Effectively Visualizing Multi-Valued Flow Data using Color and Texture

Timothy Urness[1], Victoria Interrante[1], Ivan Marusic[2], Ellen Longmire[2], and Bharathram Ganapathisubramani[2]

[1]Department of Computer Science and Engineering
[2]Department of Aerospace Engineering and Mechanics

University of Minnesota

{urness, interran}@cs.umn.edu; {marusic, ellen, bugs}@aem.umn.edu

Abstract

In this paper we offer several new insights and techniques for effectively using color and texture to simultaneously convey information about multiple 2D scalar and vector distributions, in a way that facilitates allowing each distribution to be understood both individually and in the context of one or more of the other distributions. Specifically, we introduce the concepts of:

- 'color weaving' for simultaneously representing information about multiple co-located color encoded distributions, and

- 'texture stitching' for achieving more spatially accurate multi-frequency line integral convolution representations of combined scalar and vector distributions.

The target application for our research is the definition, detection and visualization of regions of interest in a turbulent boundary layer flow at moderate Reynolds number. In this work, we examine and analyze streamwise-spanwise planes of three-component velocity vectors with the goal of identifying and characterizing spatially organized packets of hairpin vortices.

CR Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism.

Additional Keywords: flow visualization, line integral convolution, multi-variate data visualization, color, texture.

1. Introduction

Researchers in the turbulence community have long been interested in developing a deeper understanding of the key physical mechanisms in turbulent flows. Of particular interest is the question of how eddies contribute to drag, which, if understood, could enable the development of strategies to reduce drag by controlling the generation, scale and/or organization of eddies in the flow. This would have a major impact on many industries including aerospace, transportation, energy, and chemical processing. For example, reductions in turbulent skin-friction drag over aircraft would result in significant decreases in fuel usage – lowering fuel costs, increasing aircraft range and reducing pollutant and greenhouse gas emissions. Efforts to achieve a fundamental understanding of turbulent flows remain limited mainly because of a lack of understanding of the non-linear interactions that occur between the vortices that make up the motion. Since it is not computationally feasible to numerically simulate a wide range of eddy scales, researchers who are interested in probing these questions must rely on experimental methods to further their understanding.

Stereoscopic *particle image velocimetry* (PIV) is a technique that can be used to experimentally measure instantaneous components of a velocity field in a plane of a turbulent boundary layer in a moderate to high Reynolds number flow. Along with the experimentally generated vector field, values of vorticity, Reynolds shear stress, and swirl strength can be mathematically derived and are important in characterizing potential 'regions of interest'. In turbulent boundary layers, various theories have indicated that hairpin shaped vortices cause drag by producing Reynolds shear stress, and that this process may be enhanced when multiple hairpins travel together with similar speeds as a packet. In streamwise-spanwise planes parallel to the boundary layer surface, the packets can be characterized by zones of uniform but low streamwise velocity containing areas of high negative Reynolds shear stress and falling between cores of strong positive and negative vorticity [Ganapathisubramani et al. 2003]. Because the process of knowledge discovery related to this application is predicated on the ability to achieve an integrated understanding of the individual contributions of each variable and of how the variables inter-relate with each other, developing effective multivariate visualization methods is of critical importance to facilitating the understanding and analysis of results from the PIV experiments.

In this paper, we describe two techniques that we have recently developed for the purposes of more effectively visualizing multi-valued flow data. The first, and most significant, is a 'color weaving' algorithm that combines individual threads of multiple single-hue-encoded co-located scalar distributions, computed over a common vector field, to form a multi-colored line integral convolution tapestry in which multiple color combinations are represented explicitly via adjacent lines in the high resolution texture rather than through the compositing of multiple continuously defined color overlays. The second approach we describe involves the use of 'texture stitching' to achieve a more precise depiction of the spatial extents of discrete regions characterized by spatial frequency differences in a multi-frequency line integral convolution image.

Both approaches were applied to the visualization of PIV data obtained in a wind tunnel turbulent boundary layer at Reynolds number $R_\theta = 2500$ ($Re_\tau = \delta U_\tau/\nu = 1060$) at a wall-normal location nominally in the logarithmic region ($z^+ = zU_\tau/\nu = 98$). Here z is the distance normal to the wall, δ is the boundary layer thickness, U_τ is the wall shear velocity, and ν is the kinematic viscosity of the fluid. The vector fields used to generate the LIC textures were obtained by subtracting the mean streamwise velocity from the in-plane velocity components. Swirl is characterized by the quantity

$\lambda_{ci}\omega_z/|\omega_z|$, where λ_{ci} is the magnitude of the imaginary part of the eigenvalue of the local velocity gradient tensor and ω_z is the in-plane vorticity.

The paper is organized as follows. In section 2, we briefly discuss related previous work in the field of texture-based flow visualization and multivariate visualization. We then describe the 'color weaving' algorithm in section 3, and the texture stitching' algorithm in section 4. In section 5 we show results obtained using a combination of both techniques. Finally, we summarize and conclude in section 6.

