

Comparative Visualization of Ensembles Using Ensemble Surface Slicing

Oluwafemi S. Alabi,^a Xunlei Wu,^d Jonathan M. Harter,^a Madhura Phadke,^b Lifford Pinto,^b Hannah Petersen,^c Steffen Bass,^c Michael Keifer,^e Sharon Zhong,^e Christopher G. Healey, Senior Member, IEEE,^b Russell M. Taylor II, Member, IEEE^a

^aUNC-Chapel Hill, Chapel Hill, NC;

^bNorth Carolina State University, Raleigh, NC;

^cDuke University, Durham, NC;

^dRenaissance Computing Center, Chapel Hill, NC;

^eMichigan State University, East Lansing, MI

ABSTRACT

By definition, an ensemble is a set of surfaces or volumes derived from a series of simulations or experiments. Sometimes the series is run with different initial conditions for one parameter to determine parameter sensitivity. The understanding and identification of visual similarities and differences among the shapes of members of an ensemble is an acute and growing challenge for researchers across the physical sciences. More specifically, the task of gaining spatial understanding and identifying similarities and differences between multiple complex geometric data sets *simultaneously* has proved challenging. This paper proposes a comparison and visualization technique to support the visual study of parameter sensitivity. We present a novel single-image view and sampling technique which we call Ensemble Surface Slicing (ESS). ESS produces a single image that is useful for determining differences and similarities between surfaces simultaneously from several data sets. We demonstrate the usefulness of ESS on two real-world data sets from our collaborators.

Keywords: Ensemble visualization, isosurface, comparative visualization, uncertainty visualization

1. INTRODUCTION

We are collaborating with researchers studying high-energy physics, meteorology, cosmology, galaxy formation, microbiology, and biomedical data at Duke University, Michigan State University, and Sandia National Laboratories* who have developed models describing physical phenomena in their domain of interest. Each domain has generated data sets from simulations or experiments that they need to understand. These experiments or simulations are used to develop and/or validate the design of data mathematical models. Typically, the same simulation is run a number of times with different parameter settings to test the feasibility and stability of the model as well as its parameter sensitivity. This set of data is referred to as an *ensemble*. The data generated from a single simulation run is an *ensemble member*. From their data, our collaborators are interested in answering the following questions: Where do extracted surfaces from the ensemble agree, and where do they differ? When they do agree, what is the shape of the surface (at a large scale and in detail)? Where they disagree, which members are different and by how much do they differ? What are the surface features showing either similarity or the deviation among ensemble members?

In this paper we present ESS (Ensemble Surface Slicing) as a technique that has been designed for the direct visual comparison of surfaces extracted from 3D ensembles. ESS enables effective direct comparison of a handful of data sets, suitable for the exploration of the impact of varying a small number of parameters. It enables the comparison of multiple data sets over time. Our proposed technique has been demonstrated to be effective for model debugging and identifying shape variations due to parameter modifications. We show that ESS accelerates and improves the visual analysis of surfaces

Further author information: (send correspondence to Oluwafemi S. Alabi)

Oluwafemi S. Alabi: E-mail: oalabi@cs.unc.edu

*All authors cited in this paper are collaborators or investigators on the Modeling and Data Analysis Initiative (MADAI), a collaboration between Duke, Michigan State University, North Carolina State University, University of North Carolina at Chapel Hill, and Sandia National Laboratories.

extracted from ensembles. Additionally, we discuss how it satisfies the scientists’ goals of large-scale and local difference estimation, shape comprehension, and surface identification.

The display of these members in a single visualization communicates more information per unit area when compared to standard side-by-side viewing of each individual surface. When compared to classical side-by-side display, ESS removes the effect of change-blindness and also minimizes the impact on visual working memory’s limited capacity¹ by placing all relevant information in a single image. In this paper we discuss the development and application of this technique as well as the perceptual and cognitive principles guiding its efficacy.

The main contributions of this paper are:

- **ESS:** A new ensemble display technique that enables rapid identification of a handful of ensemble surfaces, rapid location of regions where surfaces are different, and determination of the interior/exterior relationship between neighboring surfaces at locations where they differ.
- **Real-world validation:** Effective application of ESS to the domains of RHIC hydrodynamics data and weather and its generalization to other ensemble domains.
- **User Feedback:** Feedback from domain scientists and users indicating the effectiveness of the technique compared to standard side-by-side display of ensembles.

2. RELATED WORK

ESS builds on principles of perceptual psychology as well as common visualization techniques to optimize the display of ensembles with more than two members. Below we describe several existing approaches to comparative visualization that influence ESS.

