Evaluation of Methods for Topographical Feature Detection Using VR Interaction Techniques

October 5, 1999

Abstract

The main goal of the project is the creation of a visualization tool that allows the investigators to easily state the validity and accuracy of a topographical feature detection methodology, by means of direct real-time interaction with the data in a virtual reality immersive environment. Calculation steering will be implemented for those methods that allow it, being able to direct the computation towards the most accurate solution while it's being developed.
Apart from real 3D navigation and direct interaction with the data, the use of a completely immersive environment will introduce a new user interface, which can be developed towards more advanced studies into new feature detection methods, not currently in development because of the visualization limitations.

Participants:

Daniel Acevedo Feliz	Principal Investigator (Computer Science Dept.)
Prabhat	Co-Principal Investigator (Computer Science Dept.)
Frederic F. Leymarie	Collaborating Investigator (Engineering Dept.)

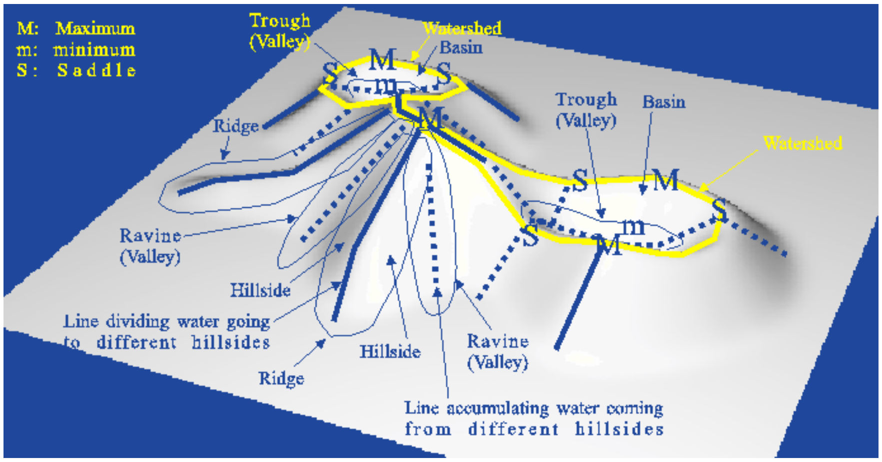

Figure 1: Illustration of typical topographic features (adapted from [3, Fig.1]).

Project Description:

Introduction:

The main goal of the project is the creation of a visualization tool that allows the investigators to easily state the validity and accuracy of a topographical feature detection methodology, by means of direct real-time interaction with the data in a virtual reality immersive environment. Calculation steering will be implemented for those methods that allow it, being able to direct the computation towards the most accurate solution while it's being developed.

Apart from real 3D navigation and direct interaction with the data, the use of a completely immersive environment will introduce a new user interface, which can be developed towards more advanced studies into new feature detection methods. The latter often are difficult to grasp because of the visualization limitations. Today, 2D static imagery remains the main method used to analyze - mostly qualitatively - the results of different methods (e.g., see Fig. 2). The difficulties in making comparisons explicit are clear (e.g., see [3, 2]), but little else has been made available thus far.[1]

Although the detected features in 2D (images) or 3D (digital elevation models; DEM's) datasets can be compared and recognized as topographical characteristics (see Fig. 1), the fields of application of these methods (also called ridge and valley detection methods) are very broad: From their direct use in topographical analysis, map generation and other GIS (Geographical Information Systems) applications, to their application in medical imaging as methods for extracting medial axes of blood vessels from coronary arteriographies or as correlation methods to match CT and MR images of the same patient [3]. The extraction of drainage patterns from DEM's or satellite images is another of the main areas where these detection methods can have a great impact [4]. The study of the drainage networks and flooding conditions of an area has classically been carried out manually or at least human-assisted, but not automatically. This is a continuously developing field with great impact in both civil and environmental engineering projects.

The use of VR interfaces to accomplish this analysis of the different methods shall be proven useful and will initiate an important movement towards the integration between the study of feature detection methods and the most advanced scientific visualization tools, normally being developed in parallel of each other [5].

Background and Significance:

The common goal of all ridge and valley detection methods is to find a way to express the underlying dataset in a more concise, understandable way. Instead of having to look at a large collection of numbers representing coordinates, heights, colors, etc. or trying to understand grey scaled images with lots of

[1] Note that no absolute reference exist to describe what a valley or ridge ought to be. Arguments and theories are still been proposed in the recent literature and no precise definition has been established as a standard (e.g. see [1] for a discussion of this topic and a good historical and mathematical perspective).

