minimizing parameterized curves or surfaces to capture image features [31]. Our particles, however, are not connected in any pre-determined topology, and are not described by any particular model of feature geometry, so they can sample codimension-one surfaces as easily as codimension-two lines.

2.3 Diffusion Anisotropy Analysis and Crease Surfaces

Diffusion tensor magnetic resonance imaging (DTI) uses a tensor model of diffusion-weighted MRI to describe directional tissue structure [4]. The coherent organization of axons in nervous tissue contributes to the anisotropic diffusion of water within it, enabling research on white matter organization in health and disease [37]. Much interest in DTI stems from using integrals of the principal eigenvector of the diffusion tensor (termed *tractography*) to approximate axonal connectivity between different brain areas [14, 5]. On the other hand, much of the neuroscientific research using DTI focuses on scalar-valued tensor invariants such as *fractional anisotropy* (FA), which are local measures of microstructural organization [56, 7]. Many applications of DTI amount to correlating statistically significant differences in FA with various neurological or psychiatric conditions. For example, Kubicki *et al.* review such methods in schizophrenia [36]. Detecting local changes in tissue properties (through measures like FA) is a fundamental neuroimage analysis task, separate from but complementary to connectivity measurements from the tractography methods typically studied in the visualization community.

To improve the statistical validity of voxel-based group studies of FA variations, Smith *et al.* developed Tract-Based Spatial Statistics (TBSS) [60]. TBSS rasterizes FA ridge surfaces in a high-resolution, group mean, tensor image to form a white matter skeleton for statistical tests. Other work has similarly analyzed the differential structure of FA to measure anatomic features in DTI. Goodlett *et al.* use an FA ridge surface strength metric to improve registration in group studies [26], analogous to previous registration work for scalar images [43]. Kindlmann *et al.* use a modification of Marching Cubes to extract polygonal models of FA ridge and valley surfaces [33]. Tricoche *et al.* study general invariant crease lines in tensor fields and detect some white matter structures as FA ridge lines [64]. Schultz *et al.* present a new approach to extracting crease surfaces that improves on previous voxel-based methods [23, 33] in speed and topological

correctness, which they apply to DTI anisotropy measures [58].

Our work relates to these previous results in two ways. First, along with lung airway segmentation (Sect. 2.1), extensions to TBSS [60] are a driving application. We plan to expand the target of the TBSS white matter skeletonization from the smooth group mean image to single-subject images. Initial experiments with this, as well as experience with anisotropy crease surfaces [33] and lines [64], have highlighted the problem with picking a single *scale* to measure derivatives: features at significantly smaller or larger scales are poorly captured or not detected at all. Second and more generally, we suspect that the essential scale-dependence of derivative measurements [20], combined with the combinatorial complexity of extending Marching Cubes-type approaches to four-dimensional scale-space lattices [23, 18], may have limited the adoption of previous crease surface extraction methods, as compared to the ubiquity of isosurfaces from Marching Cubes itself [44]. We hope our initial results with scale-space particles can expand the exploration and application of scale-space creases in real-world three-dimensional datasets.

3 METHODS

3.1 Methods overview

The particles in our system are moved, added, and deleted to minimize their collective energy. There are two energy sources: *particle-image* and *inter-particle*. A combination of constraints and *particle-image* energy (Sect. 3.3) spatially binds particles to creases, while allowing them to move along scale toward maximal feature strength. The constraints and particle-image energy are conceptually similar to the external energy terms of active contours or the data-attachment terms of level sets, which locate image features. *Inter-particle* energy (Sect. 3.4) creates uniform feature sampling. Scale-space particles lack an analog to the internal energy of active contours, or the curvature terms of level sets, which ensure smoothness in the feature representation. The particle system does not maintain any particular geometric or topological configuration, as individual particles are independent and unconnected. Rather, constraints keep particles within features, which are smooth curves or surfaces, due to their mathematical definition as well as our scale-space reconstruction (Sect. 3.2).

We minimize the energy $\mathscr{E}$ of a set of N particles $\{(\mathbf{x}_i, s_i)\}$ within the four-dimensional image $f(\mathbf{x}, s)$, where $\mathbf{x}$ is three-dimensional spatial location, and s is univariate scale, by numerically solving

$$\underset{\{(\mathbf{x}_i, s_i)\}, N}{\text{argmin}} \ \mathscr{E} = \underset{\{(\mathbf{x}_i, s_i)\}, N}{\text{argmin}} \ (1-\alpha) \sum_{i=1}^{N} E_i + \frac{\alpha}{2} \sum_{i,j=1}^{N} E_{ij}. \quad (1)$$

The particle-image energy E_i (Sect. 3.3) is determined entirely by the local properties of the image f at $(\mathbf{x}_i, s_i)$. The symmetric inter-particle energy $E_{ij} = E_{ji}$, $E_{ii} = 0$ (Sect. 3.4) is independent of location in the image domain. The balance between inter-particle and particle-image energies is determined by the parameter $\alpha \in [0, 1]$, which is roughly analogous to the user-defined weighting in level-set methods between the smoothness and data-attachment terms.

Not apparent in Eq. 1 are *constraints*, which force particles to remain within features, but which do not themselves contribute to the system energy. Alternatively, without constraints, particles can be steered by an energy that is locally minimal at the feature [63, 29]. For four-dimensional scale-space, the choice between constraints and particle-image energy exists independently for both the three-dimensional spatial domain and the one-dimensional scale domain. Following previous work on implicit surface sampling with particles, we choose to *spatially* constrain particles to lie within crease features. This is enabled by their simple mathematical definition (Sect. 3.3), combined with analytic spatial derivatives of the image (Sect. 3.2). Along the scale axis, however, we use particle-image energy to push particles toward the scales of maximal feature strength. We have found that crease strength measures (Sect. 3.3) are effective but not as stable as the spatial location of the crease itself, so we balance a particle-image energy based on crease strength with an inter-particle energy that attracts particles to each other along scale, producing a feature sampling with better spatial continuity of scale.

3.2 Interpolation in Scale-Space

We reconstruct continuous scale-space from discrete samples in space and scale. Each scale sample is some blurring of the original volume data at its original raster resolution. The range of blurring scales is sampled non-uniformly with S scales between $s = 0$ (no blurring) to $s = s_{\max}$, the largest expected feature scale, which should be known from the imaging application domain, where s is the standard deviation of a Gaussian blurring kernel. The scale-space location of a particle and the spatial reconstruction kernels determines the neighborhood of image samples that fall within the spatial convolution support and the two samples along scale that bracket the particle's scale position. We first interpolate, along scale, the spatial support of the particle (Sect. 3.2.1), and then use spatial convolution to reconstruct values and spatial derivatives (Sect. 3.2.2). Following our previous work on anisotropy creases [33, 64], tensors are interpolated component-wise, which facilitates *analytic* computation of spatial derivatives of anisotropy. Non-linear tensor interpolation methods are also possible (*e.g.* [1, 38]) but lack a similarly accurate way of computing anisotropy derivatives, and with necessarily greater computational expense.

3.2.1 Scale Interpolation

Sampling scale-space by convolving with a blurring kernel whose width equals the scale sample position is completely accurate but prohibitively slow at large scales. Alternatively, one can interpolate between a large set of images pre-blurred at varying scales, and then use a fixed-size spatial support reconstruction in space. For large three-dimensional images, however, the number of such scale samples is limited by available memory. Balancing memory efficiency with reconstruction accuracy is a challenge for volumetric scale-space.