2. Previous Work

2.1 Texture-Based Flow Visualization

Textures have traditionally been a popular and effective method for representing vector and scalar fields. In pioneering work, van Wijk [1991] introduced the concept of 'spot noise', a texture constructed from weighted and randomly positioned spots deformed in accordance with the direction of flow. Cabral and Leedom [1993] shortly afterward introduced *line integral convolution* (LIC), a versatile and widely-used technique in which intensities in an input texture are convolved along streamlines defined by an accompanying vector field to produce a texture that is highly correlated in the flow direction. Stalling and Hege [1995], in the Fast-LIC algorithm, achieved an order of magnitude increase in the efficiency of the LIC algorithm by taking advantage of coherence along streamlines. This results in the computation of the output texture being streamline oriented, not pixel oriented.

In traditional LIC images, the direction of movement in a flow is ambiguous, and animation is required to make that information explicit. However Wegenkittl et al. [1997] introduced a technique called Oriented Line Integral Convolution (OLIC) that addresses this issue. In essence, the OLIC algorithm works by taking as input a sparse texture resembling ink droplets on a page and using a ramp-like convolution kernel to smear the droplets according to the vector field, resulting in a collection of streaks in which intensity increases from tail to head. Computation time for that method was significantly reduced with the introduction of Fast Oriented Line Integral Convolution (FROLIC) [Wegenkittl and Groller 1997]. More recently, in another approach similar to OLIC, Sanna et al. [2001] propose a Thick Oriented Stream Lines (TOSL) method, in which the orientation of a flow is depicted by cyclically increasing the luminance along calculated streamlines.

Shen, Johnson, and Ma [1996] added color to LIC images through the use of simulated dye advection.

Multiple frequency input textures were first used with LIC by Kiu and Banks [1996], to incorporate indications of velocity magnitude. In their method, a single, composite input texture is constructed from multiple noise distributions, which are assigned to discrete regions defined by specific intervals of velocity magnitude, such that higher spatial frequency patterns are used in the higher velocity regions, and the LIC algorithm is run on this input. The length of the filter kernel, which defines the region of the input texture over which the convolution that determines the output intensity at each point is performed, is also varied in proportion to the magnitude of the velocity of the flow at the corresponding point.

Ware and Knight [1995] earlier proposed the use of Gabor functions to create texture-like images of flow data in which information is encoded along the perceptually significant texture dimensions of scale, orientation and contrast.

Research has also recently been done in representing scalar values with texture-based vector field representations such as LIC through the use of bump mapping [Sanna and Montrucchio 2000] and contrast enhancement [Sanna et al. 2002]. These techniques essentially enable the visualization of additional scalar values without requiring the use of color.

2.2 Multivariate Visualization

Healey and Enns [1998] have contributed methods to use texture elements on underlying 3D height fields to visualize multivariate data. They combine texture dimensions of height, density, and regularity along with perceptually uniform colors to increase the number of attributes that can be simultaneously represented.

Weigle *et. al* [2000] propose a texture generation technique, based on the layering of patches of oriented slivers, which uses orientation and luminance to encode information about multiple overlapping scalar fields.

Laidlaw et. al [1998] showed how shape, orientation, and color attributes of ellipsoids could be used to represent multivariate components in diffusion tensor images of the mouse spinal cord. In addition, they demonstrated a method for representing multi-valued data inspired by the brushing and layering techniques used in oil painting. Similarly inspired by concepts from painting, Kirby, Marmanis and Laidlaw [1999] showed how different sized icons, color, elongated ellipses, and layering could be used to portray multivariate data from 2D compressible flows.

3. Color Weaving

With few exceptions, the use of color with LIC has traditionally been limited to the simplest of color compositing operations in which a LIC texture image is in effect overlaid with a single continuous semitransparent color wash image, with the resulting effect that blacks are left black and the whites are shifted toward the specified hue at each point. While effective for conveying a single scalar distribution in the context of the flow, this post-process method does not allow for the effective simultaneous representation of multiple scalar fields, due to the perceptual difficulty of and inherent ambiguity in color decomposition (figure 1).

As an alternative, we propose a technique in which multiple colors are allowed to coexist on neighboring streamlines, resulting in multicolored images that resemble a tapestry woven with different colored threads (figure 2).