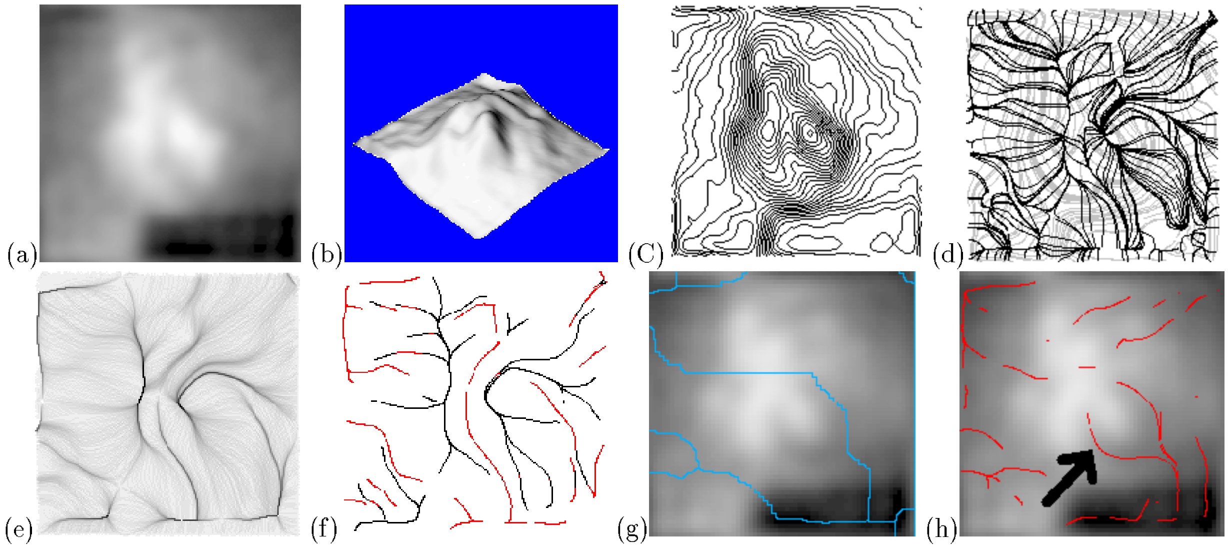

Figure 2: (a) Heart gammagraphic image. (b) Image seen as a landscape. (c) Level curves. (d) Some slopelines (black) over the level curves (grey). Observe how slopelines crowd along special slopeline segments. (e) Accumulation image where darker grey-levels denote greater accumulation. (f) Selected special slopelines segments: drainage pattern (red) and inverse drainage pattern (black). (g) Watercourses on the original gammagraphic image, detected via the algorithm in [7]. An important channel at the center of the image, on the heart, is missing since an 'opportune' maxima is not present. (h) Drainage pattern by following Rothe's definition. The channel on the heart is detected (pointed by the black arrow). Adapted from [3, Fig.10 & 11].

noise which makes the analysis harder and sometimes impossible, these methodologies extract the main geometric features automatically [6].

To understand why one refers to the name of *ridge* and *valley detection methods*, let us assume we have a dataset that represent a certain portion of land where the x and y coordinates are arranged in a grid and the z coordinate represents height. We can now create a mesh that interpolates that data and characterize its slope in every point. If we let a drop of water fall in any point of that surface, it will follow a path which has the steepest slope and eventually accumulate with other drops in a common path, that we'll call a *valley* [1].

On the other hand, there must be a point (or a set of points, that's to say, a line) on the *hills* of that dataset were if left on one side of that line, the drop will follow one path, and if left on the other side it will follow a completely different path, eventually going to a different minima of the data set. We'll call those lines the *ridges* of the dataset.

Following [2] we can classify the different methods in three main groups: Local methods, global methods and multilocal.

Local methods are based on the analysis of the mathematical representation of the dataset. They basically study the characteristics of continuity and derivativeness of the gradient functions that can be constructed from the points. They analyze a very local region of the data each time, so it's very difficult to extract continuous, meaningful global features.

Global methods take the whole dataset at once and detect topographical features all over the mesh at the same time. Watershed methodology [7] is a classic example of this. It's based on simulating the immersion of the generated interpolated mesh in water, letting it flow into through *holes* performed in some of the lower points. Which minima we make holes in is one of the decisions to be made before running the simulation. The number of holes depends on what level of detail we want.

Finally *multilocal* methods simulate the flow of water over the surface. They are called multilocal because we can choose which areas of the dataset to analyze and extract features only in those we are more interested.

System Overview:

The project will compare three different methods of ridge and valley detection. One local, one global and one multilocal [2]. The system created in this project will allow the user to chose among different DEM's to run the methods as well as which method to run in each moment. The digital models will be high resolution data either obtained from real sets of height data or from range data obtained with laser. The high resolution condition is given because of necessity of having interesting noticeable features as well as small, more local, perturbations on the mesh.

Different initial conditions have to be given depending on the method chosen. Thanks to the allowed 3D navigation over the data and complete immersion of the user in it, the interesting areas of the mesh can be easily recognized at a glance. The watershed method, for example, need to be given the interesting minima from where the immersion process will begin. This was usually done by direct analysis of the contour map, which is not intuitive at all for less trained users.

Also the possibility of scaling the data will be given to the user. Is important to highlight that different results can be obtained from the same data depending on the scaling factor in which the detection method runs on. Small tessellation of the mesh can create poorly defined features, but very highly defined meshes can lead to over-detection of non-principal ridges or valleys.