Some previous work has used linear interpolation between pre-blurred images at adjacent scale samples [34]. We introduce a more accurate method using cubic Hermite splines, based on the underlying properties of a particular blurring kernel. Consider first a continuous one-dimensional signal $f(\cdot)$ undergoing isotropic diffusion for time t, which is equivalent to convolving f with a Gaussian g.

$$L(x;t) = (f(\cdot) \star g(\cdot;t))(x) = \int g(\xi;t)f(x-\xi)d\xi, \quad (2)$$

$$g(\xi;t) = \exp(-\xi^2/2t)/\sqrt{2\pi t}. \quad (3)$$

The time-derivative of $L(x;t)$ is the second spatial derivative [35, 40]

$$\partial L(x;t)/\partial t = \tfrac{1}{2}\partial^2 L(x;t)/\partial x^2. \quad (4)$$

For discretely sampled $f[\cdot]$, we replace the continuous Gaussian $g(\cdot;t)$ with its discrete analog $K[\cdot;t]$, created by Lindeberg [39].

$$L[i;t] = (f \star K[\cdot;t])[x] = \sum_n K[n;t]f[i-n], \quad (5)$$

$$K[n;t] = \exp(-t)I_n(t), \quad (6)$$

where I_n is the modified Bessel function of integer order n. By the recursive properties of I_n and the design of $K[\cdot;t]$, the continuous derivative with respect to time is found by second central differences [39]

$$\partial L[i;t]/\partial t = \tfrac{1}{2}(L[\cdot;t] \star \begin{bmatrix} 1 & -2 & 1 \end{bmatrix})[i]. \quad (7)$$

With a change of variables from diffusion time t to spatial scale $s = \sqrt{t}$,

$$\partial L[i;s^2]/\partial s = s\,(L[\cdot;s^2] \star \begin{bmatrix} 1 & -2 & 1 \end{bmatrix})[i]. \quad (8)$$

This analytical derivative $\partial L/\partial s$ enables our cubic Hermite spline interpolation. As pointed out by Lindeberg, the same property does not hold for blurring with a discretely sampled continuous Gaussian.

We pre-compute blurrings of the original data $f[\cdot,\cdot,\cdot]$ with a set of S different three-dimensional discrete Gaussians $K[\cdot,\cdot,\cdot;s_\ell^2]$, with non-uniformly spaced scale positions $\{s_\ell\}_{\ell=1}^S$. Blurring at scale s_ℓ to produce volume $f[\cdot,\cdot,\cdot,\ell]$ is computed by

$$f[i,j,k,\ell] = \sum_{m,n,o} K[m,n,o;s_\ell^2]\,f[i-m,j-n,k-o], \quad (9)$$

$$K[m,n,o;t] = \exp(-t)^3 I_m(t)I_n(t)I_o(t). \quad (10)$$

To interpolate between $f_0 = f[i,j,k,\ell_0]$ and $f_1 = f[i,j,k,\ell_1]$ at s, we find s_{ℓ_0} and s_{ℓ_1} such that $s_{\ell_0} \le s < s_{\ell_1}$, let $\Delta = (s - s_{\ell_0})/(s_{\ell_1} - s_{\ell_0})$, and compute the 3-D analog to (8) at s_{ℓ_0} and s_{ℓ_1},

$$f_0' = s_{\ell_0}(f[\cdot,\cdot,\cdot,\ell_0] \star M)[i,j,k], \quad (11)$$

$$f_1' = s_{\ell_1}(f[\cdot,\cdot,\cdot,\ell_1] \star M)[i,j,k], \quad (12)$$

where M is the standard $3 \times 3 \times 3$ discrete Laplacian mask. Then the cubic Hermite spline between scales is

$$f[i,j,k](s) = \begin{bmatrix} \Delta^3 & \Delta^2 & \Delta & 1 \end{bmatrix} \begin{bmatrix} f_0' + f_1' + 2f_0 - 2f_1 \\ -2f_0' - f_1' - 3f_0 + 3f_1 \\ f_0' \\ f_0 \end{bmatrix}. \quad (13)$$

For comparison, linear interpolation would simply be

$$f[i,j,k](s) = (1-\Delta)f_0 + \Delta f_1. \quad (14)$$

For each particle, we interpolate along scale only those image samples that fall within the support of the spatial reconstruction kernels.
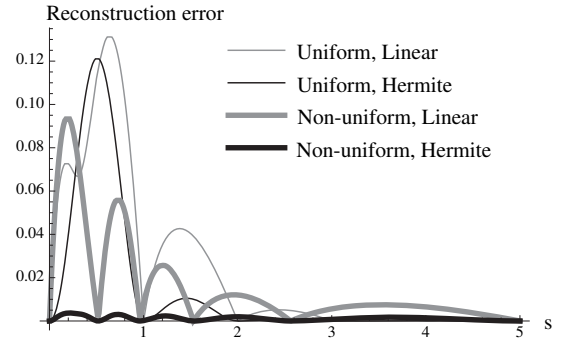

Fig. 2. Comparison of error with linear and cubic Hermite scale-space interpolation for $s \in [0,5]$, reconstructed with 6 samples, placed either uniformly or non-uniformly. Scale-space particles use Hermite interpolation with optimized non-uniform samples for greatly improved accuracy.

We measure scale-space reconstruction error as the squared difference between the true blurring kernel (the point-spread function) $K[\cdot,\cdot,\cdot;s^2]$, and an approximate kernel interpolated between $K[\cdot,\cdot,\cdot;s_{\ell_0}^2]$ and $K[\cdot,\cdot,\cdot;s_{\ell_1}^2]$, averaged over the support of K. Figure 2 depicts reconstruction error from linear and spline interpolation, with uniform and non-uniform scale samples. The locations of the non-uniform scale samples were computed by a brute-force gradient-descent minimization of the mean squared error, numerically integrated over all scales. For a given reconstruction accuracy, our new Hermite spline approach requires significantly fewer pre-blurred volumes $f[\cdot,\cdot,\cdot,\ell]$, which greatly reduces memory use and makes continuous scale-space analysis practical for large 3-D data.

3.2.2 Spatial Interpolation and Differentiation

Following scale interpolation, we compute spatial interpolation and differentiation by convolving with separable piecewise polynomial kernels. From the framework of Möller *et al.* [50], we use the six-sample support, fifth-degree polynomial, C^3 continuous approximation filter, with fourth-degree error (which exactly reconstructs cubic polynomials). The data is first de-convolved so subsequent first convolution interpolates the original data. We have used a high order of continuity and accuracy so that the creases (themselves determined by second derivatives) are smooth and ripple-free. Continuous values $f(\mathbf{x}, s)$ are reconstructed by convolving discrete data samples $f[i,j,k](s)$ with the separable filter $R(x,y,z) = r(x)r(y)r(z)$ for the chosen univariate filter $r(x)$. Partial derivatives are measured by convolving $f[i,j,k](s)$ with the corresponding partial derivative of R.

The final reconstruction step is scale-normalization of derivatives, as described by Lindeberg [41]. Intuitively, at higher blurring levels,

the data signal naturally becomes flatter, so spatial derivative magnitudes decrease. To make meaningful comparisons between derivatives at different scales, they are normalized by scale:

$$\tilde{\nabla} f(\mathbf{x}, s) = s \, \nabla f(\mathbf{x}, s), \tag{15}$$

$$\tilde{\mathbf{H}} f(\mathbf{x}, s) = s^2 \, \mathbf{H} f(\mathbf{x}, s). \tag{16}$$

We use the tilde to denote a quantity associated with a scale-normalized derivative. For DTI analysis, scale-normalization is applied to the derivatives of the individual tensor components, not to the derivatives of the derived non-linear tensor attributes like FA.

3.3 Particle-Image Energy and Crease Feature Definition

We use Eberly's definition of creases in terms of the gradient $\mathbf{g} = \nabla f$ and Hessian $\mathbf{H}$ of a smooth function f [16]. For DTI analysis, the gradient and Hessian of fractional anisotropy (FA) are computed *analytically* from the gradients and Hessians of the individual tensor components, using formulae in [33]. Let $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ be unit-length eigenvectors of $\mathbf{H}$ with corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$. Crease surfaces are defined by $\mathbf{g} \cdot \mathbf{v}_i = 0$ and crease lines by $\mathbf{g} \cdot \mathbf{v}_i = \mathbf{g} \cdot \mathbf{v}_j = 0$, with i, j set according to crease type (ridge or valley). For example, ridge surfaces are defined by $\mathbf{g} \cdot \mathbf{v}_3 = 0$: f is maximized with respect to motion along the direction of greatest negative curvature. Crease feature strength h is quantified by the magnitude of certain eigenvalues of the scale-normalized Hessian $\tilde{\mathbf{H}}$, *e.g.* $-\tilde{\lambda}_3$ is ridge surface strength since a very negative λ_3 indicates strong downward concavity. Table 1 summarizes crease features with the defining dot products, required eigenvalue sign, feature strength h, and an *approximate* crease tangent projection $\mathbf{T}$ used in particle updates (Sect. 3.5). Particle-image energy E_i (1) is defined in proportion to feature strength h by

$$E_i = e(f, \mathbf{x}_i, s_i) = -\gamma h(\mathbf{x}_i, s_i). \tag{17}$$

The negation in (17) makes particles locate higher crease strengths while decreasing their energy. Although total energy (1) has the α parameter to balance inter-particle and particle-image energy, we include $\gamma > 0$ in (17) to convert from units of the Hessian (length^{-2}) to units of energy, and to control for the dynamic range of image values. Section 3.5 describes a way to set γ automatically.