We begin by selecting several highly saturated and perceptually iso-luminant colors using the technique suggested by Kindlmann et al. [2002]. As luminance plays a primary role in how features are perceived [Ware 2000], selecting base colors that are as perceptually uniform as possible helps to achieve a final image in which similar concentrations are represented with reasonably equivalent prominence across the multiple distributions. Additionally, selecting colors that are relatively equivalently discriminable [Healey 1996] reduces the potential for ambiguous or misleading representations. We then use each base color to define a corresponding two-dimensional colormap, in which saturation increases along the horizontal axis and value increases in the vertical direction (figure 4). Each scalar distribution is associated with a unique colormap.
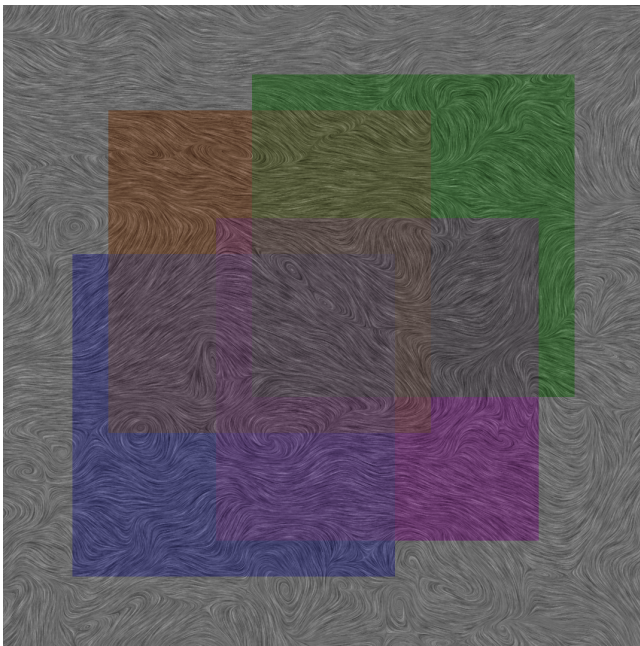
Figure 1: Four artificially defined, mutually overlapping regions, overlaid on a LIC image. The color combinations in the overlap regions are obtained by averaging in RGB colorspace.
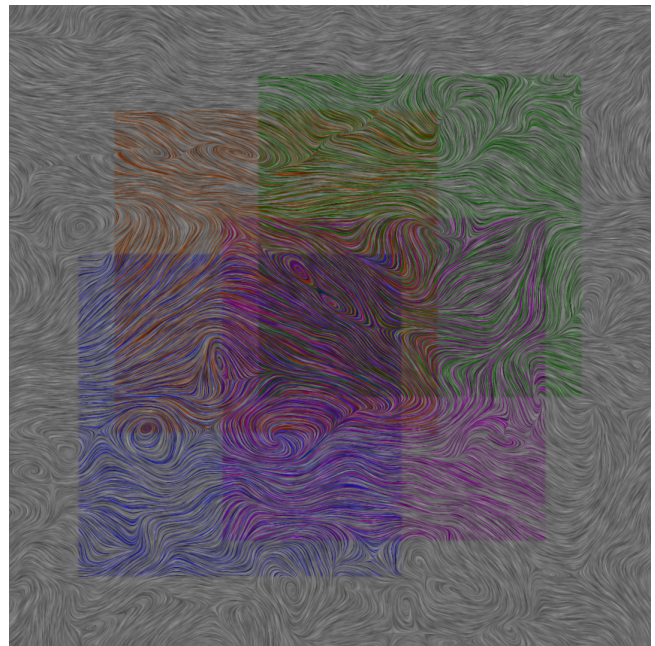

Figure 2: The same four regions, represented across the same LIC image via *color weaving*. Note the continuity of color along individual streamlines within each region, and the ability to accurately perceive combinations of component colors in the areas of high overlap (characterized by the presence of three or more layers).
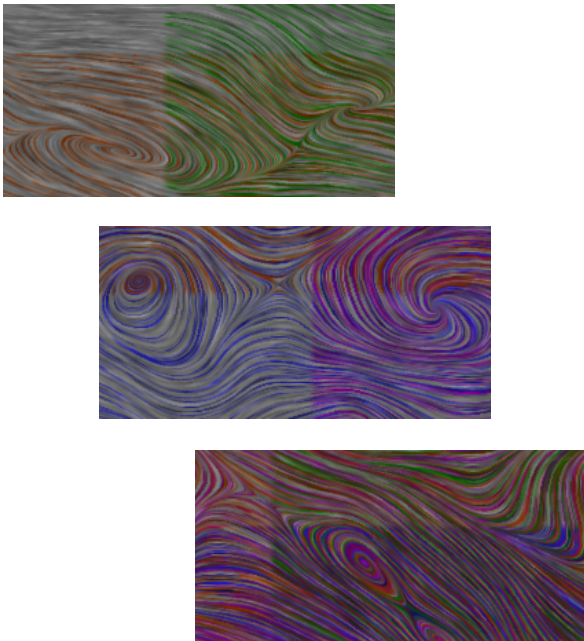
Figure 3: Three close-up excerpts from the overlap image shown in figure 2. Region combinations represented in each excerpt are: Top: background, orange only, green only and orange + green overlap. Middle: blue only, blue + orange overlap, blue + magenta overlap, and blue + orange + magenta overlap. Bottom: orange + magenta everywhere, partially overlapped by blue, partially overlapped by green, and partially overlapped by both blue and green.
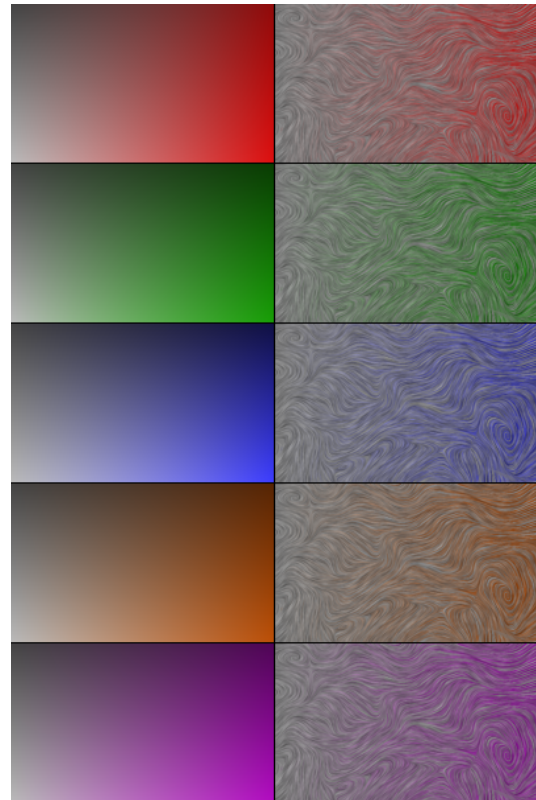

