

WEAVE: A System for Visually Linking 3-D and Statistical Visualizations, Applied to Cardiac Simulation and Measurement Data

D.L. Gresh and B.E. Rogowitz*
IBM T.J. Watson Research Center

R.L. Winslow, D.F. Scollan, and C.K. Yung †
Department of Biomedical Engineering, Johns Hopkins University School of Medicine

Abstract

WEAVE (Workbench Environment for Analysis and Visual Exploration) is an environment for creating interactive visualization applications. WEAVE differs from previous systems in that it provides transparent linking between custom 3-D visualizations and multidimensional statistical representations, and provides interactive color brushing between all visualizations.

In this paper, we demonstrate how WEAVE can be used to rapidly prototype a biomedical application, weaving together simulation data, measurement data, and 3-d anatomical data concerning the propagation of excitation in the heart. These linked statistical and custom three-dimensional visualizations of the heart can allow scientists to more effectively study the correspondance of structure and behavior.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

Keywords: visualization, medical, heart, data synthesis

1 Introduction

Several different types of visualization systems have been developed to handle different types of data. For example, scientific visualization systems have been developed for data which have an inherent spatial frame of reference, such as simulation or measurements within a volume of space. They are excellent for demonstrating patterns, correlations and structures in three-dimensional space. They also excel at representing geometric relationships extracted from the data, such as streamlines, isosurfaces, cutting planes, volume renderings, and height displacement representations. These systems (*e.g.*, [UFK⁺89], [LAC⁺92], [WS97]) allow the user to create custom visualizations for a particular data set or domain application. However, scientific visualization toolkits do not generally support quantitative comparison across multiple variables with dynamic brushing and linking.

*gresh, rogowitz@us.ibm.com

†rwinslow, dscollan, cyung@bme.jhu.edu

© 0-7803-6478-3/00/\$10.00 2000 IEEE

Interactive statistical visualization packages have been developed for tabular data. They typically provide a library of standard views, such as scatterplots, parallel coordinates, and line plots. The most sophisticated of these systems also provide mechanisms which allow the user to use color to interactively mark subsets of the data and interactively interrogate the data, quickly observing correlations between multiple variables (*e.g.* [BCW87], [BMMS91], [ID91],[Rab94], [RC94], [AW95]). However this class of statistical visualization tools is appropriate for strictly tabular or hierarchical data, and is not well suited to visualizations of spatial data, which includes not only measurements at a number of points in space, but also meta data describing the context of those points, that is, the connectivity between them, and the correct interpretation that is to be made of them.

Previous Multi-Component Systems

The value of multiple, coordinated views is well-recognized in the field of visualization [RRGK96], [WB97]. The ability to see relationships in one presentation, mark regions of interest, and dynamically see the marked regions in other presentations can lead to insight not possible with single views or with multiple, uncoordinated views. In [NS00], North and Shneiderman discuss "Snap-Together Visualization," which has some similarities to the work we describe here. In particular, they are interested in the easy creation of coordinated views of data. Users wishing to create a custom environment choose from a menu of visualizations and indicate the way in which views should be coordinated. Their focus is on a user who desires to create a *custom environment* from a set of *existing* visualization components, by specifying the presentations to display and the linkage between them. In Visage, described in [RLS⁺96], automatic brushing and linkage is implemented through the use of a common underlying data structure, and investigation of particular data items of interest is performed through "drag-and-drop" operations on atomic items of data in an object-oriented manner. Customization of presentation "panes" is possible through an exposed scripting language. Both of these emphasize the desirability of allowing the user a variety of choices for presentation styles. However both focus primarily on the set of typical presentations for "information" visualization; that is, visualization of data as might be found in databases or spreadsheets. Our focus is on creating *custom, complex, three-dimensional components* using non-programming methods, linked in a custom environment to a set of powerful statistical presentations.

The WEAVE System

The WEAVE system allows seamless integration of custom three-dimensional views based on flexible and expressive data models, with sophisticated statistical data analysis and presentation tools.

The three-dimensional views supported are completely customizable and rapidly modifiable using a visual programming environment. They are created using Data Explorer [LAC⁺92] (DX), which supports the creation of a wide range of possible visualizations, in an easy-to-modify environment. The statistical presentations are based on the Diamond technology [Rab94] which includes over 20 different representations of data, which are all automatically linked with one another. The presentations range from common plots such as histograms, pie charts and scatterplots to sophisticated presentations representing correlations between sets of variables, clustering of cases, and views of very high dimensionality data.