2.1 Comparative Visualization

Side-by-side. A generous body of work on techniques for simultaneous visualization of multiple parameters exists. Simple side-by-side views are a first approach to this class of comparative visualization. Spotting differences in traditional side-by-side visualizations of ensemble data appears trivial, but it is in fact quite challenging. Research has shown that image details across separate scenes cannot be remembered except where viewers have most recently focused their attention.² When the viewer switches attention from one image to the other, change blindness occurs. This is not a failure of the visual system per se, but rather is a result of inappropriate attention guidance due to different areas of the eye and brain responding differently to the scene once a visual interruption is introduced^{3,4}.



Figure 1. Spot the difference puzzle: There are at least four feature differences between these two images. Side-by-side visualization does not pre-attentively draw interest to areas of difference between two images, thus requiring a linear scanning of all potential features in each image. Scientists viewing side-by-side displays of feature-rich surfaces from ensembles face the same challenge.

“Spot-the-difference” puzzles that appear in the Sunday paper alongside mathematical and word puzzles testify to the fact that rapidly locating differences between images placed side by side is an inherently challenging visual task (see Fig. 1). Executing a “spot-the-difference” task on scientific data is similarly difficult due to its characteristic high feature density.

This makes it extremely difficult to rapidly identify characteristic differences when viewed side by side. The limitations of using side-by-side presentation to visualize ensembles is two-fold. First, rapidly identifying the difference(s) between data sets in an ensemble is inherently challenging. Second, beyond the problem of initial difference detection, the simple task of determining which surface is in front of the other in 3D between two views suffers from the lack of a common spatial reference frame. These two shortcomings make it difficult to conduct even a simple pairwise comparisons.

Translucency. Interrante provides an excellent description of the relevant perceptual issues for using translucency in comparative visualizations.⁵ Even though it is a common and simple method to display overlapping surfaces, it turns out that uniform translucency confounds surface shape perception away from object silhouettes. This makes any setting of uniform translucency ineffective for the display of surface features or inter-surface distances: either the outer surface is opaque and so hides all inner surfaces, or it is translucent and shape and distance between surfaces are imperceptible.

Slicing. Two methods that are closely related to ESS are planar contours and ribbons. Planar contours display the 1D curves in 3D that are the intersection of surface geometry with evenly-spaced parallel planes. The contours remove the geometry between planes while providing good overall understanding of shape without occluding interior structures. Ribbons, on the other hand retain more of the surface, removing the geometry between alternating pairs of planes. Bauer-Kirpes et. al. used a ribbon-like technique, called “barrel hoops” for the display of the outer surface in multiple surface medical applications.⁶ ESS generalizes slicing techniques to apply to more than two surfaces. This allows for it to be used for displaying small ensembles. This “ribboned” display of the surfaces provides good overall shape relationship comprehension because similar to tangent ribbons, ESS slices retain more of the surface, removing the geometry between alternating pairs of planes. This leaves axis-aligned surface patches with their internal features visible. The addition of animation of these slices over each ensemble member respectively minimizes loss of topology and ensures that the entire shape is seen.

Sparsely-Opaque Textures. Techniques have been developed that make use of textures with small, linear opaque features (sometimes in multiple orientations) to display the large-scale shape of outer surfaces while enabling the viewer to see past them to opaque inner objects. Interrante used curvature-directed strokes to display radiation dose surfaces above tumors⁵ (see Fig. 2). Weigle extended this technique to work on a pair of non-nested surfaces, rendering the interior surface opaque with nearest-approach curves dropped from opaque cross-glyphs on the exterior surface.⁷ Rheingans retiled irregular surfaces so that uniform circular and hexagonal textures could be applied⁸ (see Fig.2). Along the same lines as Interrante, Bair et. al. used a genetic algorithm with user-supplied fitness evaluation to determine the optimal parameters for texture generation to display one surface above another.⁹ This system produced an upper texture that had well-separated parallel lines of opacity along the outer surface, quite similar in appearance to those found in slicing methods. ESS extends this concept of upper texture to multiple other surfaces in a manner that enables the perception of a single surface when there is agreement among the members of an ensemble. ESS also facilitates the rapid determination of different surfaces where members differ. By adopting the use of color, as presented in research, to distinguish surfaces, ESS extends findings from Weigle and Interrante’s techniques to nested as well as non-nested geometries while providing consistency of tiling between data sets (required where ensemble sets agree).

Other comparative visualization methods attempt to counter the limitations posed by side-by-side and translucency comparative visualizations. These methods either replicate, codify or fuse data in space for a concurrent display of parameters.¹⁰ For simultaneous display of parameters, Kirby et al. combine multiple values in a 2D flow image and attempt simultaneous display.¹¹ Bokinsky uses color coding and size of simple glyphs to visualize multiple data sets on a single surface.¹² Laidlaw et al. utilized ellipsoids and brush strokes of varying intensity to visualize tensor images.¹³ Healey and Enns proposed accomplishing multiple parameter visualization by varying color and texture attributes.¹ Urness et al. use color weaving for flow visualization¹⁴ while Gosset and Chen use color blending for the presentation of fusion of multiple parameters.¹⁰ For most of these techniques, the conclusion was that with more than four parameters, the technique’s ability to convey information reduces considerably.¹⁰ ESS uses a similar sparse-surface display approach, but aims at a different goal: rather than attempting to display multiple variables on the same surface, ESS displays multiple instances of similar surfaces.

