DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507
Reprinted from
Journal of Imaging Science and Technology® 60(1): 010404-1–010404-11, 2016.

Visual Descriptors for Dense Tensor Fields in Computational Turbulent Combustion: A Case Study

G. Elisabeta Marai and Timothy Luciani
University of Illinois at Chicago, Department of Computer Science, Chicago, IL 60607
E-mail: gmarai@uic.edu

Adrian Maries
Learning Research and Development Center, Pittsburgh, PA 15260

S. Levent Yilmaz
Mathworks, Natick, MA 01760

Mehdi B. Nik
Stanford University, Center for Turbulence Research, Department of Mechanical Engineering, Stanford, CA 94305

Abstract. *Simulation and modeling of turbulent flow, and of turbulent reacting flow in particular, involve solving for and analyzing time-dependent and spatially dense tensor quantities, such as turbulent stress tensors. The interactive visual exploration of these tensor quantities can effectively steer the computational modeling of combustion systems. In this article, the authors analyze the challenges in dense symmetric-tensor visualization as applied to turbulent combustion calculation; most notable among these challenges are the dataset size and density. They analyze, together with domain experts, the feasibility of using several established tensor visualization techniques in this application domain. They further examine and propose visual descriptors for volume rendering of the data. Of these novel descriptors, one is a density-gradient descriptor which results in Schlieren-style images, and another one is a classification descriptor inspired by machine-learning techniques. The result is a hybrid visual analysis tool to be utilized in the debugging, benchmarking and verification of models and solutions in turbulent combustion. The authors demonstrate this analysis tool on two example configurations, report feedback from combustion researchers, and summarize the design lessons learned.*

INTRODUCTION

Computational simulation of turbulent combustion for gas turbine design has become increasingly important in the last two decades, due in part to environmental concerns and regulations on toxic emissions. Such modern gas turbine designs feature a variety of mixing fuel compositions and possible flow configurations,[1,2] which make non-computational simulations difficult. The focus of the computational research effort in this direction is on the development of computational tools for the modeling and prediction of turbulent combustion flows.

Received June 30, 2015; accepted for publication Nov. 4, 2015; published online Dec. 10, 2015. Associate Editor: Song Zhang.

Tensor quantities are common features in these turbulent combustion models. In particular, stress and strain tensors are often correlated to turbulent quantities—which appear unclosed in the mathematical formulation and thus need to be modeled as part of the computational simulation. Visual identification of the characteristics of such tensor quantities can bring significant insights into the computational modeling process.

However, these computational tensor fields are very large and spatially dense—a good example of the Big Data revolution across sciences and engineering. Figure 1 shows an example turbulent combustion configuration, featuring a grid size of 10^6 and 6×10^6 particles (shown as spheres); this dataset should be considered in contrast to traditional tensor datasets, which feature grid sizes in the 10^2 range. At such large scales, typical glyph encodings become cluttered and illegible. Furthermore, combustion experts seldom have an intuitive understanding of the tensor quantities. In this respect, from a tensor visualization perspective, working with these datasets poses an array of challenges. Are traditional tensor and flow representations useful in this context? Does increasing the level of complexity or expressiveness of such representations help or hinder? Is interaction speed more important than the benefits gained from complex descriptors? In this article, we address a specific application design problem. In the process of exploring the design space, we also investigate some of the larger visualization questions above, through the opportunity of a case study in the computational-combustion domain.

In this work, motivated by an ongoing collaboration with domain experts,[3] we investigate the challenges associated with the exploratory visualization of tensor quantities in turbulent combustion simulations. We first provide a characterization of the problem domain, including a data analysis. Through a case study involving five senior combustion researchers, we then iteratively explore the space of tensor visual encodings. We implement and evaluate several

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

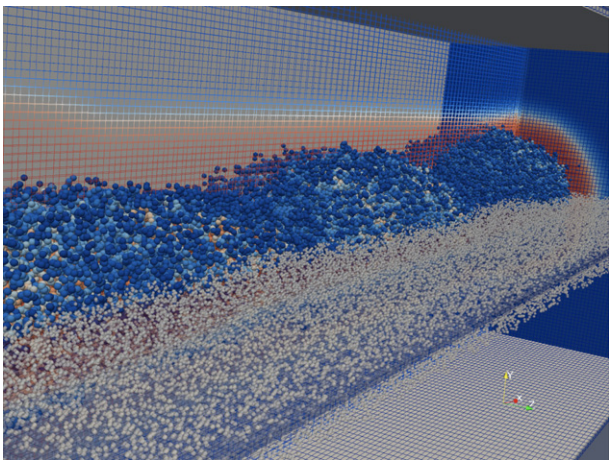

Figure 1. One timestep in an example turbulent combustion configuration. The grid size is 10^6. In this image, 6×10^6 particles are shown as spheres. This dataset should be contrasted to traditional tensor datasets, which feature sparse grids in the 10^2 range. At this scale, typical glyph encodings become cluttered.

approaches advocated by the visualization community in an interactive prototype, and we contrast these approaches with the best-of-breed visualization practices in the target domain. Based on domain expert feedback, we then focus our efforts on identifying effective visual descriptors for volume rendering of the combustion tensor data. Our contributions include a novel density-gradient descriptor and the adaptation of a machine-learning classification technique. Next, we evaluate the visual descriptors on two computational-combustion datasets of particular interest, and we show the importance of the proposed approach for debugging the numerical simulation of complex configurations. In an effort to better bridge the gap between the combustion and tensor visualization communities, we describe these tensor field datasets. Last but not least, we contribute a summary of design lessons learned from the study and from the application design process.

To the best of our knowledge, this is the first formal, exploratory case study of tensor visualization techniques in the context of very large, high-density turbulent combustion flow.

TENSORS IN TURBULENT COMBUSTION MODELING

Turbulent Combustion Modeling. A sufficiently accurate, flexible and reliable model can be used for an *in silico* combustor rig test as a much cheaper alternative to the real-life rig tests employed in combustor design and optimization. In order to achieve such a model, the methodology should be well tested and proven with lab-scale configurations.

Multiple numerical approaches exist for the generation of such computational models of combustion, most notably Direct numerical simulation (DNS), Reynolds-averaged Navier–Stokes (RANS) and Large eddy simulation (LES). DNS, RANS and LES have complementary strengths. However, all models begin by describing the compressible reacting flow via a set of partial differential equations (PDEs) that represent the conservation of mass, momentum and energy. These PDEs are a fully coupled set of multi-dimensional non-linear equations and can be posed in a variety of forms depending on the flow conditions (compressibility, scale, flow regime, etc.).[4] In this article, we exemplify the visualization of stress/strain tensors, and therefore restrict the presentation to the pertinent subset of these PDEs, namely the momentum transport equation.

