Interactive Exploration of Stress Tensors Used in Computational Turbulent Combustion

A. Maries, Md. A. Haque, S. L. Yilmaz, M. B. Nik, G. E. Marai

Abstract Simulation and modeling of turbulent flow, and of turbulent reacting flow in particular, involves 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 chapter, we discuss the challenges in dense symmetric-tensor visualization applied to turbulent combustion calculation, and analyze the feasibility of using several established tensor visualization techniques in the context of exploring space-time relationships in computationally-simulated combustion tensor data. To tackle the pervasive problems of occlusion and clutter, we propose a solution combining techniques from information and scientific visualization. Specifically, the proposed solution combines a detailed 3D inspection view based on volume rendering with glyph-based representations – used as 2D probes –, while leveraging interactive filtering and flow salience cues to clarify the structure of the tensor datasets. Side-by-side views of multiple timesteps facilitate the analysis of time-space relationships. The resulting prototype enables an analysis style based on the overview first, zoom and filter, then details on demand paradigm originally proposed in information visualization. The result is a visual analysis tool to be utilized in debugging, benchmarking, and verification of models and solutions in turbulent combustion. We demonstrate this analysis tool on three example configurations and report feedback from combustion researchers.

Adrian Maries, Md. Abedul Haque, G. Elisabeta Marai
Department of Computer Science, University of Pittsburgh, Pittsburgh, PA 15260, e-mail: {amaries, mdabedul, marai}@cs.pitt.edu

S. Levent Yilmaz
Center for Simulation and Modeling, University of Pittsburgh, Pittsburgh, PA 15260, e-mail: slyilmaz@pitt.edu

Mehdi B. Nik
Department of Mechanical Engineering and Materials Science, University of Pittsburgh, Pittsburgh, PA 15260, e-mail: smb51@pitt.edu

1 Introduction

Research into optimization of power generation systems for advanced energy and emissions performance has become increasingly important in the last two decades largely due to alarming environmental concerns and stringent regulations of NO_x and SO_x emissions, and other greenhouse gases such as CO_2 [8]. Microturbine generation systems are on the forefront of this research due to the promise of high-efficiency, ultra-clean systems that can be used to produce electrical energy as well as thermal energy in co-generation applications. There is a continuing need to develop and implement advanced computational tools for modeling and prediction of turbulent combustion for a wide range of mixing, fuel compositions and flow configurations [28, 11]. In fact, reliable and flexible computational modeling is considered a key element in achieving the objectives of modern gas-turbine design [22].

Tensor quantities are quite common in turbulence modeling, in particular in the three principal approaches to computational turbulent combustion [9]: Direct numerical simulation (DNS), Reynolds-averaged Navier-Stokes (RANS) and Large eddy simulation (LES). For example, in LES a spatially filtered form of the Navier-Stokes equations is considered, whereby only the flow features (eddies) that are larger than a characteristic filter size are resolved without approximation. Certain *subfilter* (or *subgrid*) quantities appear unclosed and need to be modeled. An important agent of such quantities is the non-linear convection term, which is defined as the *subgrid scale* (SGS) *stress tensor* [25] and the SGS scalar fluxes (for flows involving scalar transport) [9]. In eddy-diffusivity type LES models, SGS tensors are correlated to resolved quantities, such as the resolved strain tensor [21, 5].

Visually identifying the characteristics of such tensor quantities in finer details can bring significant insights into the computational modeling process. For example, DNS solutions are commonly used in verification of LES and RANS models, where the cases have typically been limited to simple canonical configurations [15]. In recent years, however, progress in DNS modeling has lead to DNS being applied to more realistic configurations with finer resolution [29, 16]. Visual exploration of tensor quantities in such complex configurations would ultimately steer the model verification and development process. Questions posed during analysis may be of the form, what artifacts do LES introduce in the simulation of a complex configuration, compared to the canonical DNS solution? How do these artifacts evolve over time? What are the regions and magnitude of error? In the long run, the visualization of tensor quantities could help turbulent combustion researchers identify regions of interest in the flow in order to design simulation schemes which are both computationally affordable and sufficiently accurate.

In this chapter we investigate the challenges associated with the exploratory visualization of tensor quantities in turbulent combustion simulations. We then propose, implement and evaluate an interactive prototype for exploring turbulent combustion tensor data that addresses some of these challenges. To the best of our knowledge, this is the first successful interactive visualization of time-varying stress tensors in the context of high-density turbulent combustion data. The result is a visual analysis tool developed through tight collaboration with researchers in computational

combustion. We evaluate the tool on several computational-combustion datasets of particular interest, and show the importance of the proposed approach for debugging the numerical simulation of complex configurations. Last but not least, we contribute a discussion of lessons learned, current limitations, and future directions of research as motivated by the driving computational application.

2 Tensors in Turbulent Combustion

2.1 Turbulent Combustion Modeling

A tensor is an extension of the concept of a scalar and a vector to higher orders. Scalars and vector are 0-th and 1-st order tensors, respectively. In general, a k-th order tensor can be represented by a k-dimensional array, e.g. a second order tensor is a 2D array (a matrix). A stress vector is the force acting on a given unit surface, and a stress tensor is defined as the components of stress vectors acting on each coordinate surface; thus it can be described by a symmetric 2-nd order tensor. The diagonal components of the stress tensor represent normal forces, i.e. compression and tension, and the non-diagonal components represent the shearing forces. The eigenvalues of the stress tensor are its principal stresses and the associated eigenvectors are the principal directions of the stress tensor. Since the stress tensor is symmetric, its eigenvalues are real. Strain tensor is a related quantity, the components of which are the components of the rate of deformation in each coordinate direction (i.e. derivatives of velocity). For the fluid flow motion, there are different so called constitutive models that relate strain to stress. In this work, we examine a Newtonian fluid, where the stress-strain relationship is linear.