2.2 Uncertainty Visualization

Scientific data sets often have associated estimates of local accuracy. Uncertainty visualization techniques present data in a manner that enables the data to be visualized along with its uncertainty. The goal is often to make the viewer aware of locations and degrees of uncertainty in their data. This desire draws a correlation between comparative visualization




Figure 2. (a) Communicating surface shape using principal-direction-driven 3D LIC stroke texture. (b) Weigle’s technique for visualizing nested surfaces by rendering the interior surface opaque with nearest-approach curves dropped from opaque cross-glyphs on the exterior surface. (c) Rheingans surface retiling technique

and uncertainty visualization—the comparison of one surface geometry to another (e.g. an uncertainty percentile surface to estimated base surface). Several researchers have developed techniques for the display of surfaces with positional uncertainty. Pang et. al. categorize a number of approaches to this problem including modulating the properties of a single surface and adorning the surface with glyphs.¹⁵ Wittenbrink et. al. address surface uncertainty by rendering fat surfaces (translucent envelope of uncertainty sliced with a cutting plane to show cross-section of difference). Grigoryan and Rheingans also add point-wise displacements to the surface to indicate the spatial region in which it might lie.¹⁶ Other methods include oscillating surfaces and line glyphs.¹⁷ Most of these techniques intentionally occlude or hinder the perception of the original surface in regions of high uncertainty (thus preventing unwarranted reliance on surface position). ESS is designed to preserve the perception of surface geometry for each member of an ensemble in regions where they differ.

2.3 Perception

2.3.1 Perceptual Features

The goal of visualization is to map features such as hue, intensity, spatial location and size to individual data elements in a manner that maximizes the perceptual salience of relationships among the relevant features. This, of course, relies heavily on the understanding of visual perception from perceptual psychology. For a visualization to be effective, the contrast-saliency relationship between its encoding features needs to be matched in importance to the purpose of the visualization.

ESS makes use of the perceptual features of luminance contrast, shape and color to communicate shape boundaries, identify salient features belonging to a particular surface, and differentiate between surfaces, respectively. Research has shown that these features “pop out” during visual search while minimizing feature interaction. Thus, rapid identification of salient information is possible with little or no increase in the time required to complete a search task as the number of distractors increase. ESS avoids luminance contrast (isoluminance) in regions where differences are not present to avoid drawing attention to them, and to avoid distorting the perception of their shape.

2.3.2 Color

One of the most important and often used perceptual features is color. Healey proposed using the CIELUV color space to choose colors that allow an observer to rapidly and accurately search a display for any given set of colors based on their linear separability, color category, and color distance.¹⁸ Based on this, ESS chooses colors along a path in perceptual color space. A parametric increment is chosen to pick discrete colors along the specified path and then each ensemble member is arbitrarily assigned an isoluminant color. The colors mapped to each surface falls outside the convex hull of colors chosen to represent all others, thus maximizing the efficacy and speed of visual search. This ensures maximum separability and distinct differentiation between datasets while maintaining isoluminance on each surface.




Figure 3. A banana, apple, grape, and orange shown using ESS. Luminance discontinuities between strips draws attention to areas of difference between neighboring surfaces.

3. ESS METHODOLOGY

ESS uniformly slices the overlapping spatial extent of all surfaces in the ensemble along one world-space axis. The surface from one member of the ensemble is visible in each slice. Which surface is visible alternates from slice to slice, iterating through the available ensemble members several times over the spatial range. The technique is shown on dis-similar surfaces in Fig. 3 and on similar surfaces in Fig. 4. In each case, it lets the user rapidly locate differences between the surfaces and estimate the magnitude of separation between surfaces. User feedback indicated that the choice of slice width enhanced certain properties of the ensemble. When fewer slices are chosen, the silhouette of the shapes in the view are rapidly processed. Larger slices, however, ensure that the features of the ensemble members are visible and not obscured by slicing artifacts. The selection of slice width is reliant on the particular query posed by the viewer. If overall shape silhouettes are of interest, thinner slices are desirable; if understanding the variation across ensemble members at boundaries of a particular surface is required, then thicker slices are desirable. When the surface normals at slice pair contours differ, a luminance discontinuity is observed between strips. Such luminance edges are easily detected visually and draw attention to regions of difference in the space. These axis-aligned contours also provide direct visual display of the separation between the surfaces; interactive control over the slicing direction enables the user to align this with the local surface normal if estimation of the 3D separation is desired.