The spatial crease constraints are implemented in terms of definitions in Table 1. If our feature were the implicit surface $f = 0$, the gradient $\mathbf{g}$ would be used in Newton-Raphson surface localization [70, 45]. Creases are local extrema, so we perform a minimization or maximization, rather than a root-finding, to locate them. We restrict gradient ascent (for ridges) or descent (for valleys) to the perpendicular space of the approximate crease tangent $\mathbf{T}$.

$$\mathbf{x} \leftarrow \begin{cases} \mathbf{x} + c(\mathbf{I} - \mathbf{T})\mathbf{g} & \text{for ridges} \\ \mathbf{x} - c(\mathbf{I} - \mathbf{T})\mathbf{g} & \text{for valleys} \end{cases} \tag{18}$$

Particles are constrained to creases by iterating (18) with an adaptive stepsize c until the length of updates to spatial position $\mathbf{x}$ falls below ε_c times the voxel size; current work uses $\varepsilon_c = 0.001$. Note that at convergence of (18), $(\mathbf{I} - \mathbf{T})\mathbf{g} = \mathbf{0}$ and $\mathbf{I} = \sum_i \mathbf{v}_i \otimes \mathbf{v}_i$ imply the $\mathbf{g} \cdot \mathbf{v}_i = 0$ crease definition. The constraint is enforced per particle per iteration, conceptually happening instantaneously relative to the particle motion that minimizes the total system energy (1). (18) does not exactly transport

	R L	R S	V L	V S
Definition	$\mathbf{g} \cdot \mathbf{v}_2 = 0$ $\mathbf{g} \cdot \mathbf{v}_3 = 0$	$\mathbf{g} \cdot \mathbf{v}_3 = 0$	$\mathbf{g} \cdot \mathbf{v}_1 = 0$ $\mathbf{g} \cdot \mathbf{v}_2 = 0$	$\mathbf{g} \cdot \mathbf{v}_1 = 0$
λ sign	$\lambda_3 \leq \lambda_2 < 0$	$\lambda_3 < 0$	$0 < \lambda_2 \leq \lambda_1$	$0 < \lambda_1$
Strength h	$-\tilde{\lambda}_2$	$-\tilde{\lambda}_3$	$\tilde{\lambda}_2$	$\tilde{\lambda}_1$
Tangent $\mathbf{T}$	$\mathbf{v}_1 \otimes \mathbf{v}_1$	$\mathbf{v}_1 \otimes \mathbf{v}_1$ $+\mathbf{v}_2 \otimes \mathbf{v}_2$	$\mathbf{v}_3 \otimes \mathbf{v}_3$	$\mathbf{v}_2 \otimes \mathbf{v}_2$ $+\mathbf{v}_3 \otimes \mathbf{v}_3$

Table 1. Summary of ridge (R) and valley (V) surfaces (S) and lines (L).

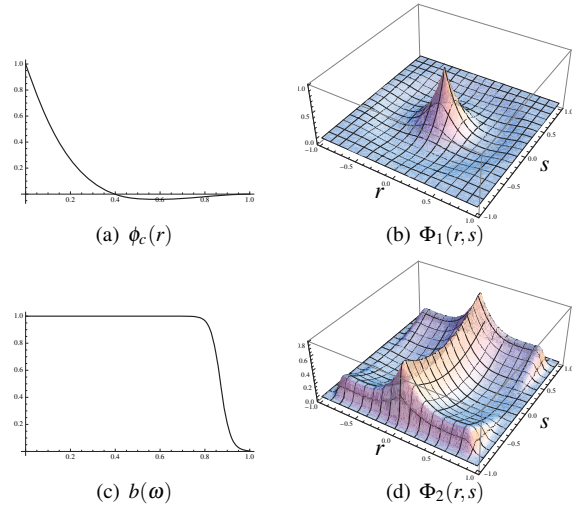(a) $\phi_c(r)$ (b) $\Phi_1(r, s)$

(c) $b(\omega)$ (d) $\Phi_2(r, s)$

Fig. 3. Inter-particle energy starts with radial profile $\phi_c(r)$ in (a) (with $w = 0.6$, $d = -0.1$ for illustrative purposes), used in scale-repelling $\Phi_1(r, s)$ in (b). Butterworth function $b(\omega)$ (with $b_{\text{ord}} = 20$, $b_{\text{cut}} = 0.87$) in (c) windows the scale-attractive $\Phi_2(r, s)$ in (d), shown with $\beta = 0.5$.

particles onto creases along a perpendicular, but it is accurate enough for the small position updates that arise in our system.

In data with any amount of noise, a crease feature that usefully models some underlying anatomic structure will, at its perimeter, gradually become meaningless as it extends into the background. To keep sampling within well-defined features, we impose a user-specified global minimum crease strength $h_{\min}$, which limits the sampled feature extent. Strength thresholding tends to separate the crease into a larger number of disjoint pieces, each with an open boundary. We currently do not explicitly locate the boundary where $h(\mathbf{x}) = h_{\min}$ but simply delete particles that fall below the strength minimum. This has implications for inter-particle spatial repulsion, discussed next.

3.4 Inter-Particle Energy

The inter-particle energy determines how particles interact with each other. Previous work has explored varying total potential [48], radius [15, 45], or orientation [62], based on local data properties. Our focus is an initial exploration of particles in scale-space, rather than the optimization of feature sampling, so currently we consider only rotationally symmetric and fixed-size spatial energy profiles. The inter-particle energy at $(\mathbf{x}_i, s_i)$ due to $(\mathbf{x}_j, s_j)$ is

$$E_{ij} = \Phi(r_{ij}, s_{ij}) = \Phi\left(\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\sigma_r}, \frac{s_i - s_j}{\sigma_s} \right), \tag{19}$$

where σ_r and σ_s are user-specified radii along space and scale. These radii are set based on the desired sampling resolutions in space and scale, likely with recognition of the intrinsic resolution of the image itself. We present two energy functions, Φ_1 for sampling the full scale-extent of a feature, and Φ_2 for localizing the scale of maximal feature strength. Both have finite support in r and s.

We start with the single spatial radial energy profile, a tunable C^2 piecewise cubic function $\phi_c(r)$ with a potential well of depth $d < 0$ at radius $r = w$, shown in Fig. 3(a).

$$\phi_c(r) = \begin{cases} 1 + \frac{3(d-1)|r|}{w} - \frac{3(d-1)|r|^2}{w^2} + \frac{(d-1)|r|^3}{w^3} & 0 < |r| < w, \\ d - \frac{2d(|r|-w)^3}{(w-1)^3} - \frac{3d(|r|-w)^2}{(w-1)^2} & w \leq |r| < 1, \\ 0 & |r| \geq 1. \end{cases} \tag{20}$$

The negative potential well in $\phi_c(r)$ serves two purposes. First, as creases are often disjoint open surfaces or curves, some attraction maintains a tighter sampling pattern than created by pure repulsion (*i.e.* $\phi(r) \geq 0 \forall r$), as with the Coulomb [29], cotangent-based [48], and Gaussian [70] profiles used previously. This is the main way in

which our crease sampling goal differentiates our inter-particle energy from previous energies for sampling closed surfaces (with the exception of the infinite-support Lennard-Jones potential, also with a potential well [62]). Second, a potential well in $\phi(r)$ allows population control to be expressed and implemented as energy minimization. As described in Section 3.5, total system energy (1) can be lowered by adding particles at the edge of a well-sampled area or by filling in holes in the distribution. We use $w = 0.6$ and $d = -0.002$ for this work, as determined by experience, but we have not found the system to be very sensitive to small changes in these values.