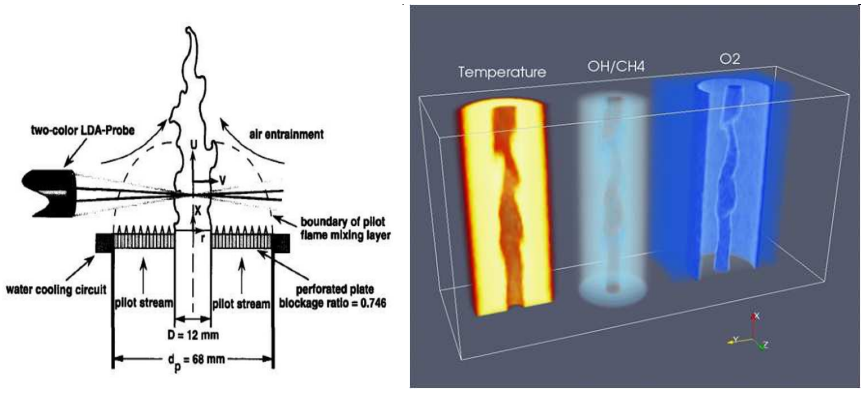

Fig. 1 The experimental setup which introduces the Bunsen burner experiment (left)[4], and example results from a corresponding computational LES model and simulation, showing the gaseous distribution at a specific time-frame (right) [44].

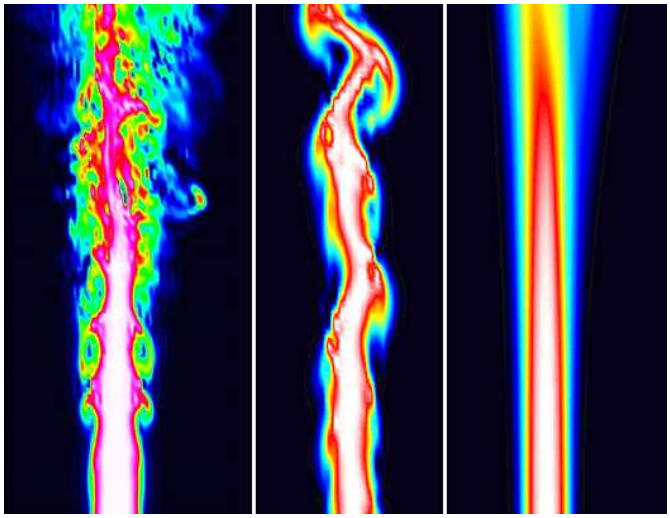

Fig. 2 DNS (left), LES (middle) and RANS (right) predictions of a turbulent jet. LES requires less computational effort than DNS, while delivering more detail than the inexpensive RANS.

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. Figure 1 shows an example configuration for a lab Bunsen burner, and example results from a corresponding computational model and simulation, showing the gaseous distribution at a specific time-frame.

Multiple numerical approaches exist for generating such computational models of combustion, most notably DNS, RANS and LES [9]. DNS consists of solving the Navier-Stokes transport equations of fluid flow and resolving all of the scales of motion. In RANS, the transport equations are time-averaged based on knowledge of the properties of flow turbulence, and the mean transport equations are solved. In LES, the spatially filtered equations of motion are considered, the solution of which portrays the large scale motion.

DNS, RANS and LES have complementary strengths. Figure 2 demonstrates the characteristically different results that can be obtained with these models. DNS provides a very accurate, model-free representation of the unsteady evolution of turbulent flows. However, applications are largely restricted by the computational power required by DNS and can be carried out only in limited and highly idealized cases [12, 29, 32]. RANS is significantly less expensive than DNS, and, as such, it is the more popular approach for engineering applications [24]. Finally, LES provides a balance between computational cost and the level of fidelity in the results. The LES methodology has been the subject of much modern research and is becoming increasingly more popular in combustion research [10].

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.). Details are beyond the scope of this chapter, and the reader is referred to many excellent books available on the subject [33, 20, 23]. In this chapter, we exemplify the visualization of stress/strain tensors, and therefore restrict the presentation to the pertinent subset of this PDEs, namely the momentum transport equation.

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

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

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. t is time, ρ is the fluiddensity, $\mathbf{u} \equiv [u_1, u_2, u_3]$ is the Eulerian fluid velocity, p is the pressure, and τ is the stress tensor defined as:

$$\tau_{ij} = \mu S_{ij} \tag{2}$$

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

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

From the computational modeling perspective, another tensor quantity of interest is the so-called turbulent stress tensor. In DNS of turbulent flow the conservation PDEs are solved in their exact form (as shown above for the momentum equation). On the other hand, in LES an alternative form is derived via the following spatial filtering operation:

$$\langle Q \rangle (\mathbf{x}, t) = \int_{-\infty}^{+\infty} Q(\mathbf{x}', t)\, G(\mathbf{x} - \mathbf{x}')\, \mathrm{d}\mathbf{x}' \tag{4}$$

where $Q(\mathbf{x}, t)$ is any turbulent quantity such as velocity or density. The filtering operation applied to the momentum equation, Eq. (1), gives:

$$\frac{\partial \langle \rho \rangle \langle u_i \rangle_L}{\partial t} + \frac{\partial \langle \rho \rangle \langle u_i \rangle_L \langle u_j \rangle_L}{\partial x_i} = -\frac{\partial \langle p \rangle}{\partial x_j} + \frac{\partial \langle \tau_{ij} \rangle}{\partial x_i} - \frac{\partial T_{ij}}{\partial x_i} \tag{5}$$

where the notation $\langle\ \rangle_L$ indicates the density weighted filter:

$$\langle Q \rangle_L = \frac{\langle \rho Q \rangle}{\langle \rho \rangle}, \tag{6}$$

Equation (5) greatly simplifies the filtered form of the PDEs and is commonly employed in LES and RANS of turbulent flow, instead of solving Eq. (1) directly. The important term in equation (5) is the turbulent stress tensor T which is defined as:

$$T_{ij} = \langle \rho \rangle \left(\left\langle u_i u_j \right\rangle_L - \langle u_i \rangle_L \left\langle u_j \right\rangle_L \right) \tag{7}$$

In LES the combustion flow solution is sought for the filtered terms *directly* by solving the simplified, filtered form of the equations (Eq. (5)). However, a closure problem originates from the non-linearity in the second term on the left-hand-side of Eq. (1) and is manifested as the unknown term $\left\langle u_i u_j \right\rangle_L$ in Eq. (5). **There is no equation in the filtered set of PDEs that corresponds to this term! Indeed, the whole literature on modeling of turbulent flow is providing closure models for T in one way or another.** [1] The same is true in RANS where ensemble averaging replaces filtering. For these reasons, the stress/strain tensors and the turbulent stress tensor are of particular interest to computational modelers.