Stress, Strain and Turbulent Stress Tensors. A tensor is an extension of the concept of a scalar and a vector to higher orders. For example, while a stress *vector* is the force acting on a given unit surface, a stress *tensor* is defined as the components of stress vectors acting on each coordinate surface; thus, stress can be described by a symmetric second-order tensor (a matrix).

The velocity stress and strain tensor fields are manifested in the transport of fluid momentum, which is a vector quantity governed by the following conservation equation:

$$\frac{\partial \rho u_i}{\partial t} + \frac{\partial \rho u_i u_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ij}}{\partial x_j} \quad \text{for } i = 1, 2, 3, \qquad (1)$$

where the Cartesian index notation is employed, in which the index $i = 1, 2, 3$ represents spatial directions along the x, y and z Cartesian coordinates, respectively, and the repeated index j implies summation over the coordinates. Here, t is time, ρ is the fluid density, $\mathbf{u} \equiv [u_1, u_2, u_3]$ is the Eulerian fluid velocity, p is the pressure and τ is the stress tensor defined as

$$\tau_{ij} = \mu S_{ij}, \qquad (2)$$

where μ is the dynamic viscosity coefficient (a fluid-dependent parameter) and S is the velocity strain tensor defined as

$$S_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right). \qquad (3)$$

For an intuitive understanding of the *fluid* stress and strain tensors (or, rather, the rate-of-strain tensor), let us consider, in contrast to combustion fluids, an elastic *solid* such as a strip of rubber. If we twist the rubber around and keep it twisted, a certain amount of deformation occurs. While this deformation does not have an intuitive equivalent in fluids, gases may dilate, rotate and shear during combustion. The rate at which these deformations occur is manifested in the velocity strain (a.k.a. rate-of-strain) tensor S defined in Eq. (3).

Similarly, the stress tensor is the force that is required to sustain a given strain rate. Considering again the solid rubber example, one can apply the same deformation on a tough rubber versus a more elastic rubber. However, the tough rubber requires more force to get the same strain. For the fluid flow motion, there are different constitutive models that relate strain to stress. In combustion, the experts examine Newtonian fluids, a special family of fluids where the stress is a linear function of the rate of strain, as spelled out in Eq. (2).

From the computational modeling perspective, the third tensor quantity of interest is the turbulent stress tensor. The

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

turbulent stress tensor T is defined as

$$T_{ij} = \langle\rho\rangle(\langle u_i u_j\rangle_L - \langle u_i\rangle_L\langle u_j\rangle_L). \quad (4)$$

We note from Eq. (4) that the role and meaning of the *turbulent stress tensor* are very different from the stress tensor. In fact, the definition above does not relate in any way to stress. However, many combustion modeling techniques rely on a "turbulent viscosity hypothesis," where T is modeled as if it is acting like a stress tensor:

$$T = \mu_t S, \quad (5)$$

where μ_t is the turbulent viscosity coefficient. Indeed, the whole literature on modeling of turbulent flow is providing closure models for T in one way or another. For these reasons, the velocity strain tensor and its derived stress tensor and turbulent stress tensor are of extreme interest to computational modelers. We note that the rate-of-strain tensor provides an immediate proxy—modulo a constant coefficient—for both the stress tensor and the turbulent stress tensor.

RELATED WORK

Several methods have been proposed for the visualization of tensor datasets. They include eigenvector colormaps, glyphs, streamlines, volume rendering and volume deformation.

Eigenvector colormaps[5] are commonly used by mechanical engineers for component-by-component visualization of vector and tensor quantities. The limitation of colormaps is that they can only display one type of information at any one time, e.g., tensor component or eigenvalue. Another approach utilizes glyphs.[5] In engineering fields, glyphs tend to be domain-based, highly specific abstractions. For example, stress hedgehogs[6] and Mohr circles[7] have been used to visualize mechanical stress in civil engineering, but are unfamiliar to combustion researchers. Streamlines (sometimes called hyper-streamlines) and streamtubes have been used to visually represent mechanical stress.[6,8] Like glyphs, streamlines are also prone to clutter and occlusions. Volume rendering of stress magnitudes,[9,10] combined with tracing of short line segments to show stress direction, has been used both in orthopedic planning[11] and to indicate the effect of stress on objects in engineering.[12] The volume-rendering approach holds promise for dense datasets, although it can primarily show scalar information; many of the existing mappings are not applicable directly to visualizing combustion flow.

Because of the complexity of a 3D tensor field—each tensor potentially has six components (three diagonal and three non-diagonal elements), three eigenvalues and three eigenvectors—tensor visualization techniques often suffer from occlusion and cluttering problems. To address these problems, researchers often use interaction and, more recently, dimension reduction via classification.[11,13–15]

METHODS

Domain Characterization

In this section we present the first contribution of this design study, an analysis of the problem domain. We collected the information below, and we also performed the followup design and evaluation through a collaboration with five senior combustion modeling researchers. The research experience levels in this group ranged from 32 years to 8 years. Four of these domain experts are well versed in the best-of-breed visualization practices in their larger field (CFD—computational fluid dynamics); two of these experts are co-authors of this article. One of the five domain experts also has several years of experience in developing commercial visualization software for modeling and simulation applications.

Data Analysis. Turbulent combustion modeling datasets are the result of computational simulations often performed at high resolution in both space and time. Even with the use of current state-of-the-art supercomputers and algorithms, the simulations can get quite demanding or even unfeasible computationally. For example, LES simulations of a realistic laboratory size flame require resolutions of the order of 10^7 grid points. Post-processing data are output at around 200 GB per timestep. A typical run using 1000 processors takes 1 month of continuous runtime. DNS, being a model-free representation of the basic equations, suffers from high computational demand. LES and RANS can address this shortcoming, however, *at the cost of modeling errors* and loss of generality. LES and RANS models allow the researcher to tackle ever larger and more realistic combustion problems, so the size of the datasets can still get prohibitively large; conventional visualization practices do not commonly apply. Visualization, however, is quite essential for the combustion community, as it can bring to light the inaccuracies in the numerical approach and can also allow researchers to analyze the validity of the models employed.

The questions posed during analysis have, fundamentally, an exploratory nature: the experts seek to determine whether a particular tensor-based modeling approach introduces any artifacts in the simulation of a complex configuration. Unlike in the case of scalar and vector fields, the experts are, however, seldom able to define beforehand the nature of these tensor artifacts.