Our first energy function Φ_1, shown in Fig. 3(b), evenly distributes particles along both space and scale.

$$\Phi_1(r,s) = \phi_c(\sqrt{r^2 + s^2}). \tag{21}$$

When using Φ_1, we set α in (1) to 1.0 (turning off particle-image energy), so that we can visualize the full extent of each feature along scale and see how feature position changes as a function of scale.

However, with Φ_2, shown in Fig. 3(d), particles *attract* each other along the scale while repelling in space, with a blending of two terms controlled by $\beta \in [0,1]$.

$$\Phi_2(r,s) = (1-\beta)\phi_c(r)b(s) + \beta b(r)b(s)s^2, \tag{22}$$

$$b(\omega) = \frac{1}{1 + (\frac{\omega}{b_{\text{cut}}})^{2b_{\text{ord}}}}. \tag{23}$$

Together with $\alpha < 1$, the particle-image energy (previous section) will draw particles along scale so they coalesce at the scale-space position of peak feature strength while maintaining uniform spatial distribution. Attraction along scale keeps spatially adjacent particles at similar scales, which makes the feature sampling more robust. We use the transfer function of a low-pass Butterworth filter [25] $b(\omega)$ (Fig. 3(c)) to enforce locality along r and s, not because of any spectral properties, but simply because it is a smooth, tunable approximation to a box function. Current work has used $b_{\text{ord}} = 20$ and $b_{\text{cut}} = 0.87$.

3.5 Initialization, Computation, and Visualization

The goal of initialization is to sample all features, wherever they may be, with at least one particle. Following the general strategy of Jalba *et al.* [29], for *every* voxel, we seed particles at multiple locations along scale (no more than the S pre-blurred volumes). From their initial seed location, particles undergo pre-constraint satisfaction (Sect. 3.3) but are deleted if they travel more than two voxel widths before convergence. Particles are then added to the system only if their support does not overlap with more than 50% of a previously added particle's support, and if their feature strength exceeds the user-specified threshold h_{min}. The result is a large number of initial particles placed unevenly on whatever creases are present in the scale-space domain. Although an abundance of particles initially slows the particle system, it guarantees the sampling of all features; the alternative of seeding fewer particles risks losing features in exchange for speed, a risk that compromises our ultimate goal of reliable feature extraction.

The mechanics of particle system computation are based on previous work [70, 48, 46]. Particles are organized into bins that cover the domain, at least $2\sigma_r$ along each edge in space, and $2\sigma_s$ along scale, ensuring that all possible neighbors of a particle are found within its own bin or in a neighboring bin. Each iteration consists of an asynchronous update of each particle and a reapplication of the feature constraint. A particle's update is found by projecting the negative gradient of its energy $-\nabla \mathcal{E}_i$ onto the approximate local feature tangent $\mathbf{T}$, and finding a step along $-\mathbf{T}\nabla \mathcal{E}_i$ that decreases energy $\mathcal{E}_i$. The important details are the energy gradients and approximate feature tangents.

The gradient of the energy $\mathcal{E}_i$ of particle $(\mathbf{x}_i, s_i)$ is

$$\nabla \mathcal{E}_i = \begin{bmatrix} \frac{d\mathcal{E}_i}{d\mathbf{x}_i} \\ \frac{d\mathcal{E}_i}{ds_i} \end{bmatrix} = (1-\alpha)\begin{bmatrix} \frac{dE_i}{d\mathbf{x}_i} \\ \frac{dE_i}{ds_i} \end{bmatrix} + \alpha \sum_{j=1}^{N} \begin{bmatrix} \frac{dE_{ij}}{d\mathbf{x}_i} \\ \frac{dE_{ij}}{ds_i} \end{bmatrix} \tag{24}$$

$$= (\alpha-1)\gamma \begin{bmatrix} \frac{\partial h(\mathbf{x},s)}{\partial \mathbf{x}}(\mathbf{x}_i, s_i) \\ \frac{\partial h(\mathbf{x},s)}{\partial s}(\mathbf{x}_i, s_i) \end{bmatrix} + \alpha \sum_{j=1}^{N} \begin{bmatrix} \frac{1}{\sigma_r}\frac{\mathbf{x}_i - \mathbf{x}_j}{|\mathbf{x}_i - \mathbf{x}_j|}\frac{\partial \Phi(r,s)}{\partial r}(r_{ij}, s_{ij}) \\ \frac{1}{\sigma_s}\frac{\partial \Phi(r,s)}{\partial s}(r_{ij}, s_{ij}) \end{bmatrix}.$$

The partial derivatives of Φ (either Φ_1 or Φ_2 of Section 3.4) are known analytically. The partial derivative of strength $h(\mathbf{x}, s)$ with respect to s is computed numerically, using central differences with step $s_{\text{max}}/1000$. For the scale-space crease feature tangent, in the spatial domain we use the approximate crease tangent $\mathbf{T}$ defined in Table 1, which has worked well even though correct tangents require third derivatives [16]; others have previously used the same approximation [3]. Along scale, we assume the feature tangent is $(\mathbf{0}, 1)$ and thus bypass computing $\partial h(\mathbf{x}, s)/\partial \mathbf{x}$. This implies that the crease is exactly stable with respect to scale, which is accurate enough when using particle-image energy and Φ_2 to localize maximal feature strength.

We evaluate the system for population control with a periodicity of PC iterations (currently PC = 10), traversing the particles once for deletion and once for addition. A particle is deleted if the total system energy would be lower without it, or if its strength falls below h_{min}. Particles are not deleted if more than half of their neighbors have just been deleted, to avoid overly aggressive culling of dense samplings. New particles are tentatively added next to existing particles that have a markedly asymmetric distribution of neighbors. Each tentative new particle undergoes PC iterations of updates in isolation, and the particle is added if the system energy is lower with it included. The negative potential well in $\phi_c(r)$ (Sect. 3.4) makes this work. If the new particle moves (during its isolated update) into the potential wells of existing particles, and thus adds to a good sampling distribution, the system energy will be lowered by including it.

When using Φ_2 to sample maximal strength creases, it is helpful to run the system until approximate convergence with $\gamma = 1.0$ and $\alpha = 0.0$, turning off all spatial interaction. Once particles are clustered around the scale of maximal strength, γ can be set with

$$\frac{1}{N}\sum_{1}^{N}\frac{\partial^2 E_i}{\partial s_i^2} = \frac{\partial^2 \Phi_2(r, s_{ij})}{\partial s_i^2} \Rightarrow -\gamma \frac{1}{N}\sum_{1}^{N}\frac{\partial^2 h(\mathbf{x}_i, s)}{\partial s^2}(s_i) = \frac{1}{\sigma_s^2} \tag{25}$$

$$\Rightarrow \gamma = \frac{-1}{\sigma_s^2 \frac{1}{N}\sum_{1}^{N}\frac{\partial^2 h(\mathbf{x}_i, s)}{\partial s^2}(s_i)}. \tag{26}$$

The idea in (25) is to equate the average *concavity* (second derivative) along scale of the particle-image energy E_i with the concavity of inter-particle energy E_{ij} so that the intuitive setting of $\alpha = 0.5$ actually balances the influence of the two energies. The second derivative of crease strength $\partial^2 h(\mathbf{x}_i, s)/\partial s^2$ in (26) is computed numerically.

While multiple parameters control the particle system, only a few need to be tuned depending on sampling goals. If using scale-attractive inter-particle energy Φ_2 to localize sampling along scale, the main parameters are: α to balance particle-image and inter-particle energies (Sect. 3.1), β to balance spatial repulsion and scale attraction (Sect. 3.4), and γ to modulate energy from feature strength (Sect. 3.3), though γ can be set automatically by (26) above. When using scale-repulsive inter-particle energy Φ_1 to sample the scale extent of features, we set $\alpha = 1$ (no particle-image energy), and β and γ are moot. The minimum feature strength h_{min} (Sect. 3.3) prevents particles from sampling insignificant features in noise. Interactively thresholding glyph visualizations of particle solutions on small test volumes can inform setting h_{min}, though future work may automate setting h_{min}. Finally, the particle radii in space σ_r and scale σ_s are set based on the desired sampling density relative to the expected characteristic size of features. The other system parameters have not been varied for our different applications, but are listed here for completeness: constraint satisfaction convergence threshold $\varepsilon_c = 0.001$ (Sect. 3.3), well position $w = 0.6$ and depth $d = -0.002$ of univariate potential $\phi_c(r)$, Butterworth order $b_{\text{ord}} = 20$ and cut-off $b_{\text{cut}} = 0.87$ (Sect. 3.4), and population control periodicity $PC = 10$ (above).