2.2 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 presents several key challenges. First, because combustion datasets are the result of computational simulations often performed at high resolutions, they tend to be particularly dense. Such high densities lead naturally to clutter and occlusion problems when visualizing the data, as well as to slow interaction when visually exploring the data. Second, 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 show the 3D context of the data. At the same time, because the researchers are typically trained to visualize their data using 2D exploration tools such as TecPlot or ParaView, it is important to still give them access to similar 2D projection tools when exploring the data. Furthermore, care should be exercised when applying tensor visualization techniques from other fields (such as diffusion tensor imaging) to combustion data; because of the high density of simulation data and the range of values these data take, many existing tensor representations do not have an intuitive equivalent in combustion flow. Finally, many exploratory tasks in this context are visual comparison tasks, where the user aims to compare multiple steps of a specific numerical simulation, or different numerical simulations. To this end, it is important to provide a quick overall sense of the tensor field structure, while also allowing the user to explore particular regions of interest in more detail.

[1] In reacting flow there are additional and generally more significant closure problems associated with chemical source terms in the mass and energy conservations equations, however this is beyond the scope of this text. The reader is referred to the literature [9, 12].

3 Related Work

Several methods have been proposed for the visualization of tensor datasets. They include eigenvector color maps, glyphs, streamlines, volume rendering and volume deformation. Most of them are used to visually represent Diffusion Tensor Magnetic Resonance Imaging (DT-MRI) data, while some have been used in mechanical engineering to display different types of quantities, such as stress or fluid flow.

The most basic method uses eigenvector color mapping. It assigns an (R,G,B) color according to the (X,Y,Z) components of the principal eigenvector and a saturation level that depends on the magnitude of the anisotropy metric [41]. Colormaps are commonly used by mechanical engineers for component-by-component visualization of vector and tensor quantities. The limitation of colormaps is that they cannot show directional information and can only display one type of information at any one time, e.g., tensor component or eigenvalue.

Another approach utilizes glyphs, which are 3D icons whose shapes, colors, textures and locations correspond in some way to the properties of the data. There are many ways these mappings can be made. Generally, the shape indicates the directions of the eigenvectors at particular voxels, while colors can point to the value of the anisotropy at those voxels [41]. Two examples of glyphs are composite glyphs [38] and superquadrics [17], both of them used in brain imaging. In the engineering field, a type of glyphs called stress hedgehogs have been used to visualize mechanical stress [13]. The disadvantage of glyphs is that, for dense datasets such as the ones generated through simulation, they lead to clutter and occlusions.

Streamlines (sometimes called hyper-streamlines) and streamtubes are often used for tracing white-matter fibers in DTI-MRI datasets. They follow the direction of the main eigenvector through the volume and can be grouped using clustering algorithms into clusters that have anatomical correspondents [40, 43]. In addition to visualizing white matter fibers, streamtubes have been used to simulate blood flow through an artery in order to detect the effect of the flow on the walls of the artery [37]. In the mechanical engineering field, hyper-streamlines have been used to visually represent mechanical stress [13]. Slavin et al. use streamtubes to visualize topological defects in nematic liquid crystals, which are revealed by discontinuities in the orientation order of molecules in the liquid [31]. Unlike glyphs, which can show directional information only at certain points in the volume, streamlines can display directional information continuously in the volume. Nevertheless, streamlines are also prone to clutter and occlusions.

Volume rendering is a method of mapping points in the volume to the screen using a transfer function which controls how transparent or opaque each point is. Previously-used methods include barycentric opacity maps, hue-balls and lit-tensors [18]. Bhalerao and Westin use a textured mapped approach called tensor splatting [1]. Volume rendering of stress magnitudes, combined with tracing of short line segments to show stress direction, can help with designing and positioning of implants in hip joint replacement planning [7]. The volume rendering approach holds promise for dense datasets, although it can primarily show scalar information.

The volume deformation method is generally used to visualize mechanical stress. It allows the viewer to infer the characteristics of the tensor field by observing the effect it has on an object. Zheng et al. present two techniques: normal vector deformation and anisotropic deformation. The former is best used at indicating the direction of the tensor field while the latter can show the compressing and shearing properties of the tensor field [42]. Because of the fact that this method requires an object onto which the deformation has to be applied, it is in general not applicable to visualizing combustion flow.

Most of the work in the field tensor visualization of has been focused on symmetric tensor fields. There has recently been some interest in visualizing asymmetric tensor fields. One example is the use of tensor lines to analyze the gradient of the velocity vector field, an asymmetric tensor field [39].

Because of the complexity of a 3D tensor field, each tensor potentially having 6 components (3 diagonal and 3 non-diagonal elements), 3 eigenvalues and 3 eigenvectors to visualize, tensor visualization techniques often have occlusion and cluttering problems. In order to fix these problems, researchers often reduce the dimensionality and/or use interaction [3, 7, 14]. Both Jianu et al. and Chen et al. provide linked views of 3D and 2D representations of brain imaging data as well as several types of interaction with the 3D and 2D models, such as selection and coloring of white-matter fiber clusters. Dick et al. adopt a focus+context technique, where the user can move a circular region over the background visualization in order to view it in more detail. Sherbondy et al. use interaction in an application that uses dynamic queries to display neural pathways between volumes of interest [30].

Interactive filtering has been used more extensively outside the tensor visualization field, in particular in information visualization [6]. In general, the idea of bridging information visualization and scientific visualization is gaining momentum [26], and several examples have been published in recent years [7, 14]. Nevertheless, the challenges in this line of research relate to the specific forms scientific data takes.

4 Methods

To the best of our knowledge, this study is the first exploratory study of tensor visualization techniques in the context of turbulent combustion flow. To address the challenges outlined in Section 2.2, we propose a prototype for interactive visualization of combustion tensor quantities. It combines glyph-based representations — used as a 2D-projection exploratory tool — with real-time volume rendering — used as a 3D-context visual anchor, and with velocity streamlines — serving as flow salience cues. To further address the problem of clutter and occlusions we implement interactive filtering techniques, allowing the user to focus on and compare specific regions of interest of the tensor field. We begin by describing the data.

4.1 Datasets

We have employed results of three simulations in this study. The first dataset is a laboratory-scale Bunsen burner flame. The second is a canonical test problem employed in turbulent reacting flow research, namely the temporal mixing layer configuration. The third and most intriguing dataset is the result of LES simulation of a turbulent jet configuration known as the Sandia-D experiment.