The desire to understand the total shape of each surface at local scale presents a challenge. ESS addresses this by animating the slices along the slice. This enables the viewer to quickly understand the shape of the members of the ensemble at any moment in time. Animation ensures that the full extent of each ensemble member is visible in the display over the course of the animation sequence. This has been implemented by adding a time-varying offset to the slice boundaries that causes each slice to move at uniform speed in the slicing direction. This overloading of the time dimension provides a means for features on the surface of all ensemble members to be understood for a single time step while simultaneously addressing the temporal nature of the data.

4. APPLICATION TO RHIC SIMULATION DATA

One of the domains where ensembles are prolific is in the field of physics. Some of our collaborators are studying hydrodynamic simulations for relativistic heavy ion collisions. Our collaborators are particularly interested in Quantum Chrono-Dynamics (QCD) which describes hadronic matter (matter susceptible to the strong interaction force).

Background. Relativistic heavy ion collisions offer the opportunity to study strongly interacting matter under extreme conditions. At high temperatures and/or densities the quantum field theory that describes the strong interaction—namely quantum chromodynamics (QCD)—predicts a phase transition to a new phase of matter in which quarks and gluons can move independently of each other—the so called quark gluon plasma (QGP). The quark gluon plasma has existed a few




Figure 4. Four gaussian blob surfaces: Viewing these surfaces side-by-side requires saccading between all four views to try to locate the region of difference. ESS simplifies this visual search by using luminance discontinuity to draw attention to discontinuities/buckling on the surface. Localized similarity is characterized by adjacent continuous color bands. In the combined view, viewers can more rapidly see that three of the surfaces are the same, with the orange one differing.

microseconds after the Big Bang in the early universe which is why heavy ion collisions are referred to as being “little bangs in the laboratory”. The driving problem in the study of QGP is that the deconfined quanta of a QGP are not directly observable. Finding a clear and unambiguous connection between transitory QGP state and what is observed as the hadronic final state is a desired goal of relativistic heavy-ion research. The data resulting from these runs with varying simulation parameters forms ensembles.

Data. Fig. 5 shows the result of ESS on QGP datasets glauber, urqmd, pcasc representing the hydrodynamic evolution of a Au + Au collision at $E_{cm}=200$ GeV as studied by Hannah Petersen and Stephen Bass of Duke University. Here, the ESS technique has been applied to an event-by-event transport + hydrodynamics description of Au + Au collisions at the highest RHIC (Relativistic Heavy Ion Collider) energies. Different approaches for the initial conditions are employed and the aim is to explore the differences in the subsequent ideal hydrodynamic evolution. The different impact parameters (b_0 and b_7) describe the centrality of the collision; this refers to either head-on (b_0) or off-axis (b_7) collisions in the x -direction. The result of this collision leads to an almond shaped asymmetry in the transverse (x - y) plane. Our collaborators are interested in seeing differences in the energy density, temperature, and net baryon density as time progresses. They want, in general, to better understand how the different initial conditions affect the evolution—especially the evolution of initial coordinate space asymmetries through the simulation. Fig. 6 shows the results from a second set of simulation runs representing a single gold-gold nuclei collision event. The ESS image produced provides insight to the effect of these initial conditions on the turbulence and momentum centrality and provides more information from which scientists can test hypotheses or debug code (e.g. when the simulation differs vastly from known and proven physical condition constraints). After watching a time-animation of this simulation using ESS, our collaborator was surprised to see that the initial conditions in one of the simulations (psasc) led to an initially larger region that suddenly disappeared from the simulation. They looked at both the physics calculations and their simulation codes grid conditions to locate the cause. The event turned out to be an issue with the simulation grid resolution, which has since been fixed.

ESS was also applied to weather fire burn simulation data (see Fig 7). The parameter varied between experiments is the surface heat flux associated with the fire. As the simulation does not resolve the actual combustion process (the simulation is at too coarse a resolution), the effect of the fire is accounted for by examining heat flux surfaces. A larger surface heat flux implies a greater fire intensity (i.e. the fire is pumping more heat into the atmosphere) and vice versa. Using ESS, Michael Kiefer and Sharon Zhong were able to identify where the surfaces are equivalent (characterized in the visualization by continuous banded isoluminant strips)(see Fig. 7). As increases in heat flux correlate to increases in potential temperature, examining the resulting potential temperature surfaces given varied heat flux initial parameters provides a quick way to measure the effect of burn intensity on the atmosphere given specific atmospheric conditions. The




Figure 5. Hydro-physics Fluid Simulation Side-by-side vs. ESS: 4 energy surfaces where $e \in [0, 3]$. Leftmost quad (clockwise from top to bottom): *pcasc*, *urqmd*, *glauber_b0*, and *glauber_b7* surfaces. In the side-by-side, it is easy to see which shapes are the same (e.g. *pcasc* and *urqmd* share similar shape while *glauber* simulation initial conditions results in very different shapes). Due to a lack of a common frame of reference, it is challenging to determine where surfaces cross each other in the side-by-side view. ESS (right image) provides understanding of relative scale, surface orientation, and large and small scale differences.