Figure 4: Left: A suite of two-dimensional colormaps. Right: the results of using each colormap to represent, over the same LIC texture, a simple scalar distribution that is increasing in value from left to right.
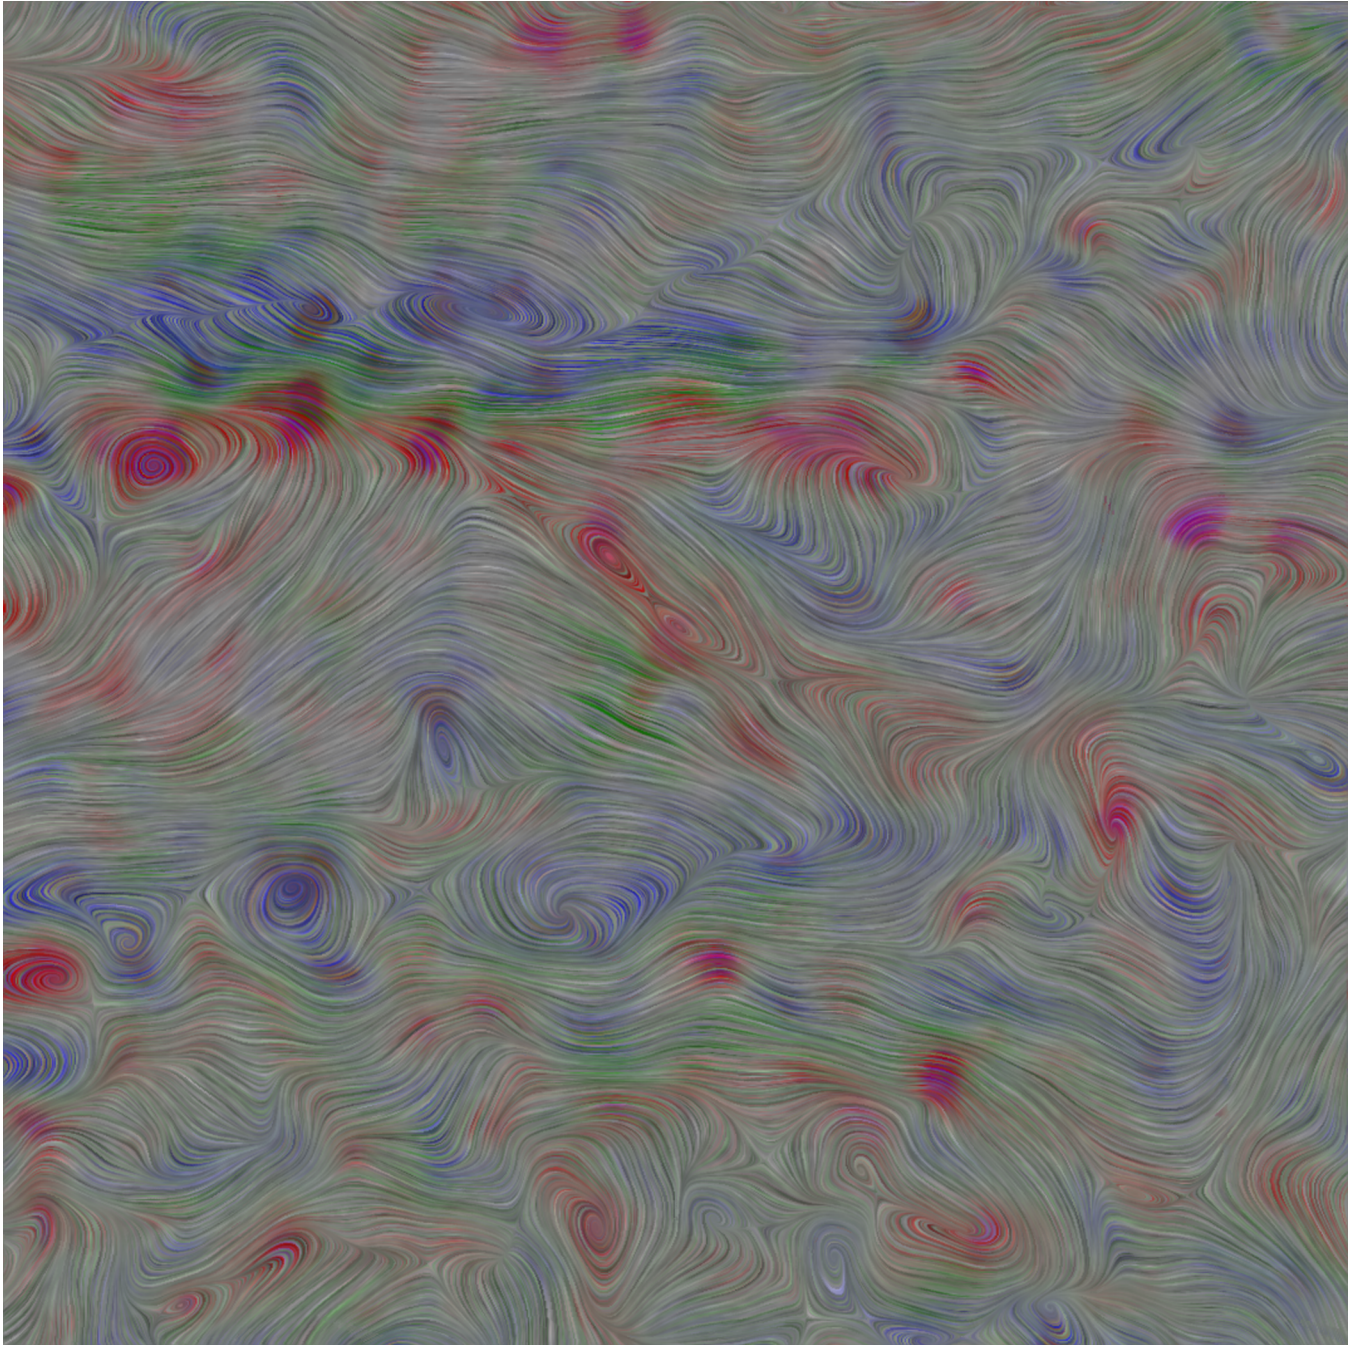
Figure 5: A composite 'color woven' image of an experimentally acquired PIV dataset in which we simultaneously highlight areas of significant positive vorticity (red), negative vorticity (blue), strongly negative Reynolds shear stress (green), and high swirl strength (orange or magenta, depending on the direction of the swirl).