The Bunsen burner is a centimeter-scale fuel jet surrounded coaxially by a hot pilot stream of burnt gas and impinging onto quiescent air. It is a high Reynolds number (24,000) configuration with a wide range of spatial scales, DNS of which is computationally unreasonable. The state of the art in DNS is the two-dimensional Bunsen burner configuration (slot-burner) with two-orders of magnitude lower Reynolds number than what is considered here. Therefore, only LES of this flame is available. The data employed in this work for the Bunsen burner are 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 101 in each direction (1M grid points).

The temporal mixing layer, on the other hand, 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 is 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 (approx. 8M grid points).

The Sandia-D dataset is very similar to the Bunsen burner, in that it is also 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 of burnt premixture of acetylene, hydrogen, and air. The difference is that the pilot flame is further surrounded by a co-flowing hot air stream. It has a slightly lower Reynolds number (22,400) than the Bunsen burner dataset, but the DNS is also 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 temporal resolution is dependent on the numerical specifications and is in the level of microseconds in the scales of this flame. The data used in this experiment is comprised of 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 streamwise direction and 160 in each of cross stream directions (approx. 5M grid points).

4.2 Glyph Representation

Representations that combine the different tensor components into a single image were of immediate interest to our collaborators in combustion research. In partic-

ular, previous studies of the Bunsen burner dataset routinely used component-by-component colormap representations (ParaView) to identify discontinuities in the tensor field, and we speculated that a combined representation of the tensor components would prove useful for this task.

As a first iteration, we pursued Westin's model of displaying the tensor as a composite glyph [38], which consists of a rod, a disc and a sphere. The eigenvectors and eigenvalues of the velocity strain tensor are calculated for each evenly spaced grid point, then mapped to the composite glyph. The direction of the main eigenvalue is deemed more important than the direction of the second and third eigenvalues, thus the rod component points in the direction of the main eigenvector. All that can be said about the second eigenvector is that is in the plane of the disc, while the third component, the sphere, does not retain the direction of the third eigenvector at all. The eigenvalues are sorted in descending order and then assigned to the rod, disc and sphere component of the glyph. Since, for many of the tensors in the Bunsen burner dataset, the three eigenvalues are very close in value, the length of the rod is set to 3 times the size of the main eigenvalue, while the radii of the disc and sphere are equal to the second and, respectively, third eigenvalues. This is done so that the direction of the main eigenvector can be distinguished more easily. Color is used to distinguish the three components of the glyph: green for the rod, magenta for the disc and red for the sphere. Figure 3 (left) shows a slice through the tensor field, each tensor being represented as a composite glyph. As anticipated, mapping the glyphs to the 3D flow lead to clutter and occlusions, even when the glyph field was subsampled by a factor of 25.

Following feedback from our collaborator, we simplified the glyph by removing the disc and sphere components of the composite glyph and only displaying the direction of the main eigenvector as well as the value of the main eigenvalue. For this purpose, the rod component was replaced with a line. Figure 3 (right) displays the same part of the tensor field as Fig. 3 (left), except it uses the piecewise linear representation of the tensors main eigenvector. 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 to deliver a sense of the 3D flow. The line glyphs were still kept in the final version of the prototype, the only difference being that an arrow was added to point in the direction of the main eigenvector. In the degenerate case, when the difference between the first two eigenvalues is less than 0.001 (heuristically determined) the corresponding glyph is displayed as a gray sphere, shown in Fig. 5 and Fig. 8.

4.3 Volume Rendering and Streamlines

To provide 3D-context to tensor quantities while reducing clutter and occlusions, we integrated volume rendering into the tensor field visualization. Volume rendering can be a very effective way of visualizing 3D volumes, although it requires mapping a color transfer function to a scalar value. The scalar quantity we turned to

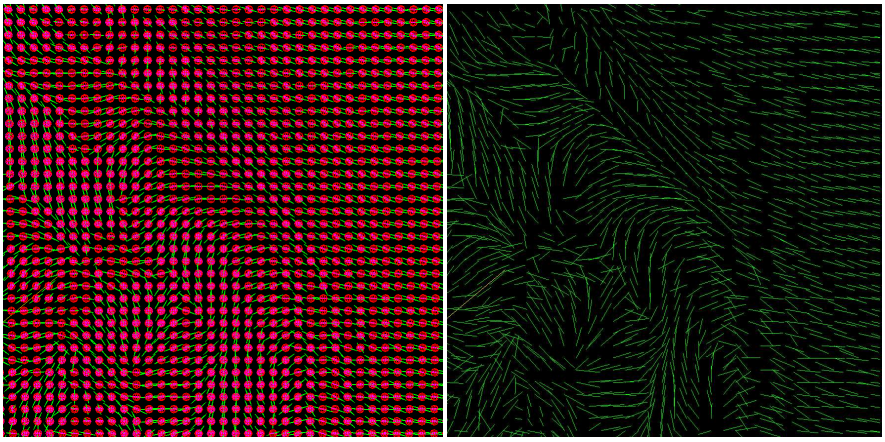

Fig. 3 Glyph (left) and piecewise-linear (right) visualization of the tensor field. In the glyph-based representation color identifies the various glyph components (green for the rod, magenta for the disc and red for the sphere.) To reduce clutter, the number of glyphs was subsampled by displaying only every fifth element along each direction, resulting in a dataset which was 1/25 of the original size. The subsampling factor has been empirically determined.

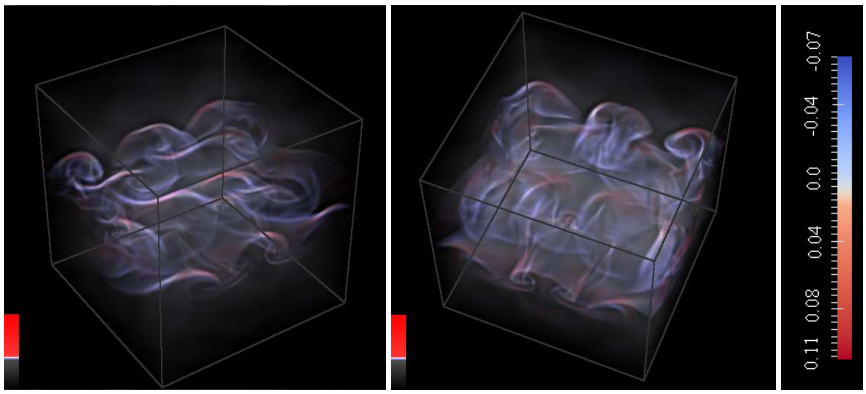

Fig. 4 Volume rendering of *divergence* of the temporal mixing layer dataset (left and middle), and color transfer function (right) used for volume rendering. Regions with saturated blue or red indicate higher magnitudes of divergence, either positive or negative. This rendering shows very clearly the 3D nature of the flow which is difficult to extract through other visualization methods.

for the volume rendering is *divergence*, which can be calculated as the trace of the strain tensor in Eq. 3. Under conditions of no mass source, divergence represents the change of density in time. The divergence reveals the 3D structure of the tensor field and effectively shows how the density of the fluid changes in different regions. Figure 4 shows two example images of divergence for the temporal mixing layer flow.