Case Study Methodology. The methodology used for this case study included informal interviews, direct observation, participation in the life of the experts' groups, collective discussions, analyses of documents and results produced within the groups, self-analysis and results from evaluation activities.

As is typical in scientific visualization, the application is primarily intended for use by a small number of skilled professionals, and is best evaluated through an exploratory, case study approach. In contrast to this approach, evaluation via formal user studies can bias the routine statistical analysis toward specific tasks, solutions and approaches—due to the lack of necessary or sufficient domain knowledge on the side of non-expert users and designers.

This report is the first formal exploratory case study on tensor visualization techniques for very large high-density turbulent combustion flows, in contrast to many previous

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

user studies conducted on visualization of scalar and vector fields. While scalar and vector field visualizations benefit from clearly defined features and tasks, tensor field analysis has largely an exploratory nature.

Challenges. While many tensor visualization techniques have been proposed and implemented in various systems, in particular in medical imaging and civil engineering, the visualization of turbulent combustion tensors is challenging on multiple fronts.

The first challenge we identify is the one of *effective visual abstractions*. By *effective* we denote visual abstractions which capture the physical or mathematical aspects of the tensors and which are intuitive to the application-domain practitioners. The visual abstraction issue is particularly difficult because the physical meaning of engineering tensors is not necessarily intuitive. Symmetric second-order tensors are used routinely as abstract quantities in the mathematical modeling of turbulent combustion, and they are considered very useful for computation. Yet, only abstractions of the tensor, such as the trace of the tensor matrix, may bear physical meaning to the domain practitioner.

The physical meaning of tensors can further greatly impact how they should be analyzed and visualized, even when the mathematical representations of these tensors are the same across domains. Examples of this include the stress tensor and strain tensor from solid mechanics, the rate-of-deformation tensor from fluid dynamics, and the diffusion tensor from medical imaging, all of which are second-order, symmetric tensors. Yet, mathematical analyses and visualizations need to be tailored to best suit the domain scientists' needs.

A second major challenge is the one of *scalability*. Because combustion datasets are the result of computational simulations often performed at high resolutions, they tend to be large scale and particularly dense. Such high densities and big volumes lead naturally to scalability, clutter and occlusion problems when visualizing the data, as well as to slow interaction when visually exploring the data.

Third, because researchers are particularly interested in the 3D structure of the flow and the possible ways to numerically decompose and simulate this flow, it is important to support *local details in the 3D global context* of the data.

Iterative Design and Evaluation

Our second contribution is an iterative exploration and evaluation of the visual design space relating the visual encodings and interaction techniques to the requirements of the tensor target domain. In this section, we describe a set of visual abstractions which aim to capture the physical or mathematical aspects of the tensors, and then focus on designing a set of specific volume-rendering descriptors.

In this study, we adopted a parallel-prototyping approach. In this approach, multiple designs are created and presented for feedback in parallel; the process has been shown to lead to improved cross-pollination of ideas and less similar designs compared with serial prototyping.[16] Following a feasibility analysis of related visualization approaches (in the third section), an analysis of the perceived visualization state of the art in the target domain, and several group discussions in the tensor visualization community,[17] we selected the following parameter axes for exploring the tensor visual descriptor design space: glyphs, streamlines and volume-rendering descriptors. Along each axis, the level of complexity of each descriptor was varied: the glyph type spanned linear, Westin's and superquadric representations; the placement, seeding and (optional) rendering style of glyphs and streamlines were varied (grid based or flow based; plain or illuminated streamlines); and visual descriptors of increasing complexity were used for volume rendering. Interactions were considered in conjunction with these visual descriptors and with the tasks outlined in the earlier section; we support manipulation, zooming and filtering operations. The application tool was implemented in C/C++ with OpenGL and CUDA for rendering and QT for the user interface.

Visual Encodings

Glyph and Streamline Descriptors. Representations that combine the different tensor components into a single image were of immediate interest to our collaborators in combustion research. In particular, previous analyses of the smallest dataset routinely used component-by-component colormap representations (ParaView) to identify discontinuities in the tensor field. We speculated that a combined representation of the tensor components as glyphs would prove useful for this task. We experimented with Westin's composite glyphs,[18] superquadric glyphs[19] and a highly simplified linear encoding. The stress tensor is symmetric, and its eigenvalues are real and positive.

Westin's glyph consists of a rod, a disc and a sphere. The eigenvectors and eigenvalues of the velocity strain tensor were calculated for each evenly spaced grid point, then mapped to the composite glyph. In the degenerate case, when the difference between the first two eigenvalues was less than 0.001 (heuristically determined), the corresponding glyph was displayed as a gray sphere.

Alternatively, to better convey shape and orientation, the tensor information was encoded using superquadric glyphs. The use of just six vertexes per quadric in the Sandia-D dataset results in about ten million vertexes to compute, store and view in real time; in reality, each quadric needs at least 20 vertexes to distinguish its shape. The superquadric glyphs can optionally be colormapped.

Finally, following user feedback, we explored a third, highly simplified, linelet encoding of the orientation and magnitude of the main eigenvectors only.

As anticipated, mapping glyphs to the 3D tensor field led to clutter and occlusions, even when the glyph field was subsampled by a factor of 25. The problem was not significantly alleviated by simplified glyph encodings. While the simplified, linear result was easier to interpret, in particular in 2D cross-sections of the field, 3D views of the representation were still illegible due to clutter, and failed

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

to deliver a sense of the 3D flow. Interestingly, the grid placement of the glyphs (as opposed to jitter) was preferred by the domain experts, because of the direct relationship of this grid to the sampling grid used to compute and simulate the fields.

Given the specific value range of the tensor fields, Runge–Kutta 4 integral paths through the main eigenvector of the stress tensor field yield only linelets of short length—at most two cells. Runge–Kutta 4 integral paths through the velocity field inspired, however, a glyph placement scheme aimed at alleviating the clutter and occlusion problems. In this scheme, the stress tensor glyphs were placed at regular intervals along the streamline paths instead of on a grid. To improve the legibility of the streamlines, we also implemented an illuminated[20] version.

Volume-Rendering Descriptors. In the target domain, volume rendering is typically used for scalar quantities like species mass concentration (combustion) or smoke concentration (CFD). It is rarely used for vector quantities like velocities or for tensor quantities. The challenge in volume rendering of tensor quantities lies in selecting the scalar variable to map to the volume rendering.