2 Modeling the Heart

The data considered here comprise two separate data sets developed at the Center for Computational Medicine and Biology in the Department of Biomedical Engineering at Johns Hopkins University. One is a finite element simulation of heart excitation. The second is a set of muscle fiber tensors inferred from imaging data. The simulation is a three-dimensional computational model of the heart (the canine ventricular cell model [WSH⁺00]), whose purpose is to investigate the basis of arrhythmia (anomalous rhythms) in heart disease. The myocardium (muscle tissue of the heart) is modeled as a set of nodes, with each node representing a volume of tissue centered about a particular point in space. A set of 32 coupled nonlinear ordinary differential equations, defining 13 distinct membrane currents [WRJ⁺99], are integrated in time at each node point to simulate cardiac activation and repolarization. For purposes of this analysis, ten different measures of cardiac activity have been selected for analysis and display at each node. These measures (and their names in italics for the purposes of the visualization presentations) are as follows: *Voltage*, or membrane potential, is the voltage difference between the inside and outside of the cell. When cardiac cells become excited (an event referred to as an action potential), this difference becomes large and positive. Membrane current (*Im*) corresponds to the current that flows between the inside and the outside of a cell. Spatial current (*Ispace*) at a point in space is defined as the total outward current from the cell at that point into all of the cells neighboring that cell. As membrane potential evolves in time, cells become excited, producing large spatial currents that will excite adjacent cells that are in front of the travelling wave.

The simulation also models the calcium and potassium ion levels in heart muscle cells. These concentrations control many cellular functions, including force generation. Every muscle fiber cell has a compartment which contains the NSR (network sarcoplasmic reticulum) and the JSR (junction sarcoplasmic reticulum). The concentration of calcium in the NSR is *CaNSR*; the concentration in the JSR is *CaJSR*. Following an action potential (the depolarizing excitatory event which triggers contraction in cardiac muscle fibers) calcium is pumped from the cytosol (the fluid portion of the cell's cytoplasm) into the NSR and is made available for release from the JSR for the next action potential. Calcium released into the space between the NSR and the cell membrane, the diadic space, is called *CaSS*. It then diffuses into the cytosol, where its concentration (*CaI*) drives muscle contraction.

Behavior of a number of potassium currents is also simulated. These currents are important for restoring the membrane potential of cardiac cells to the resting value following an action potential.

The data consist of 60 time steps following the travelling wave of excitation of the heart through the activation phase. Each time step (spaced 1 msec apart) defines the above variables at a set of approximately 110,000 node points.

The second data set consists of measurements of the diffusion tensors through magnetic resonance (MR) imaging in fixed heart tissue. It has been previously shown that the primary eigenvectors

of these diffusion tensors are aligned with the long axis of cardiac fibers [SHWF99]. Fiber orientation can be deduced from diffusion tensor imaging data as follows: In MR images, diffusion of water along the direction of an applied magnetic field gradient causes signal attenuation. Because diffusion is fastest down the long axis of a cylindrical muscle fiber, signal attenuation is greatest when the applied magnetic field gradient is along the fiber direction. By applying magnetic gradients in at least 6 noncollinear directions during image acquisition, the three-dimensional fiber orientation can be deduced at all voxels within the MR image by examining the directional dependence of attenuation. This orientation is embodied in the 6 unique components of the 3x3 symmetric diffusion tensor. We use the primary eigenvector to construct a flow field vector at each point in the three-dimensional field. Streamlines through this flow field can then represent the fiber structure. We also implemented the "tensorline" method as described by [WKL99]. This method utilizes all three eigenvectors along with a measure of the local linear anisotropy, and reduces to the streamline case for purely linearly anisotropic tensor fields. Our application allows either method; the figures reproduced here use tensorlines.