Our custom implementation of volume rendering uses ray-casting. The opacity transfer function for volume rendering was designed interactively so that high-

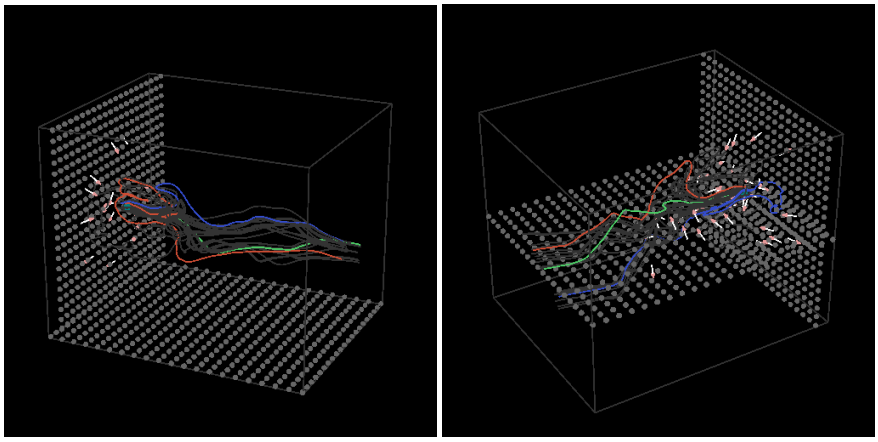

Fig. 5 Streamline and glyph visualization for the Bunsen burner dataset. The velocity field is unidirectional at the bottom of the volume (left). At the center of the volume, the velocity field is more varied (right). The streamlines emphasize the flow features in the central part of the simulation.

divergence regions had higher opacity and low-divergence regions had low opacity. This mapping highlighted the most interesting regions of the flow. The color transfer function was designed following the same principle; Figure 4 shows the color transfer function for the temporal mixing layer dataset. The user can customize transparency and color to better explore the underlying 3D stress tensor field.

To further emphasize the flow-context of the tensor data, the volume rendering was augmented with velocity streamlines — a technique borrowed from flow visualization. To create streamlines we calculated Runge-Kutta 4 integral paths through the velocity field. The starting points for the streamlines are a ”cube” of evenly spaced grid points. Figure 5 shows a combined visualization of glyphs and streamlines for the Bunsen burner dataset.

The visualization tool was implemented in C/C++ with OpenGL for rendering. Custom code was used for ray-casting.

4.4 Interactive Filtering

Early feedback from combustion researchers indicated that the glyphs, volume renderings and streamlines complement each other nicely in terms of the information displayed. However, the sheer volume of the combined information is overwhelming. To tackle this challenge, we followed a strategy based on Shneiderman’s infovis mantra [2] and Tufte’s principles [36, 35, 34]. Following Tufte’s principles, the information content of the image is maximized by combining volume rendering, streamlines and the glyph-based representations. To this combined representation we add zooming and filtering via interaction. Following further Shneiderman’s

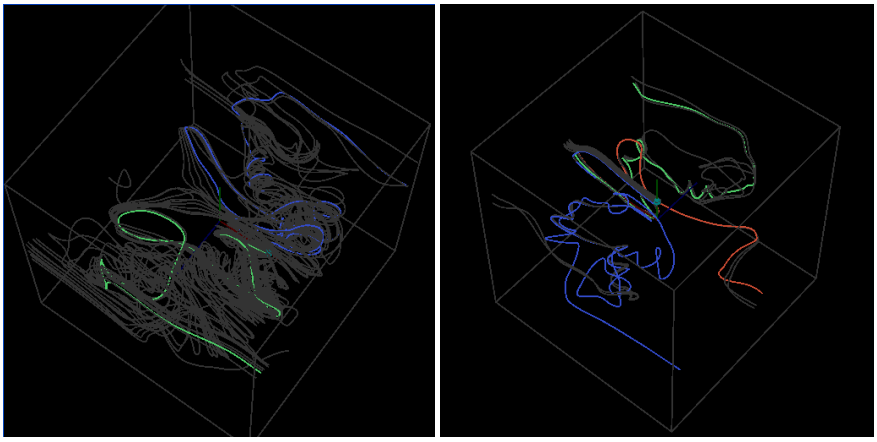

Fig. 6 Streamlines through the velocity field of the temporal mixing layer dataset. 27 streamlines are shown (left); even such small numbers clutter the image. Filtering streamlines interactively (right) helps reduce the clutter and enables the user to compare families of streamlines.

mantra, the volume rendering gives an overview of the flow and serves as a visual anchor, while the streamlines enable filtering of interesting regions, and glyph representations function as details on demand.

The visualization tool has two modes of operation, *explore mode* and *filter mode*. In the explore mode, a user can manipulate the scene and zoom in and out through mouse or keyboard interaction. Rendering the full resolution volume image during rotation or zooming slows down significantly the interaction. Therefore, during rotation or zooming, we use low-resolution volume rendering to maintain interactivity of the program: fewer rays are used for ray-casting than in the full-resolution rendering, and the remaining pixels are interpolated. We also reduce the sampling rate along a ray through the volume. The low resolution rendering produces an intensity image, while the full resolution rendering produces a color image. These modifications make the volume rendering fast and help to maintain interactive rates. The intensity volume-image during interaction helps to give an idea of how the volume is changing so that the user does not lose context between the start of the interactive step and its end.

In the filter mode, the user can focus on a sub-region of the flow by highlighting and comparing streamlines of interest. The filtering operation follows Tufte's principles and an approach originally used by the New York Times to explore stock market trends [6]. By default, all streamlines are muted gray, while mouse interaction allows the user to highlight and contrast two or more representative streamlines. Streamline seed-points can also be interactively dragged to new locations in the volume. After selecting a set of streamlines, the user can switch back to the explore mode and rotate/zoom the scene to better gauge the characteristics of the streamlines.