While the solution is being calculated, the user will be able to step into the process and follow closely the development of one of the valleys being detected (this is specially interesting in the case of using a multilocal approach, were the runoff over the surface is simulated), or look at it from far away and follow the over all behavior of the method.

At the end, a database of solutions can be created and simultaneously displayed in front of the user. The navigation tools will be very important for this direct comparison of the process.

Work Plan:

The following tables explain the work plan for the six weeks, with check points either weekly or every other week. The work braking up between the PI and the Co-PI was decided based on the skills and background of each one of them (see CV's at the end of this proposal). Of course the interaction between both investigators will be continuous and this tables reflect just their main tasks.

Week	PI
1	Data preparation
2	Data visualization and navigation
3	First results. Refining of results visualization
4	Continue with results vis. Steering process.
5	Steering process
6	Final results and presentation

Week	Co-PI
1	Algorithm's implementation design
2	First Algorithm implemented
3	Second Algorithm implemented
4	Third Algorithm implemented
5	Results database
6	Final results and presentation

Facilities and equipment:

The project will be developed in the new virtual reality facilities of the Technology Center for Advanced Scientific Computing and Visualization (TCASCV) of Brown University were a new cave environment has been installed and is currently being used as a research tool by different departments including Computer Science, Engineering, Applied Math., Geology and others.

The cave environment is currently running under IBM's UNIX based operating system (AIX)on the IBM SP 6000, but it's also ready to be connected to the existing SGI machines existing in the TCASCV.

JOT will be used as basic API to develop the system. JOT was created by the Graphics Group of the Department of Computer Science and it's ready to integrate calculations code and virtual reality user interfacing in the cave. This API is independent of the operating system were it runs, so the code created for this project could be run in either IBM, SGI or SUN architectures.

Future Applications:

The availability of this new VR visualization technique will help in the research of new methodologies for ridge and valley detection. Applications in Hydrology and Environmental studies would benefit from this technique allowing a more easy and intuitive understanding of large datasets. Studies for construction of dams and hydroelectric power stations would be carried out in a similar way, recognizing optimum and less intrusive locations. Drainage networks for lineal works (roads and highways) is another example of application.

Object recognition and matching, face or finger prints recognition are also good examples that would clearly benefit from this methodology, again because of the use of a real 3D interface with datasets that are being analyzed as 3D, characterizing them with a new, more concise language in terms of geometric features.

References

[1] Jan J. Koenderink and Andrea J. van Doorn. The structure of relief. *Advances in Imaging and Electron Physics*, 103:65–150, 1998.

[2] A. Lopez, F. Lumbreras, J. Serrat, and J. Villanueva. Evaluation of methods for ridge and valley detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):327–335, April 1999.

[3] Antonio M. Lopez. Ridge/valley-like structures: Creases, separatrices and drainage patterns. Computer Vision online: Geometric Feature Extraction Methods, 1997. Under http://www.dai.ed.ac.uk/ daidb/staff/ personal_pages/rbf/ CVonline/ LOCAL_COPIES/LOPEZ/ cvonline.html.

[4] Olivier Monga, Nasser Armande, and Philippe Montesinos. Thin nets and crest lines: Application to satellite data and medical images. *Computer Vision and Image Understanding*, 67(3):285–295, september 1997.

[5] T. Rhyne. Going virtual with geographic information and scientific visualization. *Computer and Geosciences*, 23(4):489–491, 1997.

[6] Yoshihisa Shinagawa, Tosiyasu L. Kunii, Alexander G. Belyaev, and Taketo Tsukioka. Shape modeling and shape analysis based on singularities. *International Journal of Shape Modeling*, 2(1):85–102, March 1996.

[7] Luc Vincent and Pierre Soille. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, June 1991.

Participant's CV:

Daniel Acevedo Feliz — Principal Investigator (Computer Science Dept.)

NAME: Daniel Acevedo Feliz
BIRTHDATE: October 15th, 1974.
PLACE OF BIRTH: A Coruña, Spain.
POSITION: Graduate Student, Computer Science Dept., Brown University.
EDUCATION:
- B.S in Civil Engineering in the University of A Coruña (Spain), 1997.

RESEARCH EXPERIENCE:
- Research Assistant in the Area of Continuous Media Mechanics and Structures Theory of the School of Civil Engineering of the University of A Coruña (Spain), since September 1995.
- Research Assistant in the Group of Visualization for Engineering and Urban Design of the School of Civil Engineering of the University of A Coruña (Spain), since September 1995 until August 1998.
- Research Assistant in the Graphics Group of the Computer Science Dept. of Brown University, since June 1999 until September 1999.

CAPABILITIES AND SKILLS:
- Strong Civil Engineering background including topography, hydrology, structures, and advanced numerical methods.
- Computer Science background and coding in C, C++, Java and Fortran.
- Virtual Reality background, as research assistant in the Graphics Group of the Computer Science Dept. of Brown University, specially oriented to the CAVE environment and its user interface programming. Experience with JOT (API used in this project) and with the different architectures that the project will use (SUN, SGI and IBM).