In this section we derive and propose specific visual descriptors for volume rendering of combustion tensor data, in the following order: velocity-gradient-based descriptors, a density-gradient descriptor and a classification-based descriptor obtained by adapting a machine-learning clustering technique.

Velocity-gradient-based descriptors. In the experts' opinion, volume rendering gives an idea of the extent or spread of a specific variable. Therefore, in principle, any tensor-based variable that not only has a physical meaning for tensor modeling, but also has a tendency to spread or diffuse out, is a good candidate for volume rendering. However, the challenge lies in identifying intuitive descriptors for the target domain.

Analysis of the definitions and derivations of the various tensor quantities shows that the velocity strain tensor contains a good amount of physical information. The various components of this tensor have a straightforward meaning; in particular, the deformation D can be decomposed through standard matrix decomposition into a symmetric and an anti-symmetric component as

$$D = 1/3\nabla \boldsymbol{I} + S + \Omega. \tag{6}$$

In this decomposition, ∇ is the dilatation term div u (divergence of u), which is a scalar, and $\boldsymbol{I}$ is the identity matrix; div u is zero for incompressible fluids (e.g., water at room temperature), but not in our case. Here, S is the symmetric-deviatoric component:

$$S = 1/2(\partial u_i/\partial x_j + \partial u_j/\partial x_i) - 1/3\nabla \boldsymbol{I} \tag{7}$$

and Ω is the anti-symmetric component:

$$\Omega = 1/2(\partial u_i/\partial x_j - \partial u_j/\partial x_i). \tag{8}$$

Dilatation appears here for compressible flows only.

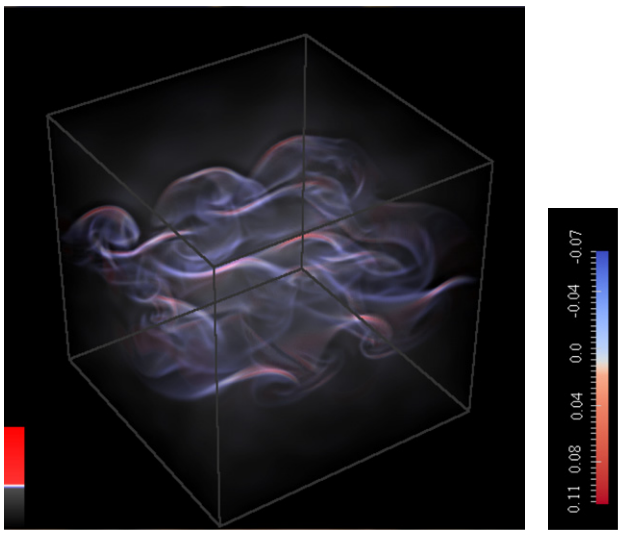

Figure 2. Volume rendering of *divergence* of the temporal mixing layer dataset, and color transfer function (right). Regions with saturated blue or red indicate higher magnitudes of divergence, either positive or negative. The rendering shows very clearly the 3D nature of the tensor field, which is difficult to extract through other visualization methods.

The respective physical interpretations for these three terms are, roughly, the dilatation (negative of compression), the shear rate and the rate of the rotation of the fluid. The velocity-gradient quantities we turned to for the volume rendering are thus the dilatation term magnitude, also known as the *divergence*, which can be calculated as the trace of the strain tensor in Eq. (3), the shear rate and the rotation of the fluid. Figure 2 shows an example image of divergence for a temporal mixing layer flow.

Density-gradient-based descriptors. We derive a second class of volume-rendering descriptors from the density gradient of combustion datasets. The density gradient relates to the stress tensor through the conservation equation Eq. (1). These descriptors can be used to generate flow visualizations in the style of Schlieren images.[21]

Schlieren (from German, meaning "streaks"), first observed by Robert Hooke in 1665, are optical inhomogeneities in transparent materials which are otherwise not visible to the human eye. The Schlieren flow photography process uses a viewing screen and a knife edge to generate a shadow pattern; this shadow pattern is a light-intensity representation of the low- and high-density regions in the flow.

Numerical Schlieren pictures can also be generated, like experimental Schlieren images, from the density-gradient field. Following Hadjadj and Kudryavtsev,[21] we derive the quantity

$$\text{Schlieren}(x, y, z) = \beta * \exp\left(-\frac{k * |\nabla \rho|}{|\nabla \rho|_{\max}}\right), \tag{9}$$

where x, y, z are the 3D coordinates at which the quantity is evaluated, max denotes the maximum values of the density gradient over the entire field, and β and k are rendering parameters. The parameter β determines the shade of

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

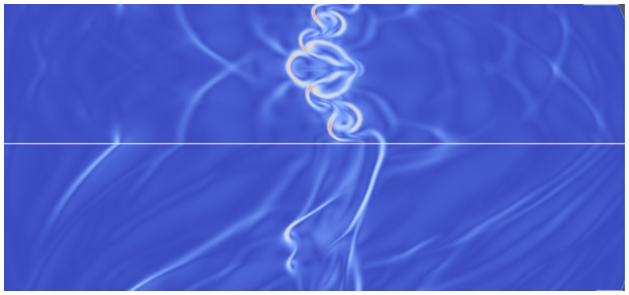

Figure 3. One timestep of a shocklet dataset, rendered using a Schlieren density-gradient-based descriptor. For this image, $\beta = 0.8$ and $k = 20$, which are experimentally determined. Note that the density-gradient descriptor captures well the shocklet boundary.

the gray color that corresponds to the zero gradient; its typical value is 0.8. The parameter k serves to amplify small gradients; in our experiments, a value of k of either 15 or 20 provided good results. Other density-gradient descriptors are also possible; variations of Eq. (9) can generate additional effects, such as numerical interferograms and shadowgraphs.[21]

Figure 3 shows the Schlieren for one timestep of a shocklet dataset—a larger, 12,900-timestep variation of the temporal mixing layer flow shown in Fig. 2. The density-gradient-based descriptor appears to capture the shocklet boundary more strongly than the velocity-based descriptors.

Classification-based descriptor. The last—and most computationally complex—visual descriptor for volume rendering we pursue attempts to account for the tensor structure as a whole. To this end, we use clustering analysis—a machine-learning technique—to group together tensor data points that are similar to one another. In this approach, clustering is performed on the six distinct values of the strain tensor. Ideally, similar tensor field points are grouped together into clusters; the optimum number of clusters is determined through repeated classification and analysis of the cluster separation. To enable application of clustering techniques at this scale, we perform a pre-clustering step[22] via subsampling of the dataset. This step enables us to obtain starting cluster centers. The pre-clustering step is followed by a K-means algorithm with L-2 norm to partition and cluster the tensor field components.