We introduce color on a streamline-by-streamline basis during the computation of the LIC image. At each pixel along a streamline we define the final image color using a 3D color table lookup. The first index defines the hue, or the choice of which 2D colormap to perform the lookup in. Within each 2D colormap, the index for the value component is defined by the grey value obtained from running the LIC computation, and the index for the saturation component is defined by the magnitude of the value of the particular scalar distribution being represented at that point. The result is to preserve the luminance pattern established by the LIC while defining the colors to be fully saturated at points where the scalar variables reaches their maxima, fading to the default LIC value as the magnitudes of the scalar variables decrease.

The streamline based fast-LIC algorithm [Stalling and Hege 1995] is critical to our implementation as the hue index is only incremented when a new streamline is calculated. In our latest implementation, the output image is created in a single pass. However special steps have to be taken to prevent color mixing, while preserving the anti-aliasing effects that are achieved when multiple streamlines are allowed to pass through each pixel. We accomplish this by recording, for each pixel in the output image, the hue index of the first streamline that was used to determine its color. When subsequent streamlines pass through the same pixel, the original hue index is used in the colortable lookup, so that only the luminance components from the multiple streamlines are blended when the results are averaged.

In areas characterized by the presence of prominent values in multiple distributions, alternate colors are visible along adjacent streamlines. We use a sparse and a consistent mapping of individual colors to individual streamlines in order to maintain continuity in the representation of each distribution and to ensure that the apparent concentration of each color remains in consistent proportion to the magnitude of the corresponding scalar distribution, regardless of the concentrations of the other scalar distributions at that point.

The success of our method depends on being able to represent multiple different streamlines across each 'point' in the multivariate distribution being portrayed. To achieve best results, it is generally necessary to up-sample the input data. For the images in this paper, a 1071x1071 input texture was used, and the maximum streamline length was 380.

4. Texture Stitching

Figure 6 illustrates the classical problem with attempting to apply a color wash to an input texture, before running LIC, to indicate the distribution of values in a scalar field associated with the vector data: the effect of the LIC is to smear out the colors, distorting the appearance of the scalar distribution in the final image and impeding efforts to accurately interpret the value of the distribution from the value of the color at any particular point. For this reason, color encoding is universally applied post-LIC, unless it is explicitly desired to use the color to demonstrate the effects of advection. Being aware of these issues with respect to the use of color, and wishing to use spatial frequency to encode the presence of discrete regions of interest in our data, we sought to develop 'texture stitching' – a post-LIC variant of the pre-LIC multi-frequency method proposed by Kiu and Banks[1996] in which it would be possible to preserve the fidelity of region boundaries implicitly indicated by spatial frequency differences in the texture pattern in the final image, while avoiding the introduction of unnecessary discontinuity artifacts.
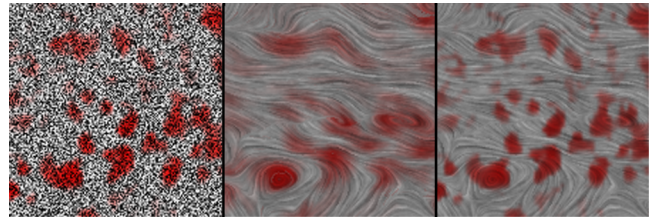