Figure 6 shows how the filtering operation can help declutter the visualization. In both modes, a user can control whether streamlines, volumes or the glyph-representations will be rendered or not. To further reduce clutter, the glyph representations are mapped to axis-oriented cutting planes, which are also controlled by the user.

Side-by-side views of multiple time-steps further facilitate the analysis of time-space relationships. While animated 3D views have certain advantages over side-by-side views (in particular with respect to display real-estate), the information visualization literature [27] and our experience indicate that side-by-side views are more effective at capturing time relationships.

5 Results and Discussion

We evaluate our analysis tool on the three datasets described in Section 2.2: a snapshot of the Bunsen burner simulation, a snapshot of the Temporal Mixing Layer simulation, and a three -step time sequence of the Sandia-D experiment. Two of the authors are computational combustion researchers who used the system and provided the following feedback.

Using the integrated exploratory system on the Bunsen burner dataset, it was noticed that the tensor field was remarkably smooth and homogeneous in the outer co-flow and the inner pilot regions. The tool was proven useful even for this elementary level of analysis, and will be further employed for more detailed LES datasets. For the mixing layer configuration, similar observations were immediately visible where the “mushroom” pattern around the shear layer at the mid-zone is distinguished well from the zero-divergence outer zones (Fig. 8). Interestingly, the analysis of the Sandia-D dataset provided a surprise — a rippling artifact surrounding the central jet — which can be seen in Fig. 7 and which showcases the advantages of the tool as a means of debugging numerical simulations. Shown in Fig. 7 are snapshots of the trace of strain tensor field at two instances in time, at 0.7 and 1.8th of a residence time, respectively; one residence time is equal to the total streamwise length divided by mean jet velocity, i.e., roughly the time it takes a fluid particle to traverse the whole length of the domain. The field shown is for the trace of the strain tensor. The regions at and near the jet core portray non-zero trace, and are indicative of the high spatial and temporal gradient in the gas-mixture density in these regions, in other words, of the active reaction zones with high compressibility. The far field trace on the other hand is much closer to zero and suggests (as expected from theory) close-to-zero compressibility of uniform and low density flow in these regions. This is especially clear in rightmost snapshot where the effects of high-frequency pressure waves are lesser compared to earlier in the simulation. This non-physical ripple effect could be attributed to numerical artifacts of the employed discretization scheme in the LES simulation, and is pronounced only in the incompressible regions, as clearly shown in the volume rendered tensor field. Overall, the researchers found the system as an exploratory tool “very good” and “cool”, and

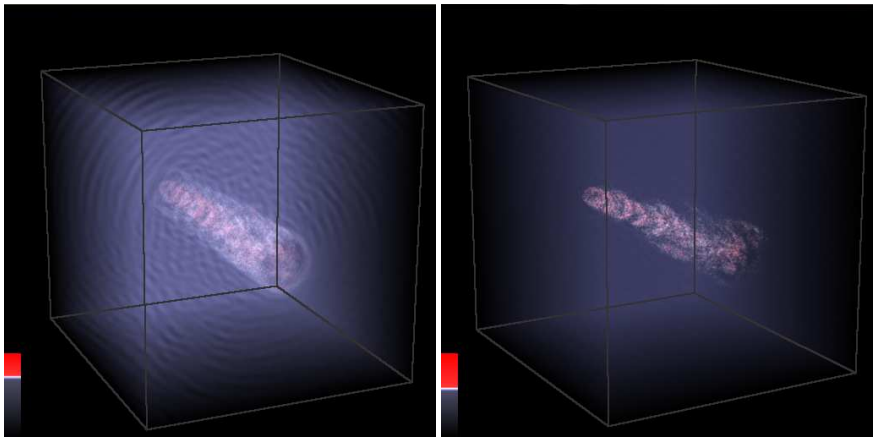

Fig. 7 Volume rendering of *divergence* of the third dataset (Sandia-D). Two timestamps of a reactive flow simulation are shown here, increasing in time from left to right. The rippling effect disappears as the simulation goes further in time. In this example, the user customized interactively the transfer function (color and opacity) to emphasize the central flow.

also commented on the potential explanatory power of the system to be utilized in comparisons of simulations against experimental data.

With respect to the various components of the visualization scheme, the glyph representation for the Bunsen burner dataset was considered somewhat useful, and only in 2D cross-sections. In 2D probes the flow patterns were "not sufficiently expressed", while in 3D the information was "hard to read" due to clutter and occlusions, as the glyphs became "intermingled and cluttered". Because the shown turbulent stress tensor field was mostly uniform throughout the domain, ellipsoid or line-style representations did not fare better; tensors did not "intuitively have direction". When exploring the more anisotropic mixing layer dataset, 2D planes of glyphs were considered useful, although only when combined with either streamlines or volume rendering to provide 3D context. Perhaps due to researchers' familiarity with 2D tools such as TecPlot, the 2D content as exposed by the glyphs was still considered important for these datasets.

The volume rendering combined with the interactive streamlines generated the most excitement. The researchers were particularly impressed by the ability to interact with the streamlines, by selecting and comparing individual streamlines and following their progress through the divergence volume. In particular, the researchers noted with the mixing layer dataset that the tangled, asymmetric streamlines in the mid-plane illustrate well the turbulent shear layer behavior where opposing streams of fuel and oxidizer meet. Both researchers made remarks about the tool'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

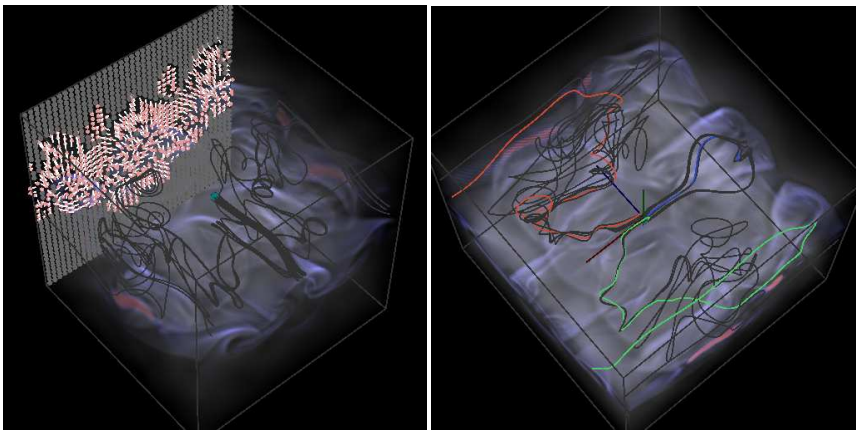

Fig. 8 Exploratory visualization tool for analyzing stress tensors used in computational turbulent combustion. The tool combines 3D inspection views based on volume renderings of divergence (shown in purple) with glyph-based representations (light gray and pink cutting planes), while leveraging interactive filtering of velocity streamlines (dark gray and color) and side-by-side views (not shown here) to clarify the structure of the tensor datasets. These snapshots highlight asymmetries in the combustion flow for a mixing layer configuration.

and, in their estimate, ten times faster than ParaView, the visual tool they had often used for volume rendering.