Glyph-based visualizations are important for inspecting scale-space particles. Visualizations can be used to debug the system implementation, observe the computation in progress to assess the appropriateness of the parameter settings, and, as we do in the following figures, to evaluate the spatial configuration of the solution and its relationship to cutting planes. While later work will compute geometric models from the particle solutions for further quantitative study, at this stage,

the visualizations offer evidence that the relevant anatomical structures are indeed sampled by the particles. Glyphs can be displayed at particle spatial locations $\{\mathbf{x}_i\}$ or, following previous illustrations of scale-space [68, 41], the scale axis can be mapped to a user-specified unit-length vector $\hat{\mathbf{s}}$, with the glyph for $(\mathbf{x}_i, s_i)$ placed at $\mathbf{x}_i - (\mathbf{x}_i \cdot \hat{\mathbf{s}})\hat{\mathbf{s}} + \rho s_i \hat{\mathbf{s}}$, with some user-specified scaling ρ. Points or other sprites would suffice for depicting the over-all system configuration, but we propose two types of glyphs, $\mathbf{D}_1$ and $\mathbf{D}_2$, which usefully reflect aspects of the eigensystem of the Hessian.

The first glyph $\mathbf{D}_1$ depicts Hessian eigenvectors and the relative magnitude of the eigenvalues. Existing tensor glyphs only express positive eigenvalues, so from the eigenvalues λ_i and corresponding eigenvectors $\mathbf{v}_i$ we find new positive eigenvalues μ_i for a tensor $\mathbf{D}_1$, which we visualize with superquadric glyphs [32].

$$\lambda' = (|\lambda_1| + |\lambda_2| + |\lambda_3|)/10, \quad \lambda_i'' = 1/\max(\lambda', |\lambda_i|), \quad (27)$$

$$\mu_i = \lambda_i''/\max(\lambda_1'', \lambda_2'', \lambda_2''), \qquad \mathbf{D}_1 = \sum_i \mu_i \mathbf{v}_i \otimes \mathbf{v}_i. \quad (28)$$

By using the reciprocal of the Hessian eigenvalues, the smallest magnitude eigenvalues λ_i, associated with the eigenvectors that approximate tangents to the crease feature, become the largest μ_i, so that the glyph visualizes the approximate local feature orientation. The glyph shapes along crease line are thus roughly cylindrical (aligned with the crease), and crease surfaces are indicated by disc-like glyphs, while the over-all size is controlled so the maximum μ_i eigenvalue is one.

The second glyph $\mathbf{D}_2$ explicitly indicates a particle's scale position s_i with glyph size along the eigenvectors *not* associated with the approximate crease tangent $\mathbf{T}$ (Sect. 3.3)

$$\mathbf{D}_2 = \mathbf{T} + s_i(\mathbf{I} - \mathbf{T}). \quad (29)$$

The base geometry is a cylinder with the axis of rotation aligned along either the principal or the minor eigenvector of $\mathbf{D}_2$, based on whether $\mathbf{D}_2$ is more linearly or more planarly anisotropic, according to measures c_l and c_p [67], as in Fig. 3 of [32]. Narrow cylinders, approximately tangent to crease lines at small scales, fatten into discs at larger scales; conversely, thin discs for crease surfaces extend into longer cylinders at larger scales.

The final component of our particle system visualization is the extraction of "connected" components. Even though there is as yet no explicit connectivity between particles, they can still be organized into equivalence classes by $(\mathbf{x}_i, s_i) \sim (\mathbf{x}_j, s_j)$ iff $E_{ij} \neq 0$. In practice, the inter-particle energy E_{ij} is non-zero for any two particles with overlapping support, even though the potential functions have zero-crossings. Particles are thus "connected" if they are near enough in space and scale to interact; disjoint features tend to be sampled by particles in different connected components. This is an imperfect means of isolating features because particles from different features can occasionally interact. We have found the connectivity heuristic, combined with sorting the components by number of particles, to help inspect the main image features sampled by the particle computation.

4 RESULTS

Figure 1 demonstrates scale-space particles on a synthetic variable-width torus (Fig. 1(a)), in which single-scale ridge sampling is clearly insufficient (Fig. 1(b), second row). Fig. 1(c) extends scale vertically, showing how the Φ_1 inter-particle energy (with zero particle-image energy) samples the complete feature manifold in scale-space, visualized with $\mathbf{D}_1$ glyphs color-mapped with crease strength. Fig. 1(d) uses $\alpha = 1$ (no inter-particle energy) to condense all the particles to the scale of maximal strength. Applying population control with Φ_2 and $\alpha = \beta = 0.5$, Fig. 1(f) shows the final result with $\mathbf{D}_2$ glyphs. Using linear instead of spline interpolation between scales, with the same number of pre-computed blurrings, significantly impedes accurate feature scale localization (Fig. 1(e)).

Fig. 4 illustrates non-scale-space properties of our particle system with a Möbius strip that emphasizes how crease surfaces are in general non-orientable [58]. The right side of Fig. 4 includes a portion of a cutting plane showing the coarse resolution of the underlying scalar
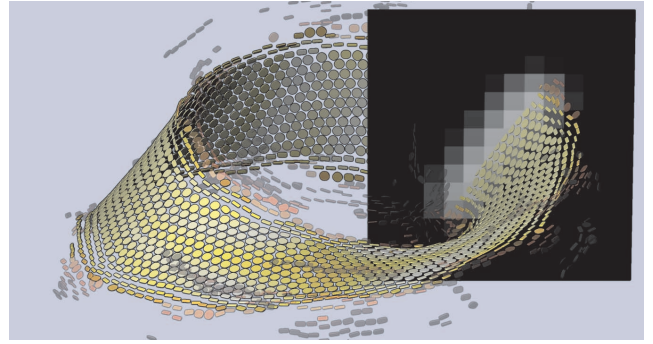

Fig. 4. Ridge surface sampling of a Möbius strip from a $40 \times 40 \times 20$ scalar volume demonstrates the shapes of $\mathbf{D}_1$ tensor glyphs and the need for thresholding feature strength (colormapped on glyphs).

data. The $\mathbf{D}_1$ glyphs show both the over-all shape of the Möbius strip and the shape of the Hessian at each particle. Through the middle of the Möbius strip, the $\mathbf{D}_1$ glyphs (28) indicate that the approximate tangent $\mathbf{T}$ (Table 1) seems to accurately reflect the actual surface tangent, which helps justify the constrained gradient ascent or descent approach of constraint satisfaction in (18). Away from the middle of the strip, the glyphs diverge most from a circular disc shape, and the ridge surface strength (indicated by the colormap) decreases. The errant, semi-transparent glyphs in Fig. 4 fell below a user-specified feature strength threshold. This illustrates how creases can be sampled even where poorly defined, and it emphasizes the need for strength thresholding to limit particles to meaningful features.

Valley lines indicating lung airways were sampled in the left lung of a 0.5mm³ resolution chest CT scan (a $230 \times 290 \times 480$ subvolume), with eight samples along scale $s \in [0\text{mm}, 8\text{mm}]$, using Φ_2 with $\alpha = 0.5$, $\beta = 0.5$, $\sigma_r = 1.2$mm, $\sigma_s = 2.5$mm. Figure 5(a) uses $\mathbf{D}_2$ glyphs color-mapped by strength to show the result of 6.4 hours of processing on 2.8 GHz Intel Core 2 Duo (single-threaded). With strength thresholding and connected component analysis, the wide range of scales in the branching structure is shown in Fig. 5(b), which includes for comparison a region-growing segmentation, shown as a semi-transparent surface. Note that the glyphs and the segmentation surface agree in their indication of both the airway location and radius over the full range of scales. Small airways are difficult to capture by region-growing methods due to limited resolution, even though they are especially valuable for quantitative study [28]. Figure 5(b) shows how particles tend to capture *greater* extents of the fine-scale airways, compared to the segmentation, while they also naturally capture even the widest airways. The cutting plane in 5(c) illustrates the underlying unsegmented data quality and resolution relative to the features successfully sampled within it.