Figure 6: An illustration of the problem with trying to use color to indicate regions of interest pre-LIC. Left: color wash applied to the input texture. Middle: results after running LIC – the region definition is not well preserved. Right: results of applying the color wash post-LIC. The goal of 'texture stitching' is to achieve the latter effect with multi-frequency texture patterns.

Following Kiu and Banks [1996], the first step in our approach is to construct a set of correlated noise texture images by low pass filtering an initial high frequency noise pattern and equalizing the intensity histograms of the results to the intensity histogram of the original. For our application we were primarily interested in using spatial frequency to indicate the locations of computed 'regions of interest' within a larger surrounding flow field. Thus we only needed to generate two noise texture patterns (high and low). To create the images in this paper, we applied a Gaussian filter of width 20 and standard deviation 2.0 to the white noise shown in figure 7(left) to achieve figure 7(right).
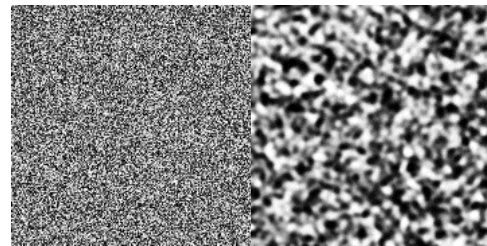

Figure 7: Samples from input textures used in our texture stitching technique. Left: the high frequency noise input texture. Right: the low frequency pattern achieved after Gaussian blurring and histogram equalization.

There is a direct correlation between the size differences of the spots in the two input textures and the filter kernel length differences that are required to achieve output textures that will appear to differ by only a uniform (isotropic) scaling factor. Although we did not intend to attempt to use filter kernel length to encode any information about the flow, we determined that we would need to use a larger filter kernel length with the low frequency input texture in order to make the lines in the low frequency output texture appear to have the same length-to-width ratio as the lines in the corresponding high frequency pattern. Since the lower frequency lines are less effective at conveying details of the flow orientation, we decided to use the low frequency texture to demarcate the regions of interest, which are characterized by uniform momentum and low velocity.

We proceeded by using the high and low frequency noise input textures to create two separate LIC images. We also created a binary mask corresponding to the results of our trial region detection algorithm [Ganapathisubramani et al. 2003] (one of the goals of the visualization effort was to determine the suitability of the results produced by our region detection method and possibly to provide insight into how it might be refined to achieve greater effectiveness). We used the binary mask to composite the results

post LIC. The results of our method are shown in figure 8, top right. Results obtained using the original multifrequency LIC method are shown in figure 8, top left, for comparison.

One issue that arose was the question of whether it might be desirable, or not, to minimize the incidence of contrast differences between the low frequency and high frequency texture regions. Contrast will inevitably be lower for LIC images obtained from higher frequency input patterns, unless there is a huge reduction in filter kernel lengths, because more different grey values will be averaged together, bringing the result closer to the mean than in the case of the low frequency pattern. Retaining the ability to equalize contrast, which can easily be done in the texture stitching approach, reserves the potential to use contrast differences to encode a different scalar distribution. Figure 8, bottom right, shows the results of performing texture stitching without contrast equalization. Here the region differentiation becomes more prominent. However, the general continuity of light and dark patterning remains consistent between the regions, which would not be the case if unrelated input textures were used (figure 8, bottom left).
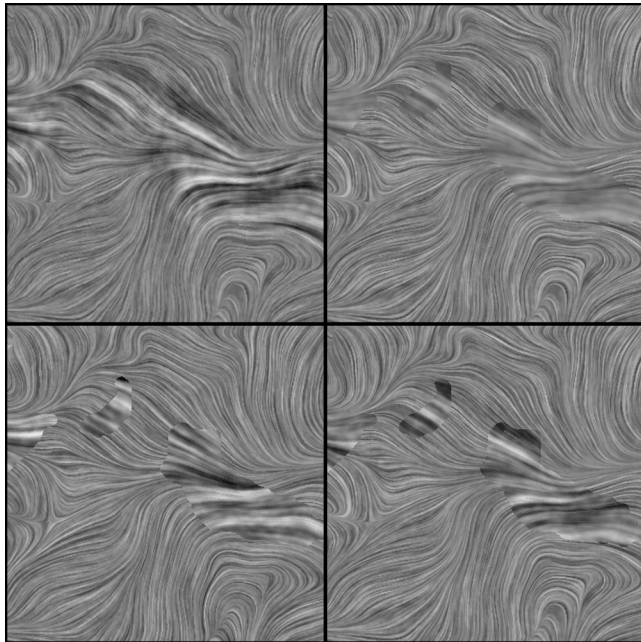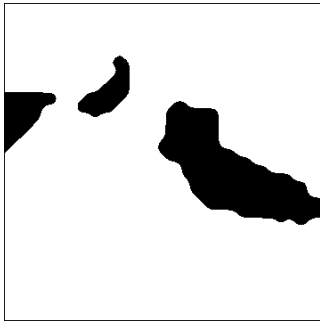

Figure 8: Top: the region of interest mask. Bottom: images are obtained, clockwise from upper left, using: the Kiu/Banks algorithm; texture stitching with histogram equalization; texture stitching plus contrast enhancement of low frequency regions; texture stitching using unrelated input patterns.