3 Implementation

WEAVE is implemented using ActiveX component technology, and the applications we have created thus far use a Visual Basic development environment. We created ActiveX components which wrapped Diamond DLLs for a number of different statistical views and linked them with an ActiveX component which wrapped a custom Data Explorer visual program. We developed special linkage and notification modules to bridge these two technologies to allow automatic interactive brushing between all components. We also developed interaction and brushing code for the three-dimensional DX views. The result is that the user can brush in any statistical presentation and see the colors reflected not only in the other statistical presentations but also in the 3-D view. Similarly, the user can brush regions in the 3-D view and see the colors reflected in the statistical presentations.

In WEAVE all components run in the same process space, so sharing of data is possible between the various presentations, which is a significant benefit for large data sets. All objects which share a data instance are automatically linked by the underlying Diamond infrastructure, so that no explicit linking is necessary other than indicating a data source (instance) for each presentation.

In order to create a linkable custom three-dimensional presentation, the visualization programmer can create a Data Explorer visualization program using its standard visual programming environment. We created a new image tool in Data Explorer which renders the image and allows user interaction. We also created a component which serves to identify objects in the rendered image (vertices, markers, posts, etc.) with specific indices from the tabular data which is to be linked with the geometric data. The rendered object can be of any desired complexity, including a number of separate objects, some of which may be associated with the tabular data, and some which are not. Brushing between the statistical presentations and the geometric presentation will automatically occur only for those parts of the geometric image containing the marking component. The designer of the visualization program decides what sort of mapping to make from the tabular data to the geometric data. In the application discussed here, we use four different mappings as described later in this section.

We used WEAVE to develop a custom application for the cardiac simulation and measurement data sets. We chose to offer a two-dimensional scatterplot, a histogram, and a parallel coordinate presentation of the simulation data linked with one of four complementary geometric views with a variety of viewing and interaction options. These geometric views were created specifically

for this application to enable the scientists to visualize particular features. One geometric presentation constructs a bounding surface of the heart and shares colors at each vertex with the statistical presentations, which contain data for every node in the myocardium. This presentation is particularly valuable for following the progression of the excitation wave across the epicardium. The second presentation shows a marking point at each vertex in the volume. This allows the scientist to see the three-dimensional structure of the marked regions. In order to see the marked volumetric regions within the context of the heart, the third presentation includes the boundary contour, which may be painted using a perceptually appropriate colormap [RT98] to represent any simulation variable. The fourth representation shows tensorlines representing the fiber structure of the heart, alone or again in conjunction with the surface of the heart. This is valuable because it allows the scientist to study the role of the heart fiber geometry in the excitation process. Because the tensorlines do not share node positions with the simulation data, we developed a general purpose tool to enable association of index values with each point on the tensorlines. We choose to use the nearest node point to a particular point on the streamline for the association with the statistical presentations.

4 Results Using the WEAVE Environment to Analyze Cardiac Data

We begin our discussion of the analysis of the cardiac simulation data with the pair of images in Figure 1, which show a geometric rendering of the boundary of the myocardium, with one of the simulation variables (in this case voltage) painted on the surface using a standard rainbow color map on the top and a perceptual segmented colormap on the bottom. These are data from 46 msec after the beginning of the excitation wave. The top of the heart is not part of the simulation and so appears to be cut away. The left ventricle is the cavity toward the front of the image, the right ventricle is the cavity toward the rear. These are typical scientific visualization-style presentations, showing one variable at a time. They allow little in the way of quantitative analysis and comparison of variables.

Figure 2 shows the prototype application interface which we created using WEAVE. The data displayed are from time step 10, which is 10 msec after the beginning of the excitation wave. A step interactor allows the investigation of other time steps of the data. The top left image is one of the three three-dimensional representations of the heart we provide. The top right image is a two-dimensional scatterplot, here showing spatial current as a function of voltage. The bottom image is a histogram plot, displaying a histogram of spatial current. Using the two-dimensional scatterplot view, a region of high spatial current has been interactively marked in red, and a region of high voltage has been marked in green and the resulting colored regions have been dynamically and automatically brushed to the other two presentations. (Yellow represents regions which have high spatial current AND high voltage) The geometric representation clearly shows the high voltage regions (in green) surrounded by the high spatial current regions in red, indicating the progression of the excitation wave: high spatial current (red) is exciting sets of cells which in turn become depolarized (green). If desired, a separate window can be brought up displaying a parallel coordinates view of the same data. Figure 3 shows such a presentation using the same coloring as in Figure 2. Figure 4 demonstrates the option of displaying only the red (high spatial current) points. Looking across this constellation of variables, we can see, for example, that Potassium currents *IKs* and *ItoI* are relatively low, reflecting the fact that the cell is not yet repolarizing.