Prabhat — Co-Principal Investigator (Computer Science Dept.)

NAME: Prabhat
BIRTHDATE: May 4th, 1977.
PLACE OF BIRTH: New Delhi, India.
POSITION: Graduate Student, Computer Science Dept., Brown University.
EDUCATION:
- B.Tech in Computer Science and Engineering from the Indian Institute of Technology, New Delhi (India), 1999.

RELEVANT COURSES:
- Computer Graphics
- Digital Image Processing (programs on Edge detection and Region based segmentation)
- Computer Vision
- Independent Study on Virtual Reality (with a module on creation of VEs from Large Terrain databases)

RELEVANT PROJECTS:
- Virtual Walkthrough on an SGI Onyx RE-2
- Interaction with Virtual Puppets in a Distributed VE

CAPABILITIES AND SKILLS:
- Computer Science and Engineering background.
- Programming experience in C, C++ and Java.
- Exposure to VR packages and programming in OpenInventor, OpenGL and IRIS Performer.

Collaboration Letters:

Frederic F. Leymarie	Collaborating Investigator (Engineering Dept.)

The issue of developing appropriate visualization tools for the better understanding and comparison of methods, aiming at eliciting shape and other geometric features, is of great importance from many standpoints, either in the pure sciences or the applied ones. Topographical representations have been used for years (centuries!) in the limited format of paper and, more recently, photographs and 2D digital images. 3D visualization and VR or immersive interaction techniques shall bring the domain of topographical analysis to a new life.

This is a very exciting project, which I believe can shed new light on existing paradigms when dealing with topographical and other surfacic features, as well as to permit to better understand the multi-dimensionality of more sophisticated, recently developed methods.

I will gladly support the investigators in this project, by providing the necessary theoretical background and algorithmic insights that I have acquired over the years in dealing with problems - mainly in geography, hydrology, vision and medical imagery - which either necessitate or can benefit from the use of sophisticated topographical representations.

Simulation and Animation of Learning and Development in Recurrent Neural Networks

PI: Asohan Amarasingham (Applied Mathematics & Cognitive Science)
Co-PI: Matthew Hutson (Cognitive Neuroscience)
Co-PI: Avi Walsky (Neuroscience)
BUSPFA Interdisciplinary Scientific Visualization Proposal
October 3, 1999

Abstract

We propose to apply contemporary computer graphics techniques for animation and scientific visualization to aid specific research problems in the theory of developing sparsely-connected, recurrent neural networks. Despite the importance of these problems for the understanding of temporal learning in both biological and artificial neural networks, the temporal dynamics of recurrent networks has long resisted an analytic solution. Hence we intend to approach the problem of characterizing network behavior empirically, by numerical simulation of appropriate mathematical models and the development of real-time visualization techniques which will enhance our ability to understand the time-dependent behavior of these high-dimensional, nonlinear systems.

1 Introduction: The Scientific Question

We address a very simple question here: can a sparse and recurrently-connected network of (very simplified) neuron-like elements learn, in an unsupervised fashion, to generate a *reliable temporal sequence* of activity patterns in response to a weak externally-originated input?

The structure of this problem arises in many contexts and we were motivated by several related considerations. One of these motivations arises out of some fairly general issues in the theory of neural networks. At the source of much of the influence which neural network and connectionist models have recently exerted in cognitive science and artificial intelligence has been a recasting of many fundamental problems in learning into the framework of statistical inference between input and output vectors. The machines built to undertake this inference have in the main consisted of feedforward networks employing supervised learning rules such as backpropagation (Rumelhart et al., 1986). Feedforward networks by themselves are unable to learn statistical relations which are linked to time. The importance of temporal statistics is particularly evident in particular examples such as language and navigation, but really fundamental to all of cognition: at a most basic level, both behavior and the environment unfold over time, and a great deal of information is contained in the statistics of the *temporal* structure of behavioral-environmental events.

There is a tradition of research which extends basic connectionist machinery to handle spatiotemporal structure, utilizing the intrinsic temporal dynamics of recurrent feedback networks as a natural tool (Amari, 1972; Anderson et al., 1977; Hopfield, 1982; Elman, 1993). A primary analogy here has been between the attractor dynamics of recurrent feedback networks and autoassociative memories (Amit, 1989).

In line with this approach, Levy and several colleagues (Levy, 1989, 1996), motivated by considerations of the physiology and function of the mammalian hippocampus, have put Hebbian learning in *sparsely*-connected feedback networks to work with the goal of forming unstable attractors for sequential recall and prediction. Using this framework, they have demonstrated that such networks are able to solve several interesting problems – such as sequence completion, sequence disambiguation, goal-finding, and several other predictive problems (for a review, see Levy 1996) – whose essential statistical structure exists in temporal as well as spatial domains.