For volume rendering, each cluster is assigned an index value, and the transfer function is set monochromatically according to each cluster. We keep track of each cluster's mean signature in order to enable cluster stabilization across the visualization of multiple timesteps.

An advantage of the approach is that it effectively segments large combustion datasets into regions of interest, thus making the effective visualization of very large datasets possible. Feature extraction also helps users to highlight and focus on regions of interest. A disadvantage of the approach is, obviously, that similarity is a relative measure: selecting too small a number of clusters for classification may miss modeling artifacts, while too large a number of clusters may lead to illegible renderings.

In our experiments, clustering was performed on a quad-core 3.33 GHz Intel i5 CPU machine with 16 GB of RAM. On average, 8M grid points took between 15 and 20 min to generate three or four clusters.

We note that, as with velocity-gradient and density-gradient descriptors, multiple other classification descriptors are possible.

Interaction Techniques

Tufte's principles[23] suggest that the information content of a visualization would be maximized by a hybrid visual encoding. To prevent the sheer volume of combined information from becoming overwhelming, we follow a strategy based on Shneiderman's info-vis mantra:[24] volume renderings are used as an overview of the flow and serve as a visual anchor, while the streamlines generated interactively enable filtering of interesting regions, and a zoom lens allows glyph representations to function as details on demand.

The hybrid prototype uses two modes of operation, *explore* and *filter*. In the explore mode, a user can manipulate the scene and zoom in and out. To speed up interaction during rotation and zooming, fewer rays and a reduced sampling rate along the ray are used compared with the full-resolution rendering. While ParaView employs a muted neutral gray background, we found that a dark background competed less with the volume-rendering overlay information.[25] In the filter mode, streamlines can be generated interactively, highlighted and compared. Streamline seed points can be dragged to new locations in the volume. To further reduce clutter, the glyph representations can be mapped either to axis-oriented cutting planes, which are also controlled by the user, or along velocity streamlines. A zoom lens was also implemented to facilitate the analysis of glyph-based representations. The magnification tool uses the stencil buffer to create a viewport into the glyphs. In both modes, a user can filter out streamlines, volumes or the glyph representations.

APPLICATION TO COMBUSTION DATA AND DISCUSSION

We have employed the results of two simulations in this study. The first and most intriguing dataset, the Sandia-D experiment, is the result of an LES simulation of a turbulent jet configuration. The second dataset is a canonical test problem employed in turbulent reacting flow research, namely a temporal mixing layer configuration. The feedback below was provided by the research groups involved in our study.

Sandia-D

The Sandia-D dataset is a centimeter-scale jet configuration with a fuel jet at the center (methane–air mixture for this dataset) surrounded coaxially by a slower-speed hot pilot flame. The pilot flame is further surrounded by a co-flowing

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

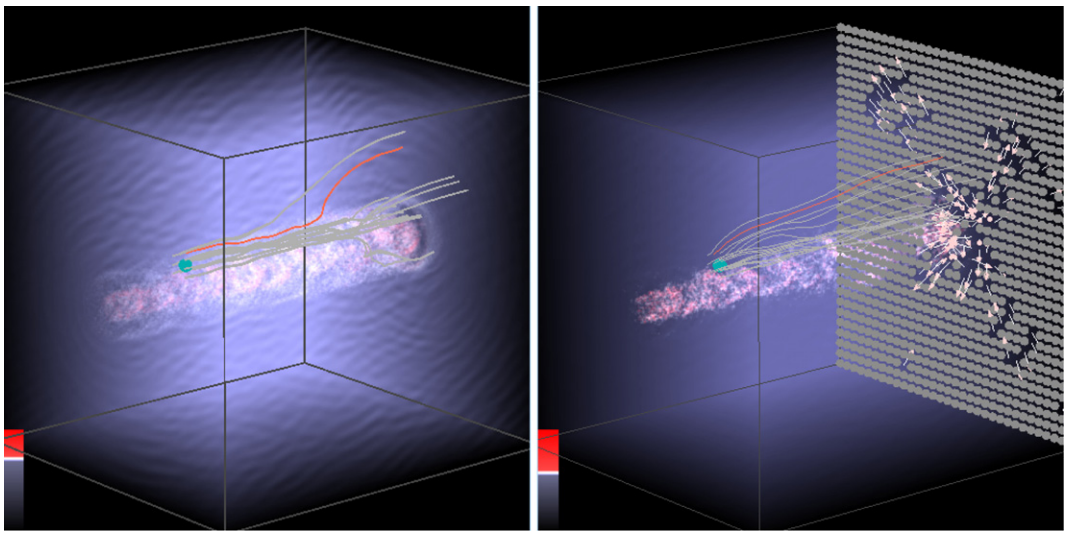

Figure 4. Velocity-gradient-based descriptor (*divergence*) of the Sandia-D dataset. Note the ripple effect in the left image, indicating a numerical artifact in the simulation; the ripple disappears in later timestamps (right).

hot air stream. It has a Reynolds number of 22,400, and so its DNS is computationally unreasonable. The LES simulation provides solutions of the turbulent fields of species as well as the velocity field as a function of 3D space and time. The data used in this experiment comprise the filtered turbulent velocity vector and the turbulent stress tensor fields taken at a snapshot in time and discretized over a uniform Cartesian grid of size 200 in the streamwise direction and 160 in each of the cross-stream directions (≈5M grid points).

The analysis of the Sandia-D dataset showcases the advantages of the volume-rendering descriptors as a means of debugging numerical simulations, along the very artifact identification and analysis tasks outlined earlier. Using the hybrid prototype, the researchers were surprised to notice a rippling artifact surrounding the central jet (Figure 4 left). This artifact had escaped previous numerical and visual analyses. The snapshots shown are for two time levels at 0.7 and 1.8 residence times, respectively, where one residence time is equal to the total streamwise length divided by the mean jet velocity. The rippling effect disappears as the simulation goes further in time. This non-physical ripple effect could be attributed to numerical artifacts of the employed discretization scheme in the LES simulation, and was pronounced only in the incompressible regions. This in-depth analysis was enabled by the hybrid prototype exploration. According to the domain specialists, the hybrid descriptors made this analysis especially clear in the right snapshot, where the effects of high-frequency pressure waves are less compared with earlier in the simulation.