Figure 6. Hydro-physics Fluid Simulation Side-by-side vs. ESS: 4 energy surfaces where $e \in [0, 3]$. In the resulting video, one surface in the simulation disappears for a couple of seconds only to reappear again. Domain scientists used this visualization to determine that the resolution of the grid inadequate. Once modified, visualization of the data provided more consistent results.

domain scientists used the ESS visualization in Fig. 7 to determine the effect of the initial conditions on the resulting heat profile and smoke dispersal of the fire.

5. USER FEEDBACK AND RESULTS

We collected feedback on our technique from our MADAI collaborators to evaluate the robustness, practicality and applicability of ESS. First, we compared our system to the current workflow of domain specialists. Then we asked the same subjects to compare ESS to side-by-side views on an image level without scene interaction. Participants were specialists familiar with working with comparative visualization of domain-independent surfaces. All three participants completed 2 tasks and viewed 16 images, with two users having 15+ years of experience and one having less than 1 year of experience.




Figure 7. Weather Fire Burn Simulation Side-by-side vs. ESS: 4 potential temperature surfaces where $t \in [287^\circ\text{F}, 345^\circ\text{F}]$. The bottom two surfaces in the left quad are almost identical; the other two surfaces (top row of quad), are the most different in the set. The ESS visualization (right image) shows both similarities as well as high/low frequency differences between all surfaces in the ensemble. The large differences near the peak in the upper part of the images are visible in side-by-side visualization, but the subtle changes in location of the features in the lower right is difficult to see. Dark gaps between sets rapidly draw attention to both differences in the ESS view.

5.1 Workflow

Our domain scientists did not have a standardized workflow for visualizing ensembles of 3D surface data. Typically, they analyze datasets using Matlab, or some other mathematical or 3D rendering software before performing a side-by-side visualization of the surfaces extracted. Otherwise the data is analyzed abstractly without direct analysis of physical shape. More recently, they have been incorporating the use of the open source ParaView visualization system to generate visualizations. This allows for a more domain-independent, intuitive and standardized workflow for data assessment. ESS functions as a plug-in to this system. Our work provides a standardized and extensible visualization tool that utilizes ParaView’s well-defined and flexible framework in a manner that enables rapid modifications to future iterations of ESS.

With our collaborators, we arrived at two classes of tasks—identification tasks and relationship tasks. Identification tasks involve identifying the number of surfaces independently discernible in an ensemble using a visualization technique as well as the duration of time required to ascertain the count. Relationship tasks revolve around the “spot-the-difference” problem: how fast can differences be observed in the visualization? Where a difference is observed, how does one surface relate to its neighbors (interior, exterior, almost identical)?

Feedback was first obtained about experience with working with 3D data and comparative visualization. ESS was then discussed and explained prior to users performing a sequence of tasks using visualizations created with our software prototype. Participants in our study were introduced to each visualization condition as well as to the manner in which regions were classified (globally, or by slice/slice-neighbor proximity). Assistance beyond explanation of properties of ESS visualizations and their respective implications was not given. This was meant to test whether the system is intuitive and rapidly deployable.

5.2 Evaluation

To demonstrate the efficacy of our technique, ESS was applied to two synthetic ensemble datasets, one composed of fruits and the other of a set of gaussian surfaces. The fruits were chosen primarily for their easily recognizable form – completing the silhouette of the shapes of the fruits is pre-attentively processed. The viewer’s attention is quickly directed to regions in the ESS visualization where there is a surface luminance discontinuity. Similarly, the gaussian blobs were chosen to show that even with more abstract shapes, the luminance discontinuity between strips draws attention to regions of difference and allows for dataset(s) that locally differ from the others to be rapidly identified.

Next, a short two-task questionnaire was administered to determine ESS’s efficiency on relationship tasks. To ensure that understanding of shape was not obfuscated by specificity to domain, synthetic series were used. Using a commonly understood shape ensemble (in this case fruits) eliminated the need for understanding information other than shape class of a given ensemble. The same study was conducted on arbitrarily generated gaussian surfaces to test our claim of ESS’s efficacy even on surfaces of arbitrary geometry (see Fig. 4). Participants were asked to identify the surface difference in

an image and then identify a surface’s relationship to the other members in the ensembles. Eight ESS images and eight side-by-side images were provided for each task. Furthermore, the amount of time it took to complete each was recorded for each participant as well as whether the task was completed accurately.