The four images in Figure 5 show a few of the alternate presentations provided to the user. All four renderings are colored under the same criteria as in Figure 2. In the first image, each point in the

simulation is shown as a point in the rendered image; we have chosen here to show *only* the red points. Because the three-dimensional image is interactively rotatable, an understanding of the volumetric shape of the colored region is easily achieved. The cloud of points can also be shown together with the heart surface, which may be optionally be colored, as shown in the second image, which colors the intracellular calcium *Ca_i* using a segmented color map. The third image shows the tensorlines computed as described in Section 2 mapped with the same colors. We also provide an option, not shown, to code the tensorlines by a measure of the confidence of the measurement; that is, by the linear anisotropy coefficient. The fourth image shows the heart boundary representation as in Figure 2.

Figure 6 shows the results of an investigation of the relationship between intracellular calcium (*Ca_i*) and calcium in the subspace (*Ca_{SS}*). Before the use of WEAVE in this application, the expectation of the researchers using this application was that calcium in the subspace would show two populations: essentially zero and some finite amount. In fact, a much more complex structure is present, with *three* populations: a population with no calcium in the subspace, and two populations having finite amounts of calcium in the subspace. Marking the two positively-valued populations with different colors, red and green, shows clearly that they overlap one another in anatomical distribution.

5 Discussion

In Section 4 we discussed the advantages of WEAVE for to a specific set of data in understanding the role of structure and function in a cardiac application. Here we discuss the specific advantages of this sort of environment as applicable to a broad range of application areas.

A variety of different visualization methods have been developed for analyzing and interacting with data. Some tools have particular strength in highlighting statistical features in multidimensional data and allowing the user to quickly compare and correlate variables. Others have strength in allowing custom, three-dimensional, realistic views of spatial data. The WEAVE environment allows an application developer to incorporate the strengths of both styles of visualization into a common, custom environment. The three-dimensional views in this environment can be created quickly (in an hour or two) and immediately incorporated into a prototype application for testing their suitability and appropriateness for the problem under consideration. All control over the characteristics and parameters to modify the three-dimensional visualization is accomplished through modification of the visual program, and no recompilation of the ActiveX control is necessary at any time. At the same time, dynamic, automatic brushing can be incorporated into the three-dimensional visualization with no special knowledge or awareness by the visualization designer. Programmers with no knowledge of C++ or other high level programming languages can create complex visualization components for custom applications, and have them automatically linked to a wide variety of statistical presentations such as histograms, 2 and 3-D scatterplots, dendograms, parallel coordinates, etc.

We have demonstrated that the weaving together of three-dimensional presentations with statistical visualizations can lead to new insights in complex, multivariate data. We have been able to observe interactions between simulation and anatomical variables which had not been anticipated. We expect this environment to be valuable in a wide range of application areas, where the goal is to find relationships, correlations, and patterns across multiple types of data.

Acknowledgements

We would like to acknowledge the contributions of David Rabenhorst, the creator of the Diamond technology, Edward Kalin, who created the initial implementation of the Diamond ActiveX components, and Greg Abram, who created the initial implementation of the Data Explorer ActiveX component. These people also contributed ideas and insight during the development of the environment discussed here.