The rippling artifact was clearly and immediately visible in the divergence descriptor, but not noticeable in either the shear or rotation descriptors, in the density-based descriptor, or in the classification descriptor, regardless of the transfer function used. While the density and classification descriptors missed the artifact, they were praised for their accurate detection of boundaries, and the latter also for its compression potential. Collectively, the descriptors were considered to be outstandingly useful in an analysis of the regions at and near the jet core, where they were indicative of the high spatial and temporal gradient in the gas-mixture density.

The glyph and streamline components (Figure 5) were considered to be useful, although not on the same scale as the volume-rendering descriptors. While the tensor itself is, in the experts' verbiage, "pretty complete, and very useful for computation," its encoding as a glyph was "not so easy to understand." Increasing levels of complexity did not help. In particular, shape changes in the superquadric encoding were hard to read at the scale of the dataset, despite the use of filtering and of the magnifying lens (sometimes at a 10,000:1 zoom level). The glyph placement on streamlines further led to "cool images," but no further insight.

The interactivity of streamlines (used as probes in the hybrid prototype) was much appreciated, and sometimes praised to excess when compared with state-of-the-art toolkits, in particular ParaView. Illuminated streamlines were considered superior; however, the advantage did not raise the same level of excitement as the interaction aspect.

Temporal Mixing Layer

The temporal mixing layer is a simple configuration where two streams of fuel and oxidizer flow over and against each other. The flow speeds are adjusted for a low Reynolds number, yielding a narrow range of length scales, and this configuration can be easily tackled with DNS and then used as a benchmark. The data for the temporal mixing layer are similarly at a snapshot in time and at the full DNS resolution over a grid of size 193 grid points in two Cartesian directions and 194 in the other (≈8M grid points).

For the mixing layer configuration, observations were immediately visible where the "mushroom" pattern around the shear layer at the mid-zone was distinguished well

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

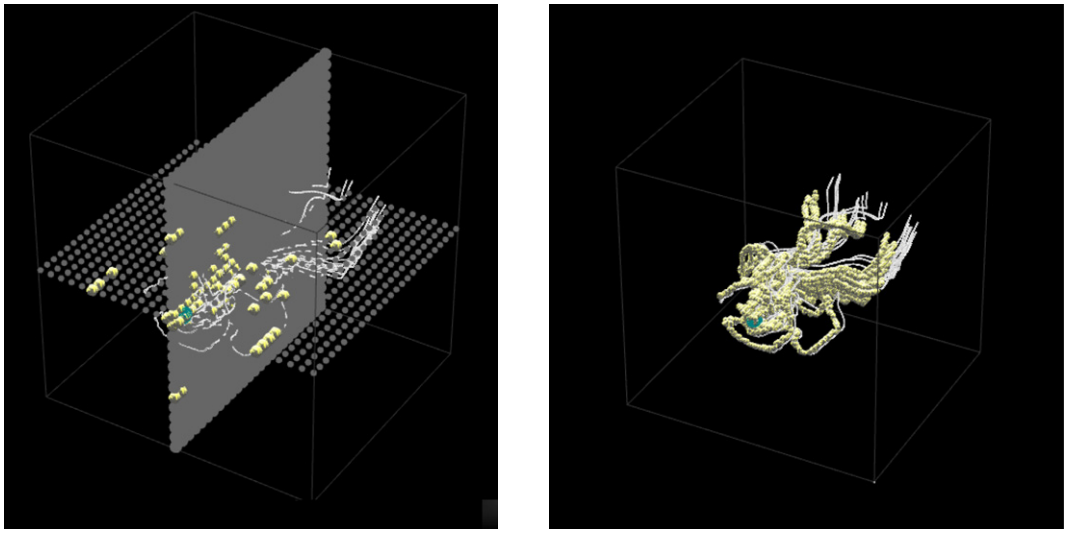

Figure 5. Glyph- and streamline-based exploration of the Sandia-D dataset, timestep 1. Left: glyphs placed on cutting planes. Right: glyphs placed along streamlines. Neither representation captures the rippling artifact.

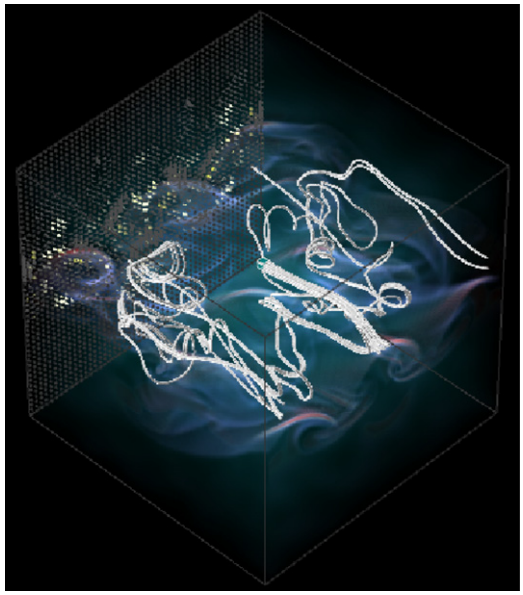

Figure 6. Hybrid tensor visualization of the mixing layer configuration. In this example, the user customized the transfer function interactively (color and opacity) to emphasize the median area of the flow corresponding to the two mixing layers. This snapshot highlights asymmetries in the combustion flow; the asymmetries are emphasized by streamlines.

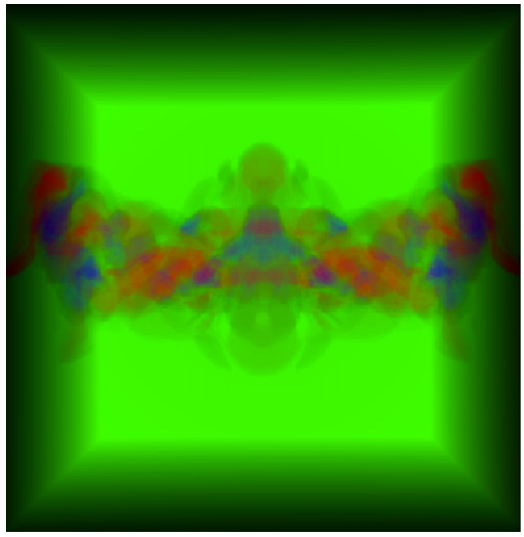

Figure 7. Classification-based descriptor with four color-coded clusters for the same 8M point mixing-layer combustion dataset.

from the zero-divergence outer zones (Figure 6). The density-gradient (Schlieren) descriptor was particularly appreciated over the velocity-gradient ones, for its ability to emphasize the shocklet boundary. The classification descriptor (Figure 7) also generated significant excitement for its ability to correctly identify and track regions of interest in the dataset.