The main drawback of texture stitching, compared to the multi-frequency LIC approach taken by Kiu and Banks, is that it allows region boundaries to be noticeable in the final image. Hence the texture stitching approach will not be suitable for applications in which one hopes to approximate a continuous series by a finite set of different spatial frequency patterns, which was the target application for Kiu and Banks.

5. Combination of Techniques

The color weaving and texture stitching techniques can be used in combination to visualize multiple scalar distributions across a flow field and emphasize regions of interest in the data. The first step in the process of merging the two techniques is to create high and low frequency color weave component images using the different spatial frequency input patterns.

Using a low spatial frequency input texture does not change the resolution of the streamlines calculated. If the same technique is used for hue index selection in the low and high frequency texture cases, the result will be an appearance of high frequency color changes within a lower frequency luminance-dominated pattern, as seen in figure 9. The low frequency color weave image can also be combined via texture stitching, as shown in figure 10, with the high frequency image color weave image computed earlier, using the region of interest mask partially shown in figure 8.
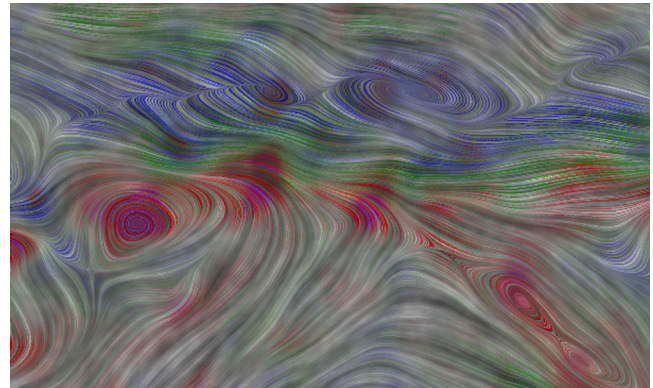

Figure 9: A close-up view of a portion of a color weave image created using a uniformly low frequency input pattern.
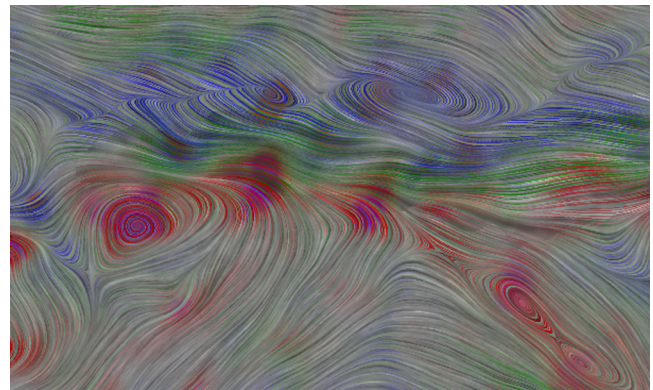

Figure 10: The same excerpt, created using a combination of color weaving and texture stitching. The low spatial frequency pattern indicates the probable extent of a coherent packet of hairpin vortices.

6. Summary and Discussion

We have presented two techniques that can facilitate the effective visualization of multi-valued flow data using color and texture.

'Color weaving' provides an alternative to traditional color compositing by allowing multiple colors to be closely interwoven, via the assignment of distinct separate hues to individual streamlines, rather than blended. In order to allow each color to encode multiple values in a continuous distribution, we let the saturation of the color at each point vary according to the value in the corresponding scalar distribution. We achieve consistent combinations of colors by assigning color indices to streamlines in an alternating manner that depends on the order in which they are encountered in a deterministic walk through the pixel grid. The result is a multicolored LIC image that resembles a tapestry woven with different colored threads.

'Texture stitching' allows preservation of faithful region boundaries in multi-frequency LIC through the use of post-LIC merging of selected adjacent regions. We first obtain separate LIC textures based on correlated high and low frequency noise input patterns, then combine the results using a binary mask to force adherence to pre-defined boundary curves.