Overall, the integrated prototype compared favorably against existing visualization software that offers 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 and its being 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 scientists.

It is important to emphasize that the crux of the study in this paper is not the utility itself but rather the application of the aforementioned visualization algorithms in the context of turbulent reacting flow. Currently, to the best of our collaborators' knowledge, commonly employed visualization strategies for serious analysis in this field is limited to 2D cross sectional contours, or rather general 3D techniques which do not necessarily highlight the significant features required, in particular in the context of visual comparison tasks.

Much of the researcher feedback was directed towards potential extensions of the tool, as well as future application to other datasets. In particular, interest was expressed with respect to loading, exploring and comparing multiple datasets. A suggestion for future work, which came from both evaluators, was the implementation of a "field calculator" — an interface to allow the on-the-fly calculation of various scalar, vector and tensor quantities of interest. The "calculator" would include operators to add scalar fields or find a derivative or gradient for a certain field before displaying it.

In terms of limitations, all datasets we analyzed are canonical datasets of relative modest size, by computational combustion standards. If the dimensions of the volume were increased significantly, the current volume rendering visualization would require sub-sampling to stay interactive; more performant volume renderers would be necessary. Furthermore, in our approach we explored only two glyph representations. Alternative representations may help clarify whether glyph size and shape have a significant influence on the visualization. Also, since the glyphs are evenly spaced throughout the volume, there is no guarantee that they capture the important features of the volume. One possible direction for future work my be the use of glyph packing [19]. In terms of assisting combustion researchers in designing simulation schemes which are both computationally affordable and sufficiently accurate, techniques for segmentation of combustion tensor fields, perhaps interactively, hold particular promise.

One step further in terms of future research, exploratory visualization of massive combustion tensor datasets poses additional particularly significant challenges. Combustion simulations use thousands of CPUs to generate snapshots with millions of grid points; copying the data on a server for visualization is simply not feasible. Novel, memory-efficient exploratory visualization techniques will be necessary for such datasets.

6 Conclusion

In this chapter, we examined the challenges associated with tensor-field visualization in the context of turbulent combustion calculations and we proposed a prototype tool that can be used to visually explore combustion datasets. Our approach leverages interactive filtering and flow salience cues to clarify the structure of the tensor datasets, while effectively addressing the problems of occlusion and clutter. Side-by-side views of multiple timesteps facilitate the analysis of time-space relationships. The resulting framework enables an analysis style based on the overview first, zoom and filter, then details on demand paradigm originally proposed in information visualization. The result is a visual analysis tool to be utilized in debugging, benchmarking, and verification of models and solutions in turbulent combustion. We evaluated the proposed tool on three examples of turbulent reacting flow. Feedback from combustion researchers indicates that the tool is very useful in the exploration of turbulent combustion simulations, and emphasizes the urgent need of the field for visual analysis tools. However, significant research is still required to arrive at standard visualization in this application domain.

Acknowledgements Acknowledgements to the Pitt Visualization Research group and to NSF-IIS-0952720.

References

1. Bhalerao A, Westin C-F (2003) Tensor Splats: Visualising Tensor Fields by Texture Mapped Volume Rendering, Proc. Sixth Int'l Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI '03), pp. 294-901.
2. Card, S. K., Mackinlay, J. D., and Shneiderman, B.: Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann, (1999)
3. Chen W, Ding Z, Zhang S, MacKay-Brandt A, Correia S, Qu H, Crow J A, Tate D F, Yan Z, and Peng Q (2009) A Novel Interface for Interactive Exploration of DTI Fibers. IEEE Transactions on Visualization and Computer Graphics 15, 6, 1433-1440.
4. Chen, Y.; Peters, N.; Schneemann, G. A.; Wruck, N.; Renz, U.; Mansour, M. S. The detailed flame structure of highly stretched turbulent premixed methane-air flames Combust. Flame, 1996, 107, 223-226.
5. Ciofalo, M.: Large Eddy Simulation : A Critical Survey of Models and Applications. In: Advances in Heat Transfer, Vol 25, pp. 321-419. Academic Press, New York, NY (1994)
6. Cox, A., Xaquin, G. V., Leonhardt, D.: How This Bear Market Compares. (2008) http://www.nytimes.com/interactive/2008/10/11/business/20081011_BEAR_MARKETS.html
7. Dick C, Georgii J, Burgkart R, and Westermann R (2009) Stress Tensor Field Visualization for Implant Planning in Orthopedics. IEEE Transactions on Visualization and Computer Graphics 15, 6, 1399-1406.
8. Dunn-Rankin, D. (eds.): Lean Combustion: Technology and Control. Academic Press, New York, NY (2008)
9. Fox, R. O.: Computational Models for Turbulent Reacting Flows. Cambridge University Press, Cambridge, UK, 2003.
10. Givi, P.: Filtered Density Function for Subgrid Scale Modeling of Turbulent Combustion. AIAA J. 44(1), 1623 (2006)
11. Hack, R. L. and McDonell, V. G.: Impact of Ethane, Propane, and Diluent Content in Natural Gas on the Performance of a Commercial Microturbine Generator. J. Eng. Gas Turb. Power, 130(1), 011509 (2008)
12. Haworth, D.: Progress in probability density function methods for turbulent reacting flows. Progress in Energy and Combustion Science, (In Press, Corrected Proof), 2010
13. Jeremic B, Scheuermann G, Frey J, Yang Z, Hamann B, Joy K I, and Hagen H (2002) Tensor visualizations in computational geomechanics. International Journal for Numerical and Analytical Methods in Geomechanics, vol. 26, Issue 10, pp.925-944.
14. Jianu R, Demiralp C, and Laidlaw D H (2009) Exploring 3D DTI Fiber Tracts with Linked 2D Representations. IEEE Transactions on Visualization and Computer Graphics 15, 6, 1449-1456.
15. Johnsen, E., Larsson, J., Bhagatwala, A. V., Cabot, W. H., Moin, P., Olson, B. J., Rawat, P. S., Shankar, S. K., Sjgreen, B., Yee, H., Zhong, X. and Lele, S. K.: Assessment of high-resolution methods for numerical simulations of compressible turbulence with shock waves Journal of Computational Physics, (In Press, Corrected Proof), 2010
16. Kang, S., Iaccarino, G., Ham, F. and Moin, P.: Prediction of wall-pressure fluctuation in turbulent flows with an immersed boundary method. Journal of Computational Physics. **228**, 6753-6772 (2009)
17. Kindlmann G (2004) Superquadric Tensor Glyphs. In Proceedings of IEEE TVCG/EG Symposium on Visualization 2004. 147-154.
18. Kindlmann G, Weinstein D, and Hart D (2000) Strategies for Direct Volume Rendering of Diffusion Tensor Fields. IEEE Transactions on Visualization and Computer Graphics 6, 2, 124-138.
19. Kindlmann, G, Westin, C-F (2006) Diffusion Tensor Visualization with Glyph Packing. IEEE Transactions on Visualization and Computer Graphics 12, 5 (September 2006), 1329-1336.
20. Kuo, K., Principles of Combustion. John Wiley & Sons Inc., 2005
21. Lesieur, M. and Metais, O.: New Trends in Large Eddy Simulations of Turbulence. Ann. Rev. Fluid Mech. **28**, 4582 (1996)