The scale-repulsive inter-particle energy Φ_1 permits visualizing the complete scale-space extent of a feature. Figure 6(a) shows this for a single airway vessel segment, while colormapping onto round glyphs the valley line strength. This type of visualization informed our choice of inter-particle energy rather than constraint satisfaction for feature localization along scale, because the visualization showed the high spatial variability of feature strength. Figure 6(b) uses the scale-attractive inter-particle energy Φ_2 with low values of α and β (so that particle-image energy dominates inter-particle energy, which is mostly spatially repelling), to produce what is essentially a plot of the highly variable scales of maximal strength along the feature. Higher values of α and β in Fig. 6(d) (*i.e.* scale-attractive inter-particle energy dominates the particle-image energy) regularize the spatial variability of scale. There are noticeable but not dramatic improvements in the resulting sampled feature uniformity, shown in Fig. 6(c).

Figure 7 shows single-subject FA ridge surface analysis in a $[128 \times 128 \times 105]$ 1.5mm³ resolution DTI scan, sampled six times along $s \in [0\text{mm}, 5\text{mm}]$. Colormapping by scale illustrates how features near the cortical surface are better sampled at fine scales, while larger scales are better for structures deep inside the brain. The computation re-

(a) Full result
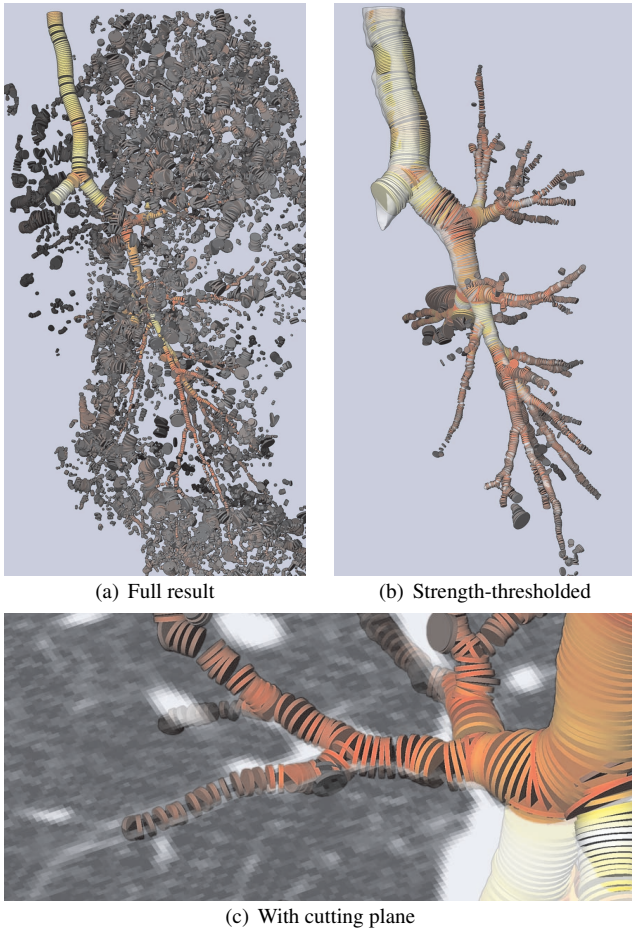
(b) Strength-thresholded

(c) With cutting plane

Fig. 5. Valley line analysis of lung CT to find airways. The complete result (a) is cleaned up with strength thresholding and connected-component analysis (b). Cutting plane (c) shows underlying unsegmented data relative to sampled features.

(a) Φ_1, mapping strength

(b) Φ_2; $\alpha = 0.1$, $\beta = 0.2$

(c) α, β comparison, mapping scale

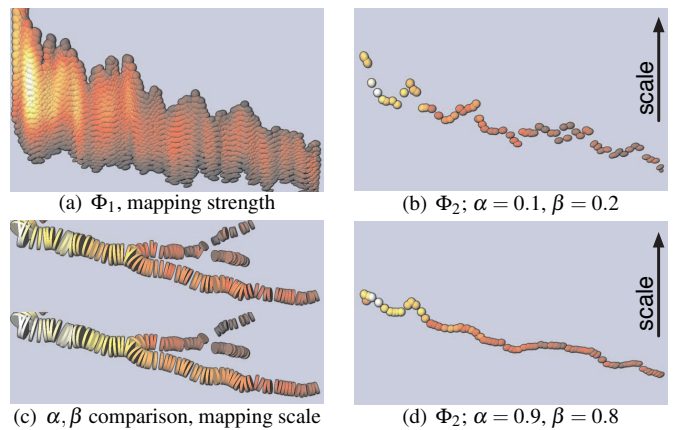(d) Φ_2; $\alpha = 0.9$, $\beta = 0.8$

Fig. 6. Scale-repelling inter-particle energy Φ_1 results in (a), colormapped by strength, provide context for evaluating two samplings (b) and (d) from different parameters for scale-attractive inter-particle energy Φ_2. 3-D visualization in (c), colormapped by scale, puts results from settings in (b) above those in (d).
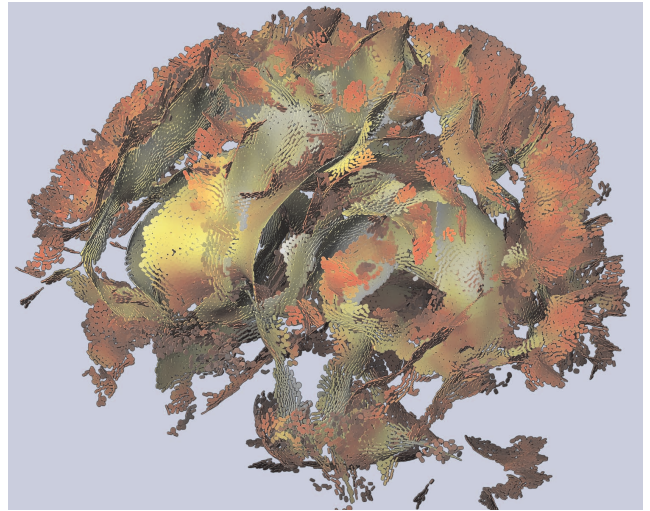

Fig. 7. Whole-brain scale-space FA ridge surfaces, colormapped by scale (brighter colors for higher scales).

quired roughly 8 hours, using Φ_2 with $\alpha = 0.7$, $\beta = 0.5$, $\sigma_r = 2$mm, $\sigma_s = 2.2$mm. Figure 8 demonstrates why scale-space analysis is useful for brain DTI: the corpus callosum (the thick connection between the hemispheres) is not reliably represented by an FA ridge surface sampled at a single scale, appropriate for structures closer to the cortical surface, while its interior shape is well sampled even at the thickest part by the scale-space particles.

Complementing FA ridge surfaces are the FA ridge lines in Fig. 9, shown with $\mathbf{D}_2$ glyphs that indicate the sampled lines as variable-width tubes. While we are still understanding which of these lines correspond to major white matter pathways, Fig. 9(b) shows that certain major tubular pathways, such as the Fornix and Cingulum Bundles, seem to be well captured. Given the generic nature of scale-space particles, we feel this represents a useful contribution to brain DTI analysis, given that other methods have been designed specifically to create models of just these pathways [51].

5   CONCLUSIONS AND DISCUSSION

Scale-space particles combine approaches from graphics, visualization, and vision to create a tool for robustly sampling and visualizing structure in real-world volume datasets. From computer vision, we leverage Lindeberg's discrete Gaussian [41] to create a memory-efficient scale-space reconstruction, and use scale-normalized strength measures [41] for crease feature [16] localization along scale [69, 35, 66]. The basic mechanics of particle-based feature sampling were derived from computer graphics [62, 70, 27] and scientific visualization [48, 46], and the whole system was debugged, monitored, and evaluated here with tensor glyph visualizations [32]. Our results sug-

gest that many possible biomedical applications of particle systems have yet to be explored, and that scientific visualization provides an effective framework for those explorations. Without being particularly adapted for either lung or brain analysis, our particles capture airways from CT as well as white matter structure from DTI, due to the particles' ability to gracefully handle features that span a continuous range of scales. The increasing resolution of medical scanners will decrease the lower limit on feature scale, creating more demand for better feature sampling and visualization methods that, like scale-space particles, work naturally over a wide range of scales.