The results of the study suggest the strength of ESS in conveying spatial and inter-member orientation information. According to feedback, this was difficult to achieve using side-by-side visualization. Using ESS, the users were able to accurately and rapidly detect and compare surfaces in the ensemble that did not locally conform to the shape of other members in the ensemble. In ensembles with many non-uniform surfaces, providing a global similarity or difference estimate proved most difficult as this situation resulted in a rather noisy image. However, this still showed significant improvement on side-by-side views. On average across all tasks, task completion using ESS took 28% of the time required to complete the same task using side-by-side. No errors were made in either the ESS or the side-by-side conditions.

Table 1. Side-by-side vs. ESS: Results Achieved in Time(seconds) Per User

Image	User 1 time	User 2 time	User 3 time	Type	Mean	ESS:SBS
1	2.5	2.2	6.2	SBS	3.57	0.35
	1.0	1.8	1.0	ESS	1.27	–
2	2.5	2.0	2.0	SBS	2.23	0.35
	0.5	1.4	0.5	ESS	0.80	–
3	2.5	2.6	2.3	SBS	2.46	0.27
	0.5	1.0	0.5	ESS	0.67	–
4	2.5	3.2	1.8	SBS	2.10	0.30
	0.5	0.9	0.5	ESS	0.63	–
5	2.5	2.4	7.1	SBS	3.93	0.20
	0.5	1.1	0.5	ESS	0.80	–
						0.28

Compared to their earlier workflow, the users indicated that the largest advantage of our technique stemmed from its ability to rapidly see and recognize which slices in a particular region differed from its neighbors. While this is more obvious for large differences, even small scale differences on slice boundaries were pre-attentively assessed. Thus, ESS provides users with a rapid method for performing visual analytics. Our findings encourage the use of this technique as a standard and simple way to gain understanding of ensembles. Furthermore, its implementation on the standardized visualization platform (ParaView) allows for other standard visualization techniques to be applied.

6. DISCUSSION

Besides the domains of hydro-physics, weather, and biomedicine, other fields such as geography require the simultaneous display of surfaces and understanding how they relate to each other spatially on global and local levels. We introduce a new technique called Ensemble Surface Slicing (ESS) to visualize regions of similarity and difference among surfaces extracted from ensemble simulations. ESS communicates inter-surface positioning and relationships while minimizing surface distortion. Drawing inspiration from Escher drawings in conjunction with Bauer-Kirpes et al.’s “barrel hoops” used for display of dosage surfaces in medical applications,⁶ ESS displays spatial relationships of surfaces generated from ensemble data sets.

6.1 Advantages

The feedback collected from users indicated that they spent less time to complete a given task using ESS as opposed to side-by-side. All users appreciated the ability to quickly compare shapes in a visualization. Domain specialists also indicated that a major advantage of ESS is that the coloring of slices coupled with the per time step animation of slices over the entire surface of an ensemble member enabled interpolation of overall shapes of ensemble members in regions where portions of the shape were not being shown at every instance in time. As such, there was no real need for reference to a legend until a particular difference was noted. When such a case arose, the user could simply use the color to identify which member is of interest at that particular instance in time. The ability to see and compare multiple slices in a single image and the ability for ESS to plug into current workflow (as a standalone application or as a plug-in for ParaView) received positive reviews from our domain scientists and other users. The users indicated that prior techniques of side-by-side

visualization, or even more advanced techniques like checkerboard visualization (and its variants) would have been very time consuming for completing the requested tasks.

Besides being a rapid way to identify differences between members of an ensemble, ESS can be applied to the visualization of localized uncertainty. It may be used to display the uncertainty in a surface by first creating N sample surfaces whose distribution in space matches that of the uncertainty (perhaps the 25%, 50%, ... percentile surfaces) and then displaying the set of uncertainty sample surfaces using ESS. Uncertainty is a local measure so the understanding of global shape is often not as important as understanding the local shape variation. This plays to our technique's unique ability provide rapid local visual analysis.

Users were able to conduct rapid difference detection in at least 4–5 data sets in one view. Compared to side-by-side views, this is an improvement that saves time as well as screen space. Our domain specialists in the fields of hydro-physics and weather have used ESS to recognize important features in their data sets that led to the refinements and improvements of their simulation code. As mentioned earlier, ESS was able to rapidly display both large-scale and subtle differences in hydro-data data by identifying unexpected behavior in the energy surface extracted from a RHIC gold nuclei collision simulation. Similarly, ESS provided a means for meteorologists studying fire burn simulation to verify the physical implications of varying parameters.

6.2 Limitations

To enable the comparative display of multiple surfaces in the same space, ESS uses more pre-attentive channels than standard side-by-side display. Beyond the luminance channel normally required to display surface shape, it uses the hue channel to distinguish between surfaces; this channel cannot then be used to display data attributes on the surface. Additionally, when animating ESS on time-varying data sets, the surfaces from the next time step are not rendered until all the invisible slices of a surface have been shown. Without prior understanding of this property of ESS, this slow transition could hinder user's visual detection of transient features.