22. Peters, N.: Turbulent Combustion. Cambridge University Press, Cambridge, UK (2000)
23. Poinsot, T., Veynante, D.. Theoretical and Numerical Combustion, 2nd Ed., R.T. Edwards, Inc., 2005
24. Pope, S.B.. Turbulent Flows, Cambridge University Press, 2000
25. Pope, S. B.: Advances in PDF Methods for Turbulent Reactive Flows. In: Andersson, H. I. and Krogstad, P. A. (eds.) Advances in Turbulence X, CIMNE, pp. 529536 (2004)
26. Rhyne, T., Tory M., Munzner, T., Ward, M., Johnson, C., Laidlaw, D.. Information and scientific visualization: Separate but equal or happy together at last? IEEE Visualization (Panel Proceedings), pages 611614, 2003.
27. Robertson, G., Fernandez, R., Fisher, D., Lee, B., Stasko, J.. Effectiveness of animation in trend visualization. IEEE Transactions on Visualization and Computer Graphics, 14(6):13251332, 2008.
28. Richards, G. A., McMillian, M. M., Gemmen, R. S., Rogers, W. A., and Cully, S. R.: Issues for Low-Emission, Fuel-Flexible Power Systems. Prog. Energy Combust. Sci., 27(2), 141169 (2001)
29. Sankaran, R., Hawkes, E. R., Chen, J. H., Lu, T. and Law, C. K.: Structure of a Spatially Developing Turbulent Lean Methane-Air Bunsen Flame. Proc. Combust. Inst. **31**, 1291-1298 (2007)
30. Sherbondy A, Akers D, Mackenzie R, Dougherty R, and Wandell B (2005) Exploring Connectivity of the Brain's White Matter with Dynamic Queries. IEEE Transactions on Visualization and Computer Graphics 11, 4, 419-430.
31. Slavin V, Pelcovits R, Loriot G, Callan-Jones A, and Laidlaw D (2006) Techniques for the Visualization of Topological Defect Behavior in Nematic Liquid Crystals. IEEE Transactions on Visualization and Computer Graphics 12, 5, 1323-1328.
32. Vreman, A., van Oijen, J., de Goey, L. and Bastiaans, R.: Direct numerical simulation of hydrogen addition in turbulent premixed Bunsen flames using flamelet-generated manifold reduction. International Journal of Hydrogen Energy. **34**, 2778-2788 (2009)
33. , Tannehill, J.C., Anderson, D.A., Pletcher, R.H.. Computational Fluid Mechanics and Heat Transfer, 2nd Ed., Taylor & Francis, 1997
34. Tufte, E. R.: The Visual Display of Quantitative Information. Graphics Press, Cheshire, CT (2001)
35. Tufte, E. R.: Visual Explanations: Images and Quantities, Evidence and Narrative. Graphics Press, Cheshire, CT (1997)
36. Tufte, E. R.: Envisioning Information. Graphics Press, Cheshire, CT (1990)
37. Wenger A, Keefe D F, Zhang S, and Laidlaw D H (2004) Interactive Volume Rendering of Thin Thread Structures within Multivalued Scientific Data Sets. IEEE Transactions on Visualization and Computer Graphics 10, 6, 664-672.
38. Westin C-F, Maier S E, Mamata H, Nabavi A, Jolesz F A, and Kikinis R (2002) Processing and visualization for diffusion tensor mri. Medical Image Analysis, 6(93-108).
39. Zhang, E., Yeh, H., Lin, Z., Laramee.,R. S. (2009) Asymmetric Tensor Analysis for Flow Visualization. IEEE Transactions on Visualization and Computer Graphics 15, 1 (January 2009), 106-122.
40. Zhang S, Demiralp C, and Laidlaw D H (2003) Visualizing Diffusion Tensor MR Images Using Streamtubes and Streamsurfaces. IEEE Transactions on Visualization and Computer Graphics 9, 4, 454-462.
41. Zhang S, Kindlmann G, and Laidlaw, D H (2004) Diffusion Tensor MRI Visualization. In Visualization Handbook. Academic Press.
42. Zheng X and Pang A (2002) Volume deformation for tensor visualization. In Proceedings of the Conference on Visualization '02. 379-386.
43. Zhukov L and Barr A H (2002) Oriented tensor reconstruction: tracing neural pathways from diffusion tensor MRI. In Proceedings of the Conference on Visualization '02. 387-394.
44. Yilmaz, S. L.; Nik, M. B.; Givi, P.; Strakey, P. A. Scalar Filtered Density Function for Large Eddy Simulation of a Premixed Bunser Burner J. Propul. Power., 2010, 26, 84-93