There are many possible directions for future research. The refinement of particle systems for analyzing the scale-space structure of non-scalar data is especially exciting. It may be possible to reliably generate useful models of major white matter pathways solely from the differential structure of a single diffusion tensor invariant. This is counter-intuitive, given that we are not analyzing what might seem to be the most important aspect of the tensor, the principal eigenvector that is the basis of fiber tractography. On the other hand, we are implicitly taking advantage of significantly more tensor information by blurring the tensor field and then measuring anisotropy (rather than vice versa). Exploiting the non-linearity of anisotropy in this way mirrors earlier work on measuring tensor field *coherence* by blurring and then measuring anisotropy [67]. Future work can explore the non-linearity of anisotropy with respect to scale in more detail. The system

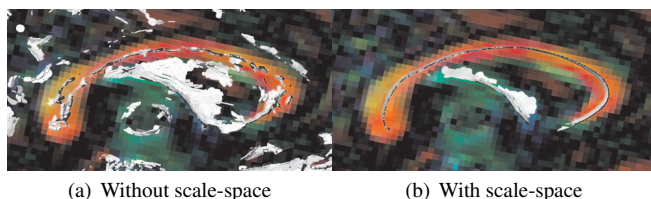(a) Without scale-space   (b) With scale-space

Fig. 8. The corpus callosum is a particularly large-scale feature that benefits from scale-space sampling (b) and is not reliably captured by a ridge at a single scale (a).

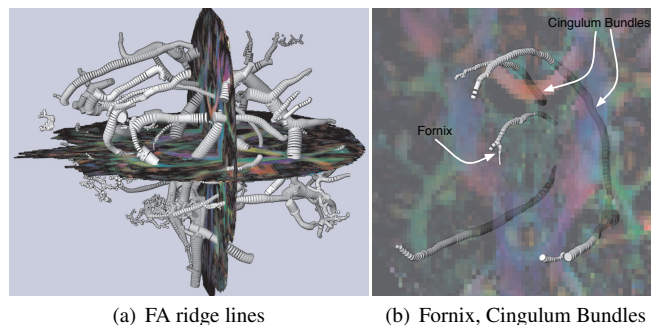(a) FA ridge lines   (b) Fornix, Cingulum Bundles

Fig. 9. FA ridge lines in scale-space (a) capture some white matter features known to be more tubular, such as the cingulum bundles (b).

described here may in fact be the most practical tool to help explore the broader question of the precise anatomic relationship between FA creases extracted in scale-space, and the shape of major pathways as computed from tractography.

Another direction for future work is the reduction of the number of particles by tuning the sampling density to local feature properties. Previous work [45] did this with a medial axis of the closed surface, but we see no obvious analog to medial axes for creases: crease lines are codimension-two, and crease surfaces can be non-orientable [58]. It is tempting to make the spatial particle radius vary in proportion to the scale position (larger particles at higher blurring levels), but it is unclear whether scale actually puts an upper bound on feature curvature, especially for creases of non-linear anisotropy measures.

Finally, there are many ways of optimizing our proof-of-concept implementation. Though we now err on the side of seeding too many particles, a secondary strength threshold (lower than the $h_{\min}$ in Sect. 3.3), evaluated at starting seed locations, could accelerate initialization and reduce the initial population. The current bottleneck in our system is actually in the later stage of computing particle motion, after population control has handled feature oversampling. Constraining particles to creases is expensive due to the many convolutions needed to reconstruct values and derivatives (Sect. 3.2). Simply copying the image values within the kernel support from the volume data into a convolution buffer accounts for a large part of the cost; it would be alleviated by more efficient use of the memory hierarchy through volume bricking [53]. Given the data-parallel nature of convolution and constraint satisfaction, multi-threading would also be beneficial, as would a GPU-based approach. Correctly parallelizing inter-particle interactions in the context of the global data structure used for spatial binning will require some care.

ACKNOWLEDGMENTS

REFERENCES

[1] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, August 2006.

[2] S. Aylward, E. Bullitt, S. Pizer, and D. Eberly. Intensity ridge and widths for tubular object segmentation and description. In *Proc. Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, pages 131–138, 1996.

[3] S. R. Aylward and E. Bullitt. Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction. *IEEE Trans. on Medical Imaging*, 21(2):61–75, Feb 2002.

[4] P. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysics Journal*, 66(1):259–267, 1994.

[5] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractograpy using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.

[6] D. Bauer and R. Peikert. Vortex tracking in scale-space. In *Proc. IEEE TVCG/EG Symposium on Visualization*, pages 140–147, 2002.

[7] C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - A technical review. *Nuclear Magnetic Resonance in Biomedicine*, 15:435–455, 2002.

[8] F. J. Bossen and P. S. Heckbert. A pliant method for anisotropic mesh generation. In *Proc. 5th Intl. Meshing Roundtable*, pages 63–74, 1996.

[9] S. Burgeth, S. Didas, L. Florack, and J. Weickert. A generic approach to diffusion filtering of matrix-fields. *Computing*, 81(2–3):179–197, Nov 2007.

[10] S. Busking, A. V. Bartroli, and J. J. van Wijk. Particle-based non-photorealistic volume visualization. *The Visual Computer*, 24:335–346, 2007.

[11] J. Cates, P. Fletcher, M. Styner, M. Shenton, and R. Whitaker. Shape modeling and analysis with entropy-based particle systems. In *Proc. IPMI 2007*, volume 4584 of *Lecture Notes in Computer Science*, page 333. Springer, 2007.

[12] J. Cates, M. Meyer, P. T. Fletcher, and R. Whitaker. Entropy-based particle systems for shape correspondence. In *Proc. Workshop on Mathematical Foundations of Computational Anatomy (MICCAI)*, pages 90–99, 2006.

[13] C. Chefd'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision*, 20(1–2):147–162, 2004.

[14] T. E. Conturo, N. F. Lori, T. S. Cull, E. Akbudak, A. Z. Snyder, J. S. Shimony, R. C. McKinstry, H. Burton, and M. E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proc. National Academy of Sciences*, 96:10422–10427, August 1999.

[15] P. Crossno and E. Angel. Isosurface extraction using particle systems. In *Proc. IEEE Visualization*, pages 495–498, 1997.

[16] D. Eberly. *Ridges in Image and Data Analysis*. Kluwer Academic Publishers, 1996.

[17] D. Eberly, R. Gardner, B. Morse, and S. Pizer. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*, 4(4):353–373, 1994.

[18] M. Fidrich. Iso-surface extraction in 4D with applications related to scale space. Technical Report 2833, Institut National de Recharche en Informatique et en Automatique (Sophia-Antipolis), March 1996.

[19] L. M. J. Florack and L. J. Astola. A multi-resolution framework for diffusion tensor images. In *Proc. CVPR Workshop on Tensors in Image Processing and Computer Vision*, pages 1–7. IEEE, June 2008.

[20] L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink, and M. A. Viergever. Scale and the differential structure of images. *Image and Vision Computing*, 10(6):376–388, 1992. Special issue: Information Processing in Medical Imaging 1991.

[21] L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink, and M. A. Viergever. Linear scale-space. *Journal of Mathematical Imaging and Vision*, 4(4):325–351, December 1994.

[22] A. F. Frangi, W. J. Niessen, R. M. Hoogeveen, T. van Walsum, and M. A. Viergever. Model-based quantitation of 3-D magnetic resonance angiographic images. *IEEE Trans. on Medical Imaging*, 18(10):946–956, 1999.

[23] J. D. Furst, S. M. Pizer, and D. H. Eberly. Marching cores: A method for extracting cores from 3D medical images. In *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 124–130, June 1996.

[24] J. M. Gauch and S. M. Pizer. Multiresolution analysis of ridges and val-

leys in grey-scale images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(6):635–646, June 1993.

[25] R. Gonzalez and R. Woods. *Digital Image Processing*. Addison-Wesley Publishing Company, Reading, MA, 2nd edition, 2002.

[26] C. Goodlett, B. Davis, R. Jean, J. H. Gilmore, and G. Gerig. Improved correspondence for DTI population studies via unbiased atlas building. In *Proceedings MICCAI 2006*, Lecture Notes in Computer Science 4190, pages 260–267, Copenhagen, Denmark, October 2006.