Overall, the volume descriptors combined with the interactive streamlines (Figure 8) generated remarkable excitement. The researchers noted that the tangled, asymmetric streamlines in the mid-plane illustrate well the turbulent shear layer behavior where opposing streams of fuel and oxidizer meet. Illuminated streamlines were again appreciated, yet again not on the same scale as the interaction rates. The experts remarked repeatedly about the prototype's ability to focus on the interesting region of the volume (e.g., "[compared to this, in other tools] interactive selection is a beast"). The researchers also commented on the resolution and interactivity of the volume rendering, which was eight times more dense and, in their estimate, ten times faster than ParaView, the visual tool they had often used for volume rendering.

As before, increasingly complex glyphs, as well as their placement along streamlines (Figure 9) were deemed "cool," but did not lead to additional insight. We note that the domain experts were unable to discern shape variation among superquadric glyphs, even with filtering and under extreme magnification. The experts asked, in fact, for the

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

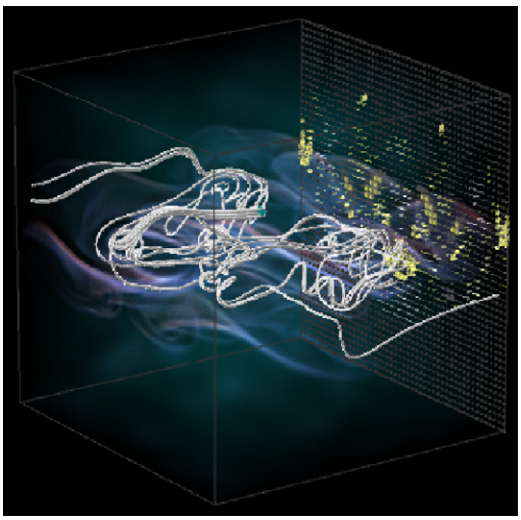

Figure 8. Exploratory visualization of dense tensor fields used in computational turbulent combustion modeling: hybrid rendering using a velocity-gradient volume rendering, plane-seeded glyphs and illuminated streamlines.

Figure 9. Visual velocity-gradient descriptor for an 8M point mixing-layer combustion dataset. The descriptor is augmented with streamline-placed superquadric glyphs. While the glyphs here and elsewhere are downsampled by a factor of 100, it should be noted that the scale of the datasets renders them illegible and of limited use, even when significantly magnified.

more complex glyphs to be colormapped into different categories; voiding the point of complex glyphs based on shape.

We note that the hybrid application compared favorably with existing visualization toolkits that offer a reduced set of similar, though less interactive, visualization features, such as ParaView, VisIt, Ensight and TecPlot. The essential advantages of using our utility for turbulent tensor visualization were its performance, the specific volume descriptors proposed and the fact that it was tailored to this specific application. The interactive rendering rates, the real-time selection of seed points for the streamline data, and the overall easy flow of interaction were major points repeatedly emphasized by the combustion hcientists.

Discussion

Following our design study performed in the application field, we reflect on the visualization insights gained. The design and evaluation process was completed through tight collaboration with combustion researchers. At its end we conclude that a hybrid approach works best for the target domain.

In the state-of-the-art practices in the target domain, interviews with the domain experts indicate that glyph visualization is nonexistent, aside from arrow glyphs. 2D contour plots, iso-surfaces and streamlines are readily available and popular, but are rapidly losing ground to volume rendering. Volume rendering is very popular for scalar quantities, in particular pressure and temperature, but also for derived quantities like vorticity, divergence, curl or swirl, although descriptor usage guidelines are nonexistent. Animations of flow vectors, iso-surfaces and fluid particles are common. Multi-views, although not linked, are routinely used to spot correlations between multiple scalar or vector fields.

Through our application development, we found that the use of traditional tensor descriptors such as ellipsoid or superquadric glyphs in the combustion domain is fraught with legibility problems. Elegant visual abstractions which capture the physical or mathematical aspects of the tensors were not necessarily intuitive to the application-domain practitioners. Increasing glyph complexity appeared to further hinder the experts' understanding of the data.

Glyph occlusion and clutter were additional major concerns, given the density of the tensor fields. In both of our dataset applications we found that the placement of glyphs along streamlines was not particularly beneficial. When the glyphs are evenly spaced throughout the volume, there is no formal guarantee that they capture the important features of the volume. The use of glyph packing[26] may alleviate this problem, but the legibility challenge outlined earlier remains. From our study, we conclude that in the application-domain glyphs are best used as details on demand, and when so, they should be mapped to cross-cutting planes.

Streamlet representations were not applicable to the combustion tensor data: due to the numerical range of the tensor field, they were essentially reduced to the linear glyphs explored earlier. However, the expert feedback indicates that streamlines through the velocity field continue to be useful, and that streamlines can be used as probes to highlight the structure of the flow. In both dataset applications, we found that interaction with these probes was far more important to the domain experts than the actual complexity of the streamline rendering. We also found, through our field foray, that iso-surface-based renderings are falling into less use in the application domain, and are being replaced by volume-based renderings.

In terms of volume-based rendering, the challenge lies in identifying appropriate visual descriptors for the tensor

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

fields. In this work, we proposed three classes of visual descriptors: velocity-gradient based, density-gradient based and classification based. Many other descriptors are possible. Each of the three classes we explored has complementary strengths: the velocity-gradient descriptors, in particular divergence, were best able to identify artifacts in the data; the density-gradient descriptor was best able to identify boundaries of interest; and the machine-learning descriptor was able to identify regions of interest through the data. Overall, we find that the proposition of relevant visual descriptors requires a solid understanding of the science behind the target domain.

Finally, interactivity and responsiveness of the visual descriptors are found to be of significant importance; for instance, basic interaction aspects are reported to make a significant difference in the effectiveness of streamlines as a probing tool. The overall use of volume renderings to provide global context was deemed highly relevant, while local probes are still required for detailed information extraction. Clutter and occlusion problems are shown to further benefit from an info-vis mantra and from an interactive filtering approach.

CONCLUSION

In this article we investigated, through a case study approach, the challenges posed by the application of tensor visualization research to the domain of turbulent combustion. In partnership with combustion researchers, we characterized the target domain, and iteratively explored the design space of tensor visual encodings through a parallel-prototyping approach. We further identified three sets of effective visual descriptors for volume rendering of the tensor data: velocity-gradient based, density-gradient based and classification based. We implemented a hybrid prototype using these visual encodings, evaluated their relative strengths on two examples of turbulent reacting flow and reported feedback from the target users in the application domain. The value of visualization as a comparison and debugging tool is well documented in this work.