Both of these methods were developed to facilitate simultaneous visualization of multiple derived quantities in experimentally acquired, stereo PIV data of wall-bounded turbulent flow at moderately high Reynolds numbers. However, both methods would also prove useful in visualizing PIV or stereo PIV data in a variety of turbulent or other three-dimensional flows. Our ultimate goal through this work is to enable researchers to obtain a succinct, meaningful visual summary of the contents of a dataset through providing techniques that allow the creation of images in which the important features of multiple scalar distributions can be understood both independently and in the context of multiple other distributions. Color weaving provides a method for accurate display of derived quantities that may overlap — for example, wall-normal vorticity and Reynolds shear stress, or wall-normal vorticity and swirl strength, as shown in Figure 5. The method thus allows for accurate interpretation of the flow structure in the overlapping regions, e.g. correlation of out-of-plane vorticity and velocity or identification of vortex cores within shearing zones. Texture stitching is important in accurately displaying the extent of specific regions identified within a flow field, particularly in zones where velocity gradients may be strong. In the current paper, this effect was employed to delineate the boundaries of a vortex packet identified with a specific feature extraction algorithm [see Ganapathisubramani et al. 2003]. In other studies, it could be employed to accentuate individual vortices, streaming or straining zones, etc.

In the future, we would like to extend our texture stitching approach beyond LIC to the case of multiple different texture types, finessing the problem of maintaining continuity between adjacent texture patches through the use of novel texture synthesis methods that operate on the basis of multiple considerations in the flow.

7. Acknowledgments

References

CABRAL, B. AND LEEDOM, C. 1993. Imaging Vector Fields Using Line Integral Convolution, In *Proceedings of ACM SIGGRAPH 93*, ACM Press/ACM SIGGRAPH, New York, James T. Kajiya, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 263–269.

GANAPATHISUBRAMANI, B. LONGMIRE, E. AND MARUSIC, I. 2003. Characteristics of Vortex Packets in a Turbulent Boundary Layer, *Journal of Fluid Mechanics*, *478*, 35–46.

HEALEY, C. 1996. Choosing Effective Colours for Data Visualization, In *Proceedings of IEEE Visualization '96*, 263–270.

HEALEY, C. AND ENNS, J. 1998. Building Perceptual Textures to Visualize Multidimensional Datasets", In *Proceedings of IEEE Visualization '98*, 111–118.

KINDLMANN, G., REINHARD, E. AND CREEM, S. 2002. Face-based Luminance Matching for Perceptual Colormap Generation, In *Proceedings of IEEE Visualization 2002*, 309–406.

KIRBY, R., MARMANIS, H. AND LAIDLAW, D.H. 1999. Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting, In *Proceedings of IEEE Visualization '99*, 333–340.

KIU, M.-H. AND BANKS, D. 1996. Multi-Frequency Noise for LIC", In *Proceedings of IEEE Visualization '96*, 121–126.

LAIDLAW, D., AHRENS, E., KREMERS, D., AVALOS, M., JACOBS, R. AND READHEAD, C. Visualizing Diffusion Tensor Images of the Mouse Spinal Cord, In *Proceedings of IEEE Visualization '98*, 127–134.

SANNA, A. AND MONTRUCCHIO, B. 2000. Adding a Scalar Value to 2D Vector Field Visualization: the BLIC (Bumped LIC), In *Eurographics 2000 Short Presentations Proceedings*, 119–124.

SANNA, A., MONTRUCCHIO, B., MONTUSCHI, P. AND SPARAVIGNA, A. 2001. Visualizing vector fields: the thick oriented stream-line algorithm (TOSL), *Computers and Graphics*, *25*, 5, 847–855.

SANNA, A. MONTRUCCHIO, B. ZUNINO, C. AND MONTUSCHI, P. 2002. Enhanced Vector Field Visualization by Local Contrast Analysis, In *Eurographics/IEEE TCVG Symposium on Data Visualization 2002*, 35–41

SHEN, H.-W., JOHNSON, C.R. AND MA, K.-L. 1996. Visualizing Vector Fields Using Line Integral Convolution and Dye Advection, In *Proceedings of the ACM Symposium on Volume Visualization '96*, 63–70.

STALLING, D. AND HEGE, H.-C. 1995. Fast and Resolution-Independent Line Integral Convolution, In *Proceedings of ACM SIGGRAPH 95*, ACM Press/ACM SIGGRAPH, New York, Robert Cook, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 249–256.

WARE, C. AND KNIGHT, W. 1995. Using Visual Texture for Information Display, *ACM Transactions on Graphics*, *14*, 1, 3–20.

WARE, C. 2000. *Information Visualization: Perception for Design,* Morgan Kaufman.

WEGENKITTL, R., GROLLER, E. AND PURGATHOFER, W. 1997. Animating Flowfields: Rendering of Oriented Line Integral Convolution, In *Proceedings of IEEE Computer Animation '97*, 15–21.

WEGENKITTL, R. AND GROLLER, E. 1997. Oriented Line Integral Convolution for Vector Field Visualization via the Internet, In *Proceedings of IEEE Visualization '97*, 309–316.

WEIGLE, C., EMIGH, W., LIU, G., TAYLOR, R., ENNS, J. AND HEALEY, C. 2000. Oriented Sliver Textures: A technique for local value estimation of multiple scalar fields, In *Proceedings of Graphics Interface 2000*, 153–162.

VAN WIJK, J. 1991. Spot Noise – Texture Syntheses for Data Visualization, In *Computer Graphics (Proceedings of ACM SIGGRAPH 91), 25,* 4, ACM, 309–318.