There exists a trade-off between the number of ensemble data sets that can be displayed and the size of the features that can be perceived in regions where the ensemble surfaces differ (all features are preserved in regions where the ensemble data sets agree). If it is important that all features on slice that differs from its neighbors are seen, then ESS can be combined with an interactive toggle or magic lens technique that enables the user to select the display of just that ensemble member's surface in the region of difference. This incorporation of the magic lens technique may diminish the limitation of ESS slices having only two differing adjacent neighbors. Additionally, differences are only shown at slicing plane positions and even with shifting of slice planes and animation of slices, the problem still remains especially if there is high-frequency/high-amplitude difference present in members of the ensemble.

7. FUTURE WORK

We plan to immediately begin testing the ESS technique on a wider range of ensemble data, including yeast-mitotic-spindle data from microbiology and data from simulations of pediatric airway surgical plans from biomedicine. We also plan to carry out a full user study to quantitatively determine the speed and error rates for ESS task completion (e.g. identification and relationship tasks), and how they scale with the number of members in an ensemble.

ACKNOWLEDGMENTS

The authors wish to thank the collaborators on the MADAI project who have contributed time, support, and feedback on the visualizations. This work was supported by NSF's Cyber-Enabled Discovery and Innovation Program through grant no. 0941373, Sandia National Laboratories contract number 979573, and NIH grant no. 5-P41-EB002025.

REFERENCES

- [1] Healey, C. and Tateosian, L., "Perceptually-based brush strokes for nonphotorealistic visualization," *ACM Transactions on Graphics* **23**, 8–10 (January 2004).
- [2] Wolfe, J. K., N., and Dahlen, K., "Post-attentive vision," *Journal of Experimental Psychology: Human Perception and Performance* **26**(2), 693–716 (2000).

- [3] Healey, C. G., Booth, K. S., and Enns, J. T., “Visualizing real-time multivariate data using preattentive processing,” *ACM Trans. Model. Comput. Simul.* **5**, 190–221 (July 1995).
- [4] Vogel, E. K., Woodman, G. F., and Luck, S. J., “Storage of features, conjunctions, and objects in visual working memory,” *Journal of Experimental Psychology: Human Perception and Performance* **27**, 92–114 (January 2001).
- [5] Interrante, V., Fuchs, H., and Pizer, S. M., “Conveying the 3d shape of smoothly curving transparent surfaces via texture,” *IEEE Transactions on Visualization and Computer Graphics* **3**, 98–117 (1997).
- [6] Bauer-Kirpes, B., Schlegel, W., Boesecke, R., and Lorenz, W. J., “Display of organs and isodoses as shaded 3-d objects for 3-d therapy planning,” *International Journal of Radiation Oncology*Biophysics*Physics* **13**(1), 135 – 140 (1987).
- [7] Weigle, C. and II, R. M. T., “Visualizing intersecting surfaces with nested-surface techniques,” *Visualization Conference, IEEE* **0**, 64 (2005).
- [8] Rheingans, P., “Opacity-modulating triangular textures for irregular surfaces,” in [*Proceedings of the 7th conference on Visualization '96*], *VIS '96*, 219–ff., IEEE Computer Society Press, Los Alamitos, CA, USA (1996).
- [9] Bair, A., House, D., and Ware, C., “Perceptually optimizing textures for layered surfaces,” in [*Proceedings of the 2nd symposium on Applied perception in graphics and visualization*], *APGV '05*, 67–74, ACM, New York, NY, USA (2005).
- [10] Malik, M. M., Heinzl, C., and Groeller, M. E., “Comparative visualization for parameter studies of dataset series,” *IEEE Transactions on Visualization and Computer Graphics* **16**, 829–840 (September 2010).
- [11] Kirby, R. M., Marmanis, H., and Laidlaw, D. H., “Visualizing multivalued data from 2d incompressible flows using concepts from painting,” (1999).
- [12] Bokinsky, A. A., *Multivariate data visualization with data-driven spots*, PhD thesis (2003).
- [13] Laidlaw, D. H., Ahrens, E. T., Kremers, D., Avalos, M. J., Jacobs, R. E., and Readhead, C., “Visualizing diffusion tensor images of the mouse spinal cord,” (1998).
- [14] Urness, T., Interrante, V., Marusic, I., Longmire, E., and Ganapathisubramani, B., “Effectively visualizing multi-valued flow data using color and texture,” *Visualization Conference, IEEE* **0**, 16 (2003).
- [15] Pang, A. T., Wittenbrink, C. M., and Lodh, S. K., “Approaches to uncertainty visualization,” *The Visual Computer* **13**, 370–390 (1996).
- [16] Grigoryan, G., Rheingans, P., and Society, I. C., “Point-based probabilistic surfaces to show surface uncertainty,” *IEEE Transactions on Visualization and Computer Graphics* **10**, 2004 (2004).
- [17] Wittenbrink, C., Pang, A., and Lodha, S., “Verity visualization: Visual mappings,” tech. rep., University of California at Santa Cruz (1995).
- [18] Healey, C. G., “Choosing effective colours for data visualization,” in [*In Proceedings Visualization 96*], 263–270 (1996).
- [19] Bauer, B., Jolicoeur, P., and Cowan, W., “Visual search for colour targets that are or are not linearly separable from distractors,” *Vision Research* **36**, 1439–1466 (May 1996).
- [20] Vázquez, P.-P., Feixas, M., Sbert, M., and Heidrich, W., “Viewpoint selection using viewpoint entropy.,” in [*VMV'01*], 273–280 (2001).
- [21] Wyszecki, G. and Stiles, W. S., [*Concepts and Methods, Quantitative Data and Formulae, 2nd Edition*], John Wiley and Sons, Inc., New York, New York (1982).
- [22] Ware, C. and Beatty, J. C., “Using color dimension to display data dimensions,” *Human Factors* **30**(2), 127–142 (1988).
- [23] Hagh-Shenas, H., Interrante, V., Healey, C. G., and Kim, S., “Weaving versus blending: a quantitative assessment of the information carrying capacities of two alternative methods for conveying multivariate data with color,” in [*Proceedings of the 3rd symposium on Applied perception in graphics and visualization*], *APGV '06*, 164–164, ACM, New York, NY, USA (2006).
- [24] Miller, J. R., “Attribute blocks: Visualizing multiple continuously defined attributes,” *IEEE Computer Graphics and Applications* **27**, 57–69 (2007).
- [25] Ebert, D., [*Extending Visualization to Perceptualization: The Importance of Perception in Effective Communication of Information*], Elsevier Academic Press, Burlington, MA (2004).