In a further effort to bridge the gap between the tensor visualization and combustion communities and to facilitate the development of future visual benchmarking tests, we described in this article the two large datasets we used for validation. Finally, we analyzed the findings of our case study in the context of the larger visualization community. We summarized the design lessons learned through our collaboration with combustion experts, with particular emphasis on the need for global/local visualization, uncluttered visual encodings and fluid interaction. Given the prevalence of Big Data across sciences, we hope that this work will serve as a foundation for discussion and for future applications and contributions.

ACKNOWLEDGMENTS

The authors wish to thank Peyman Givi for his unwavering support, feedback and guidance on combustion modeling. They further thank the reviewers for their carefully considered feedback and suggestions. This work was supported by NSF CAREER IIS-1541277, and by NSF CBET-1250171.

REFERENCES

[1] R. L. Hack and V. G. McDonell, "Impact of ethane, propane, and diluent content in natural gas on the performance of a commercial microturbine generator," J. Eng. Gas Turbines Power **130**, 011509 (2008).

[2] N. Peters, *Turbulent Combustion* (Cambridge University Press, Cambridge, UK, 2000).

[3] A. Maries, M. A. Haque, S. L. Yilmaz, M. B. Nik, and G. E. Marai, "Interactive exploration of stress tensors used in computational turbulent combustion," *New Developments in the Visualization and Processing of Tensor Fields*, edited by D. Laidlaw and A. Villanova (Springer, 2012).

[4] J. C. Tannehill, D. A. Anderson, and R. H. Pletcher, *Computational Fluid Mechanics and Heat Transfer*, 2nd ed. (Taylor & Francis, 1997).

[5] S. Zhang, G. Kindlmann, and D. Laidlaw, "Diffusion tensor MRI visualization," *The Visualization Handbook*, edited by C. Johnson and C. Hansen (Academic Press, 2004).

[6] B. Jeremic, G. Scheuermann, J. Frey, Z. Yang, B. Hamann, K. I. Joy, and H. Hagen, "Tensor visualizations in computational geomechanics," Int. J. Numer. Anal. Methods Geomech. **26**, 925–944 (2002).

[7] P. Crossno, D. H. Rogers, R. M. Brannon, D. Coblentz, and J. T. Fredrich, "Visualization of geologic stress perturbations using Mohr diagrams," IEEE Trans. Vis. Comput. Graph. **11**, 508–518 (2005).

[8] V. Slavin, R. Pelcovits, G. Loriot, A. Callan-Jones, and D. Laidlaw, "Techniques for the visualization of topological defect behavior in nematic liquid crystals," IEEE Trans. Vis. Comput. Graphics **12**, 1323–1328 (2006).

[9] G. Kindlmann, D. Weinstein, and D. Hart, "Strategies for direct volume rendering of diffusion tensor fields," IEEE Trans. Vis. Comput. Graphics **6**, 124–138 (2000).

[10] A. Bhalerao and C.-F. Westin, "Tensor splats: visualising tensor fields by texture mapped volume rendering," *Sixth Int'l. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI'03), Montreal, Canada*, (2003), pp. 294–901.

[11] C. Dick, J. Georgii, R. Burgkart, and R. Westermann, "Stress tensor field visualization for implant planning in orthopedics," IEEE Trans. Vis. Comput. Graphics **15**, 1399–1406 (2009).

[12] X. Zheng and A. Pang, "Volume deformation for tensor visualization," *VIS'02: Proc. on Visualization'02* (IEEE Computer Society, Washington, DC, USA, 2002), pp. 379–386.

[13] W. Chen, Z. Ding, S. Zhang, A. MacKay-Brandt, S. Correia, H. Qu, J. A. Crow, D. F. Tate, Z. Yan, and Q. Peng, "A novel interface for interactive exploration of DTI fibers," IEEE Trans. Vis. Comput. Graphics **15**, 1433–1440 (2009).

[14] R. Jianu, C. Demiralp, and D. Laidlaw, "Exploring 3d DTI fiber tracts with linked 2d representations," IEEE Trans. Vis. Comput. Graphics **15**, 1449–1456 (2009).

[15] A. Kratz, B. Meyer, and I. Hotz, "A visual approach to analysis of stress tensor fields," Technical Report 10–26, ZIB, Takustr. 7, 14195 Berlin (2010).

[16] S. P. Dow, A. Glassco, J. Kass, M. Schwarz, D. L. Schwartz, and S. R. Klemmer, "Parallel prototyping leads to better design results, more divergence, and increased self-efficacy," ACM Trans. Comput.–Hum. Interact. **17**, (2010) 18:1–18:24.

[17] D. H. Laidlaw and A. Vilanova, *New Developments in the Visualization and Processing of Tensor Fields* (Springer, 2012).

[18] C. F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis, "Processing and visualization for diffusion tensor MRI," Med. Image Anal. **6**, 93–108 (2002).

[19] G. Kindlmann, "Superquadric tensor glyphs," *Proc. IEEE TVCG/EG Symposium on Visualization 2004* (IEEE, Piscataway, NJ, 2004), pp. 147–154.

[20] M. Zöckler, D. Stalling, and H.-C. Hege, "Interactive visualization of 3d-vector fields using illuminated stream lines," *Proc. 7th Conf. on Visualization '96* (IEEE Computer Society Press, Los Alamitos, CA, USA, 1996), p. 107–ff.

DOI: 10.2352/ISSN.2470-1173.2016.1.VDA-507

[21] A. Hadjadj and A. Kudryavtsev, "Computation and flow visualization in high-speed aerodynamics," J. Turbul. **6**, N16 (2004).

[22] A. McCallum, K. Nigam, and L. H. Ungar, "Efficient clustering of high-dimensional data sets with application to reference matching," *Proc. Sixth ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining, KDD '00* (ACM, New York, NY, USA, 2000), pp. 169–178.

[23] E. R. Tufte, *Envisioning Information* (Graphics Press, Cheshire, CT, 1990).

[24] *Readings in Information Visualization: Using Vision to Think*, edited by S. K. Card, J. D. Mackinlay, and B. Shneiderman (Morgan Kaufmann, 1999).

[25] J. Chuang, D. Weiskopf, and T. Moller, "Hue-preserving color blending," IEEE Trans. Vis. Comput. Graphics **15**, 1275–1282 (2009).

[26] G. Kindlmann and C.-F. Westin, "Diffusion tensor visualization with glyph packing," IEEE Trans. Vis. Comput. Graphics **12**, 1329–1335 (2006).