It would certainly be useful to understand how and which types of sequence codes are likely to develop as a function of architectural parameters and external input characteristics in these networks, but a satisfactory theoretical understanding of how these networks learn remains largely an unsolved problem. Successful approaches here, and in recurrent networks in general, typically employ equilibrium and mean-field type assumptions (e.g., Amarasingham & Levy, 1998) to understand network behavior prior to or following learning, but a coherent account of on-line Hebbian learning has remained elusive. Our own empirical observations have led us to view code formation as emerging from a competitive balance between the intrinsic temporal dynamics of the network and the imposed structure of weak but regular external input. The present work hopes to add to the insights of these empirical observations. Here, we want to minimize the external input and see what the interaction between on-line Hebbian learning and the network's dynamics produces. What happens if we just provide a weak but regular spark of external input every once in a while and let the network relax on its own, but with learning turned on? Will a reliable sequence of activity patterns learn to follow the external input probe? Will the network fall into an attractor? Will it remain statistically inert?

Levy and Sederberg (1997) have proposed that this problem is tied to the brain's solution of trace classical conditioning. Trace conditioning is a variant on classical conditioning in which the conditioned stimulus (CS; e.g., a tone) is followed by a stimulus-free delay preceding onset of the unconditioned stimulus (UCS; e.g., an eyepuff to a rabbit). Hence the animal must learn not only the CS-UCS association, but also keep track of the *amount* of time following offset of the CS in order to properly anticipate the onset of the UCS air puff with an eyeblink.

Levy & Sederberg approach this problem in a manner similar to the toy problem we investigate here: a weak external input signifying the CS sparks the sparsely-connected network, the network relaxes for a period of time signifying the delay, and another external input finally marks the UCS. Over time, if the network can learn to produce a reliable and temporally precise sequence of activity patterns in response to the CS external input, this sequence can function as a clock with which the animal can anticipate the UCS.

There are some generic neural computation issues present here as well. Some of the motivation for neural network research is that artificial neural networks may provide a reasonable first-order approximation of real brains. But brain regions are not fully-connected and they do not process information in a single shot in analogy to a feedforward neural network; rather they exhibit a rather serious paucity of connections, and processing occurs over time in the midst of considerable feedback. Both of these observations seem to hold in varying degrees pretty much anywhere we choose to go in the brain, and they naturally suggest that the utility of 'neural net' approaches to understanding the brain would be substantially enhanced if we had a better understanding of how sparsely-connected recurrent networks behave.

Consideration of feedback networks as brain models ties into one of the most foregrounded issues in contemporary neuroscience: the role of time in the representation space of neural codes. This issue is particularly evident in the familiar rate-versus-timing code debate: do neurons transmit their 'interpretation' of presynaptic events via changes in average firing rate over some coarsely-determined temporal interval, or does the precise distribution of spike times convey information (Rieke et al., 1997; Abeles, 1991; Shadlen & Newsome, 1998)?

Motivated by electrophysiological observations of precise and reliable spatiotemporal firing patterns, Abeles (1991, 1993) has argued for the timing code side of the debate and formulated a *synfire chain* model of cortical circuitry. Synfire chains are characterized by subpopulation neuronal groups activating in a temporally precise and reliable sequential fashion. The idea of synfire chains in cortex have received a lot of interest of late because they provide a functional context for observations of precise spike patterns, and because the idea of multiple, interacting synfire chains suggests the intriguing possibility of a larger scale organization for neural processing. Bienenstock (1995,1996) , in particular, has related synfire chain dynamics to the "binding problem" in cognition (Treisman, 1996), and has argued that the dynamical binding of coactive synfire chains may provide a way for the brain to represent compositional hierarchies.

The scientific goal of this project is to study how such reliable temporal patterns might form from an initially unstructured feedback network via self-organizing (Hebbian-type) processes.

2 Scientific Visualization

The visualization problem is easy to state: the network consists of N neurons, typically on the order of thousands, connected by cN^2 time-varying synaptic weights. Time development consists of a long run of discrete time steps divided into epochs for the sake of organizing periodic external input. At a given time step, each neuron is each either active or not active. Hence at each time step, the network is a stochastic process which lives in a binary vector space $Z(t) \in \{0,1\}^N$. The visualization question is then: how best to visualize $Z(t)$ as a function of time?

The goal of this project is to construct an animation of $Z(t)$ along with a simulator to sample from the underlying (discrete, finite-dimensional) stochastic differential equation. The animation ideally would provide researchers with a 'visual intution' of the time-dependent evolution of the network by combining visualization of $Z(t)$ over time with information (augmented with color) identifying the external input and properties of the evolving connectivity matrix.

3 Research Plan

The research will be conducted over a period of six weeks. The first two weeks will be largely exploratory, as the researchers will meet to determine a strategy for data simulation, and alternative parameters for the visualization.Neural network parameters will be drawn from the literature, and the second through fourth weeks of the project half of the investigation team will write software and run these simulations. The other half of the investigating team will write software for animation of this data during this same period. In the fifth week of the project these two aspects of the project will be combined, and all researchers will work on implementing the animation with the simulated data.This will continue into the sixth week along with documentation of the results.