- [26] Healey, C. G., “Formalizing artistic techniques and scientific visualization for painted renditions of complex information spaces,” in [*Proceedings of the 17th international joint conference on Artificial intelligence - Volume 1*], 371–376, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2001).
- [27] den Berg, R. V., Cornelissen, F. W., and Roerdink, J. B. T. M., “Perceptual dependencies in information visualization assessed by complex visual search,” *ACM Trans. Appl. Percept.* **4**, 3:1–3:21 (February 2008).
- [28] Sailer, S. L., Chaney, E. L., Rosenman, J. G., Sherouse, G. W., and Tepper, J. E., “Treatment planning at the university of north carolina at chapel hill,” *Seminars in Radiation Oncology* **2**, 267–273 (October 1992).
- [29] Gabrielides, N. C., i. Ginnis, A., and Kaklis, P. D., “Constructing smooth branching surfaces from cross sections,” in [*Proceedings of the 2006 ACM symposium on Solid and physical modeling*], *SPM '06*, 161–170, ACM, New York, NY, USA (2006).
- [30] Kijewski, P., “Three-dimensional treatment planning,” *Radiation, Oncology, Technology, and Biology* **10**, 10–33 (1994).
- [31] Gosset, N. and Chen, B., “Paint inspired color mixing and compositing for visualization,” in [*Proceedings of the IEEE Symposium on Information Visualization*], 113–118, IEEE Computer Society, Washington, DC, USA (2004).
- [32] Ion, A., Artner, N. M., Peyré, G., Kropatsch, W. G., and Cohen, L. D., “Matching 2d and 3d articulated shapes using the eccentricity transform,” *Comput. Vis. Image Underst.* **115**, 817–834 (June 2011).
- [33] Group, P. T. C. W., “Three-dimensional display in planning radiation therapy: A clinical perspective,” *International Journal of Radiation, Oncology, Biology and Physics* **21**, 78–89 (1991).
- [34] Taylor II, R. M., “Visualizing multiple scalar fields on the same surface,” *IEEE Computer Graphics and Applications* **22**, 6–10 (March-April 2002).
- [35] Stokking, R., Zubal, I. G., and Viergever, M. A., “Display of fused images: Methods, interpretation, and diagnostic improvements,” *Seminars in Nuclear Medicine* **33**(3), 218–227 (2003).
- [36] Sisneros, R., Johnson, C. R., and Huang, J., “Concurrent viewing of multiple attribute-specific subspaces,” *Computer Graphics Forum* **27**, 783–790 (January 2008).
- [37] Wolfe, J. K., N., and Dahlen, K., “Post-attentive vision,” *Journal of Experimental Psychology: Human Perception and Performance* **26**(2), 693–716 (2000).
- [38] Alred, G. J., Brusaw, C. T., and Oliu, W. E., [*Handbook of Technical Writing*], St. Martin’s, New York (2003 (seventh edition)).