References

- [AW95] C. Ahlberg and E. Wistrand. Iveen: An information visualization and exploration environment. In *Proceedings of InfoViz '95*. 1995.
- [BCW87] R.A. Becker, W.S. Cleveland, and A.R. Wilks. Dynamic graphics for data analysis. *Statistical Science*, 2:355–395, 1987.
- [BMMS91] A. Buja, J.A. McDonald, J. Michalk, and W. Stuetzle. Interactive data visualization using focusing and linking. In *Proceedings of IEEE Visualization '91*. 1991.
- [ID91] A. Inselberg and B. Dimsdale. Parallel coordinates: A tool for visualizing multivariate relations. In A. Klinger, editor, *Human-Interactive Systems*. Plenum Press, 1991.
- [LAC⁺92] Bruce Lucas, Gregory D. Abram, Nancy S. Collins, David A. Epstein, Donna L. Gresh, and Kevin P. McAuliffe. An architecture for a scientific visualization system. In *Proceedings of IEEE Visualization '92*, 1992.
- [NS00] Chris North and Ben Shneiderman. Snap-together visualization: A user interface for coordinating visualizations via relational schemata. *Proc. Advanced Visual Interfaces 2000*, May 2000.
- [Rab94] David A. Rabenhorst. Interactive exploration of multidimensional data. *Proceedings of the SPIE Symposium on Electronic Imaging*, February 1994.
- [RC94] Ramana Rao and Stuart K. Card. The table lens: Merging graphical and symbolic representations in an interactive focus + context visualization for tabular information. In *Proceedings of CHI'94, ACM Conference on Human Factors in Computing Systems*. 1994.
- [RLS⁺96] S. Roth, P. Lucas, J. Senn, C. Gombert, M. Burks, P. Strofollino, J. Kolojejchick, and C. Dunmire. Visage: a user interface environment for exploring information. In *Proc. Information Visualization*. IEEE, 1996.
- [RRGK96] B.E. Rogowitz, D.A. Rabenhorst, J.A. Gerth, and E.B. Kalin. Visual cues for data mining, 1996.
- [RT98] B.E. Rogowitz and L.A. Treinish. Data visualization: the end of the rainbow. *IEEE Spectrum*, Dec. 1998.
- [SHWF99] D. F. Scollan, A. Holmes, R. L. Winslow, and J. Forder. Histologic validation of reconstructed myocardium microstructure from high resolution mr diffusion tensor imaging. *Am. J. Physiol.*, 275:H2308–H2318, 1999.
- [UFK⁺89] C. Upson, T. Faulhaber, D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A. van Dam. The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4), July 1989.
- [WB97] Pak Chung Wong and R. Daniel Bergeron. 30 years of multidimensional multivariate visualization. In *Scientific Visualization: Overviews, Methodologies, Techniques*. IEEE Computer Society, 1997.
- [WKL99] David Weinstein, Gordon Kindlmann, and Eric Lundberg. Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *Proceedings of Visualization '99*, 1999.
- [WRJ⁺99] R. L. Winslow, J. J. Rice, S. Jafri, E. Marban, and B. O'Rourke. Mechanisms of altered excitation-contraction coupling in canine tachycardia-induced heart failure ii: Model studies. *Circ. Res.*, 84:571–586, 1999.
- [WS97] Bill Lorensen William Schroeder, Ken Martin. *The Visualization Toolkit*. Prentice Hall Computer Books, 1997.
- [WSH⁺00] R.L. Winslow, D. F. Scollan, A. Holmes, C. K. Yung, J. Zhang, and M. S. Jafri. Electrophysiological modeling of cardiac ventricular function: From cell to organ. *Ann. Rev. Biomed. Eng.*, 2:119–155, 2000.

Figure Captions

Figure 1: Typical three-dimensional visualizations of heart simulation data, with a rainbow colormap and a segmented colormap applied to the voltage variable. While relative levels of the specific variable colored may be qualitatively compared, this visualization is not well suited for simultaneously comparing multiple variables in a quantitative way.

Figure 2: The cardiac application interface. Using the two-dimensional scatterplot presentation, high values of voltage have been colored green, and high values of spatial current have been colored red. Regions that are high in both variables appear yellow. Using WEAVE, coloring in any presentation is automatically reflected in all presentations.

Figure 3: Linked parallel coordinates presentation, available as an additional window in the application shown in Figure 2.

Figure 4: Parallel coordinates presentation with only red cases shown, allowing the analyst to see the values of the simulation variables when spatial current is high, right before depolarization.

Figure 5 Four different representations of the anatomical view. In all cases, the red, green, and yellow coloring has been set as in Figure 2. Top: a cloud of marked node positions (with only red points shown); Second: marked positions combined with a colormapped surface rendering; Third: tensorlines representing fiber direction; and Bottom: a heart boundary representation with vertices sharing colors with marked cases.

Figure 6 The two-dimensional scatterplot presents C_{ai} (intracellular calcium) as a function of Ca_{SS} (calcium in the diadic space). Three populations appear. The two populations with positive values of C_{ai} have been marked with red and green. The three-dimensional model of the heart clearly shows that these two populations coexist in the same anatomical regions.