References

Abeles, M. (1991). *Corticonics: neural circuits of the cerebral cortex*. Cambridge; New York: Cambridge University Press.

Abeles, M., Bergman, H., Margalit, E., & Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, *70*(4), 1629–1638.

Amarasingham, A. & Levy, W. B. (1998). Predicting the distribution of synaptic weights and cell firing correlations in a self-organizing, sequence-prediction model. *Neural Computation*, *10*(1), 25–57.

Amari, S.-I. (1972). Learning patterns and pattern sequences by self-organizing nets of threshold elements. *IEEE Transactions on Computers*, *21*(11), 1197–1206.

Amit, D. J. (1989). *Modeling Brain Function: the world of attractor neural networks*. Cambridge: Cambridge University Press.

Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: some applications of a neural model. *Psychological Review*, *84*, 413–451.

Bienenstock, E. (1995). A model of neocortex. *Network*, *6*(2), 179–224.

Bienenstock, E. (1996). Composition. In A. Aertsen & V. Braitenberg (Eds.), *Brain Theory – Biological Basis and Computational Principles* (pp. 269–300). Elsevier Science B.V.

Elman, J. L. (1993). Learning and development in neural networks: the importance of starting small. *Cognition*, *48*, 71–99.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.*, *79*, 2554–2558.

Levy, W. B. (1989). A computational approach to hippocampal function. In R. D. Hawkins & G. H. Bower (Eds.), *Computational Modeling of Learning in Simple Neural Systems* (pp. 243–305). Orlando, FL: Academic Press.

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, *6*(6), 579–590.

Levy, W. B. & Sederberg, P. B. (1997). A neural network model of hippocampally mediated trace conditioning. In *Proceedings of the 1997 IEEE International Conference on Neural Networks*, volume I, (pp. 372–376).

Rieke, F., Warland, D., van Steveninck, R. R., & Bialek, W. (1997). *Spikes: Exploring the Neural Code.* Cambdridge, MA: MIT Press.

Rumelhart, D. E., McClelland, J. L., & the PDP Research Group (1986). *Parallel distributed processing: explorations in the microstructure of cognition.* Cambridge, Mass.: MIT Press.

Shadlen, M. N. & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience*, *18*(10), 3870–3896.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, *6*, 171–178.

Visualizing Heat, Salinity, and Water Flows for a Holistic Understanding of Salt Fingers

Paul S A Reitsma

October 5, 1999

Primary Investigator: Paul S A Reitsma
Other Investigators: Daniel Keefe, Song Zhang
Oceanography Consultant: Jordan Dawe (University of Washington)
Visualization Consultant: David Laidlaw

1 Abstract

While the images used to illustrate salt fingers in systems with double diffusion instability are adequate, the application of scientific visualization techniques should allow creation of more intuitive and illustrative methods of viewing these phenomenon, particularly for teaching applications.

2 Introduction

2.1 Background

The hot sun of the tropics causes significant evaporation from the surface of the ocean, increasing the salinity of the surface water. This, in turn, increases the density of the water, and the resulting salinity gradient pulls the surface water down. The sun also warms the surface, however, lowering the density and creating a temperature gradient that keeps the surface water bouyant. Both excess salt and heat tend to diffuse across these gradients, resulting in increasing instability in the double diffusion system until some equalizing phenomenon takes place.

Small local variations provide the resolution to this: a small amount of warm, salty water near the surface sinks slightly into the cooler water below, and begins diffusing heat into it. Since heat diffuses approximately 100 times faster than salt, the surface water becomes increasingly cooler while retaining essentially the same salinity. This increases its density and accelerates its plunge into the cooler subsurface waters. As this is occuring, the fresher water around in is gaining heat and, hence, being bouyed upwards. This effect happens through the entire interface layer in a rough grid of cells up to three centimetres across and several times that length. The appearance is similar to that of meshed aqueous 'fingers' of roughly human size transporting salt from the surface; hence the name of the phenomenon.

2.2 Importance

These conditions which lead to the formation of salt fingers are present over very large areas of tropical oceans, including roughly 25% of large areas of the tropical Atlantic, and so understanding them is important to understanding the oceans in those areas.

Moreover, layers of liquids generally don't mix well on their own. Turbulent wave action provides some mixing, but not very quickly. The gravity-fed heat pump mechanism of salt fingers, however, mixes salt from the surface around twenty times as quickly as waves. This has the effect of pumping large quantities of salt into the deeper parts of the ocean where it tenuously balances the colder, fresher water flowing from the north. Understanding salt transfer, then, is of surprising importance in understanding the large-scale features of the ocean.

Further, very similar phenomenon have been detected causing flaws in cooling metals and are theorized to be present in stars, among other places. The hope is that an improved visualization tool for salt fingers would aid investigation into other similar effects as well.