[27] P. Heckbert. Fast surface particle repulsion. In *ACM SIGGRAPH 1997 Course Notes (New Frontiers in Modeling and Texturing)*, pages 95–114, August 1997.

[28] J. C. Hogg, F. Chu, S. Utokaparch, R. Woods, W. M. Elliott, L. Buzatu, R. M. Cherniack, R. M. Rogersa, F. C. Sciurba, H. O. Coxson, and P. D. Par'e. The nature of small-airway obstruction in chronic obstructive pulmonary disease. *New England Journal of Medicine*, 350(26):2645–2653, June 2004.

[29] A. C. Jalba, M. H. F. Wilkinson, and J. B. T. M. Roerdink. CPM: A deformable model for shape recovery and segmentation based on charged particles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1320–1335, 2004.

[30] X. Jiao and P. Alexander. Parallel feature-preserving mesh smoothing. In *Proc. Int. Conf. Comput. Sci. Appl.*, pages 1180–1189, 2005.

[31] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour model. *International Journal of Computer Vision*, 1(4):321–331, January 1988.

[32] G. Kindlmann. Superquadric tensor glyphs. In *Proc. IEEE TVCG/EG Symposium on Visualization*, pages 147–154, May 2004.

[33] G. Kindlmann, X. Tricoche, and C.-F. Westin. Delineating white matter structure in diffusion tensor MRI with anisotropy creases. *Medical Image Analysis*, 11(5):492–502, October 2007.

[34] T. Klein and T. Ertl. Scale-space tracking of critical points in 3D vector fields. In *Proc. Workshop on Topology-Based Methods in Visualization*, pages 35–49, 2005.

[35] J. J. Koenderink and A. J. van Doorn. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.

[36] M. Kubicki, R. McCarley, C.-F. Westin, H.-J. Park, S. Maier, R. Kikinis, F. A. Jolesz, and M. E. Shenton. A review of diffusion tensor imaging studies in schizophrenia. *Journal of Psychiatric Research*, 41:15–30, 2007.

[37] D. LeBihan. Looking into the functional architecture of the brain with diffusion MRI. *Nature Reviews Neuroscience*, 4(6):469–480, June 2003.

[38] 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. *Journal of Mathematical Imaging and Vision*, 25(3):423–444, 2006.

[39] T. Lindeberg. Scale-space for discrete signals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(3):234–254, March 1990.

[40] T. Lindeberg. Scale-space theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics*, 21(2):224–270, 1994.

[41] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. In *Proceedings CVPR 1996*, pages 465–470, June 1996.

[42] L. Linsen, T. V. Long, P. Rosenthal, and S. Rosswog. Surface extraction from multi-field particle volume data using multi-dimensional cluster visualization. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1483–1490, 2008.

[43] A. M. López, D. LLoret, and J. Serrat. Creaseness measures for CT and MR image registration. In *Proceedings CVPR*, pages 694–699, June 1998.

[44] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics*, 21(4):163–169, 1987.

[45] M. Meyer, R. M. Kirby, and R. Whitaker. Topology, accuracy, and quality of isosurface meshes using dynamic particles. *IEEE Trans. on Visualization and Computer Graphics*, 12(5):1704–1711, Sept–Oct 2007.

[46] M. Meyer, B. Nelson, R. M. Kirby, and R. Whitaker. Particle systems for efficient and accurate high-order finite element visualization. *IEEE Trans. on Visualization and Computer Graphics*, 13(5):1015–1026, Sept–Oct 2007.

[47] M. Meyer, R. Whitaker, R. M. Kirby, C. Ledergerber, and H. Pfister. Particle-based sampling and meshing of surfaces in multimaterial volumes. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1539–1546, Nov–Dec 2008.

[48] M. D. Meyer, P. Georgel, and R. T. Whitaker. Robust particle systems for curvature dependent sampling of implicit surfaces. In *Proc. Shape Modeling and Applications (SMI)*, pages 124–133, June 2005.

[49] J. Miller and J. Furst. The maximal scale ridge: Incorporating scale into the ridge definition. In *Proc. Scale Space Theory in Computer Vision*, volume 1682 of *Lecture Notes in Computer Science*, pages 93–104. Springer, 1999.

[50] T. Möller, R. Machiraju, K. Mueller, and R. Yagel. Evaluation and design of filters using a Taylor series expansion. *IEEE Trans. on Visualization and Computer Graphics*, 3(2):184–199, 1997.

[51] L. O'Donnell, C. Westin, and A. Golby. Tract-based morphometry for white matter group analysis. *NeuroImage*, 45(3):832–844, Apr 2009.

[52] A. Pang and K. Smith. Spray rendering: Visualization using smart particles. In *Proc. IEEE Visualization*, pages 283–290, 1993.

[53] S. Parker, M. Parker, Y. Livnat, P.-P. Sloan, C. Hansen, and P. Shirley. Interactive ray tracing for volume visualization. *IEEE Trans. on Visualization and Computer Graphics*, 5(3):238–250, 1999.

[54] R. Peikert and M. Roth. The "parallel vectors" operator - a vector field visualization primitive. In *Proc. IEEE Visualization*, pages 263–270, 1999.

[55] P. Perona and J. Malik. Scale-space and edge detection using anisotropy diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[56] C. Pierpaoli and P. J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 33:893–906, 1996. Erratum in Magn. Reson. Med. 1997;37(6):972.

[57] A. Rösch, M. Ruhl, and D. Saupe. Interactive visualization of implicit surfaces with singularities. *Eurographics Computer Graphics Forum*, 16(5):295–306, 1997.

[58] T. Schultz, H. Theisel, and H. Seidel. Crease surfaces: From theory to extraction and application to diffusion tensor MRI. *IEEE Trans. on Visualization and Computer Graphics*, (to appear), 2009. doi:10.1109/TVCG.2009.44.

[59] I. Sluimer, A. Schilham, M. Prokop, and B. Van Ginneken. Computer analysis of computed tomography scans of the lung: A survey. *IEEE Trans. on Medical Imaging*, 25(4):385–405, 2006.

[60] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. J. Behrens. Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. *NeuroImage*, 31:1487–1505, 2006.

[61] W. Y. Su and J. C. Hart. A programmable particle system framework for shape modeling. In *Proc. Shape Modeling and Applications (SMI)*, pages 114–123, 2005.

[62] R. Szeliski and D. Tonnesen. Surface modeling with oriented particle systems. *Computer Graphics (Proc. ACM SIGGRAPH)*, 26(2):185–194, 1992.

[63] R. Szeliski, D. Tonnesen, and D. Terzopoulos. Modeling surfaces of arbitrary topology with dynamic particles. In *Proceedings CVPR*, pages 82–87, 1993.

[64] X. Tricoche, G. Kindlmann, and C.-F. Westin. Invariant crease lines for topological and structural analysis of tensor fields. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1627–1634, 2008.

[65] J. Weickert, C. Feddern, M. Welk, B. Burgeth, and T. Brox. PDEs for tensor image processing. In J. Weickert and H. Hagen, editors, *Visualization and Image Processing of Tensor Fields*. Springer, Berlin, 2005.

[66] J. Weickert, S. Ishikawa, and A. Imiya. Linear scale-space has first been proposed in Japan. *Journal of Mathematical Imaging and Vision*, 10(3):237–252, May 1999.

[67] C.-F. Westin, S. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108, 2002.

[68] A. Witkin. Scale-space filtering: A new approach to multi-scale description. In *Proc. IEEE Intl. Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 9, pages 150–153, March 1984.

[69] A. P. Witkin. Scale-space filtering. In *Proc. Intl. Joint Conference on Artificial Intelligence*, pages 1019–1021, Karlsruhe, Germany, 1983.

[70] A. P. Witkin and P. S. Heckbert. Using particles to sample and control implicit surfaces. *Computer Graphics (Proc. ACM SIGGRAPH)*, 28:269–277, 1994.

[71] C. Xu and J. L. Prince. Snakes, shapes, and gradient vector flow. *IEEE Trans. on Image Processing*, 7(3):359–369, 1998.

[72] S. Yamakawa and K. Shimada. High quality anisotropic tetrahedral mesh generation via packing ellipsoidal bubbles. In *Proc. 9th Intl. Meshing Roundtable*, pages 263–273, 2000.