3 Proposal

Current display methods for salt fingers, while generally adequate, are perhaps not as informative as they could be. This project will make images that not only more accurately but more intuitively convey the underlying mechanics of salt finger activity; a particular goal is better representation of the actual diffusion of heat and salt as well as their relative rates.

One common method of illustrating salt finger activity is simply to form salt fingers in the lab with the warmer, saltier water dyed. This allows salt fingers to be seen directly. Unfortunately, while an intuitive illustration of salt fingers, this type of presentation provides little information on heat diffusion, little quantitative information, and virtually no information emphasis.

The other common method of illustrating how salt fingers work is a schematic diagram. While this gives a much clearer presentation of the mechanics involved, all but the most rudimentary information on diffusion and motion is lost.

The goal of this project, then, is to combine the best features of these display methods, providing both the intuitiveness of the picture with the informational content of the diagram. Painterly information representation techniques perhaps based on those used by Laidlaw [1] to visualize fluid flows offer promise for retaining the flowing nature of the picture while encoding the data of the diagram.

The resources of the project will initially be devoted primarily towards securing a data source for creating the visualizations; this will most likely involve implementing a moderately simple fluid dynamics simulator with reference to those already created by oceanographers and in consultation with the project oceanography consultant. This should be complete after approximately three weeks. While implementation is under way, research into and design of the visualization methods to be used will also take place, culminating shortly after the simulator is complete. At this point, the chosen visualization paradigm will be implemented and after approximately four and one half weeks, the visualizations will be available. Ad hoc user studies will take place at this point, along with feedback solicitation from oceanographers, particularly at the University of Washington and Dalhousie University. Finally, most of the last week will be devoted to preparing a final report of the project's results.

4 Applications

Primarily, a salt finger visualization would be useful as a teaching tool to introduce people, particularly non-oceanographers, to the mechanics of salt finger formation. Additionally, however, a sufficiently flexible visualization would be useful in investigating how and why salt fingers form under varying conditions; while such experiments can be done in the lab for some initial conditions, a visualization would be much faster and much more general. Ideally, the visualization would be general enough to accept different state equations, allowing it to be used to investigate the formation of salt finger-like structures in other environments, such as stellar atmospheres.

Additionally, while only salt finger-like images will be created in this project, the techniques for combining the best features of pictures and diagrams which we will investigate should apply to many physical sciences where such illustrations are used.

5 Related Work

Computer simulations and pictures of salt fingers have already been done by the oceanography community, such as figure 2; additionally, a version of the phenomenon is simple to reproduce in the lab by layering dyed, warm, salty water over colder, fresher water; figure 1 is a movie of this.

What makes the proposed work a little more novel than a reimplementation of old ideas is the exploration of intuitiveness and expressivity in displaying salt fingers. Most current images are more akin to schematic diagrams, with large, sparse arrows very roughly showing the flow of heat and coloured regions representing the movement of water, but with little or no sense of salt diffusion. The proposed visualization aims to be more expressive in the details of the phenomenon, particularly at the interfaces where the diffusion is actually taking place. The hope is that this will provide a more intuitive visualization for explaining and understanding the mechanism than either photographs of actual dyed salt fingers or simple schematic drawings. Additionally, the techniques developed for melding pictures and diagrams should find use in other fields.

6 Timeline

6.1 Work Schedule

- October 19: final proposal
- October 26: design or aquire simple fluid simulator
- November 2: compile and get first core dump from fluid simulator
- November 9: fluid simulator running
- November 16: tie simple display onto simulator
- November 23: upgrade simple display to interesting display
- November 30: ad hoc user studies and subsequent system alterations
- December 7: prepare abstract and final presentation

- December 13: final presentation

6.2 Deliverables

- October 19: final proposal completed
- November 9: fluid simulator running
- November 23: preliminary visualization available
- December 7: user study results
- December 13: final presentation

7 Referenced Work

1. R. Michael Kirby, H. Marmanis, D. Laidlaw. Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting. Visualization '99 Proceedings, October 1999.

Curriculum Vitae of Paul Reitsma

Box 1910, Department of Computer Science
Brown University
Providence, RI 02912
psar@cs.brown.edu
http://www.cs.brown.edu/∼ psar/
Education

- BSc, Combined Honours in Computer Science and Mathematics, University of British Columbia, with Honours, 1998

Relevant Experience

- Research Assistant, Electronic Games for Education in Math and Science (E-GEMS), University of British Columbia, January to April 1997
  - Web-based educational game development for grades 5-7
- Research Assistant, TRIUMF National Laboratory, January to April 1995
  - Assisting Neutrino Group with software simulations

Honours

- Andries van Dam Fellowship

Skills

- Strong computer science background and programming experience
- Strong mathematical background
- Good technical writing background
- Basic grounding in most physical sciences

Name: Daniel Keefe
Birthdate: Dec. 27 1976
Place of Birth: USA
Current Status: Graduate Student / Researcher Brown University, Providence RI
Address: Box 1910, Brown University, Providence RI 02912
Telephone: (401) 863-7662
Education:

- 1999 B.S. in Computer Engineering, Tufts University, Medford MA

Research and/or Professional Experience:

- NASA Langley Research Center - Summer Research Program 1998. Worked in NASA's Data Visualization and Animation Lab at NASA Langley.
- Programming experience at Duke University Medical Center and CogniSyst, Inc.

Related Papers:

- Keefe, Daniel (1998). Advanced Computer Graphics Techniques for Accurately Reconstructing 3-D Flow Visualizations from Multiple Wind Tunnel Images and Measurements. NASA LARSS Program. Awarded IEEE Prize for best student paper 1998. unpublished.

Capabilities and Skills:

- Advanced Oil Painting experience with instructors from the Museum of Fine Arts School, Boston MA.
- Experience with Scientific Visualization of airflow in windtunnels.
- Computer Engineering/ Computer Science background.

CV of Song Zhang

Box 1910, Department of Computer Science
Brown University
Providence, RI 02912
sz@cs.brown.edu
http://www.cs.brown.edu/sz/

Education:

- Brown University, PhD Candidate.
- Nankai University, P.R.China, 1996, BS in Computer Science.

Research Experience:

- Research assistant, Brown Graphics Group, Sept. 1998-present.
  Research in Diffusion tensor field visualization.
- Research assitant, Nankai University, Sept. 1996-May. 1998.
  Developing a commercial OCR(Optical Character Recognition) product, TypeReader.

Relevant Courses and course projects in Brown:

- cs123, a recursive raytracer.
- cs224, photorealistic rendering using photonmap, mini-sketch in Java3D, nonphotorealistic rendering, photorealistic rendering using metropolis method

Relevant skills:

- Experience in diffusion tensor visualization.
- Programming skill in C++, Java
- Knowledge in Graphics

Curriculum Vitae of Jordan Dawe

School of Oceanography
Box 357940
University of Washington
Seattle, WA 98195-7940
freedryk@ocean.washington.edu

Education

- BSc, Honours in Physics, University of Victoria, 1999
- M.S., Physical Oceanography, University of Washington, expected 2001
  - M.S. Advisor: Dr. Luanne Thompson

Relevant Experience

- Lab Assistant, Dr. Christopher Garrett, University of Victoria, May to July 1999
  - Development of Fluid Dynamics/Oceanography Labs for 4th year Undergrads
- Research Assistant, School of Oceanography, University of Washington, September 1999 to present
  - Modeling the circulation and heat flux of the North Pacific.

Skills

- Strong physical science background
- Good fluid dynamics and oceanography background
- Programming experience in MATLAB, C and C++ on Windows and Unix based systems
- High school teaching experience
- Good writing/communication skills

Quantifying the Benefits of Artistic Methods in CFD Visualization

J. Reiter, J. LaViola, R.M. Kirby[*], D.H. Laidlaw[†]

Abstract

Finding useful and accurate techniques for presenting many fields of vector and scalar data in a single visualization is essential for developing good CFD visualization techniques. Recent work borrows ideas from artists to produce such visualizations. The authors wish to quantify the benefits of this approach over traditional visualizations.

[*]Brown University Center for Fluid Mechanics
[†]Brown University Department of Computer Science

1 Goals

A typical CFD problem requires understanding the relationships among half a dozen data quantities, both scalar and vector. It simply is not possible to apply traditional visualization techniques to produce a single picture for that many different quantities at one time. One recent novel technique for solving this problem is to present each field of data as a layer of strokes in a painting, leveraging knowledge about painting to guide stroke style and placement. While this technique appears promising there have been no studies to measure its effectiveness. We propose to undertake such a quantative analysis of artistic methods in CFD visualization. The proposed research would focus on the understanding of correlations among sets of variables in traditional and artistic visualizations–and attempt to quantify the differences between the two approaches.

1.1 Scalar - Scalar Correlations

The simplest type of correlation to glean from a visualization is one between two scalar quantities–for CFD this might be the relationship between vorticity and turbulent charge. In this case the traditional visualization approach is to generate two pictures, each of which is colored according to a single variable's value throughout the domain. The scientist then has two pictures describing the behavior of a fluid in one situation and is required to mentally assimiate the information to understand correlations.

The artistic approach here would produce a single image colored with one scalar value. The other scalar would be rendered as a collections of strokes on top of the background. Here the scientist is presented with only one image for all of the relevant data. However, putting more data into a single image increases the potential for confusion and error. The proposed research would measure that error and allow an objective analysis of whether artistic techniques are useful for CFD.

1.2 Scalar - Vector Correlations

Vector quantities are very common in CFD and are far more difficult to deal with. Even in cases where a strict mathematical relationship exists between two quantities, like vorticity and velocity, it can be hard to see the correlation. This is only complicated by human perceptual issues which traditional techniques totally ignore. An artistic rendering of a vector field would take advantage of human perceptual characteristcs to convey velocity properly: instead of using uniformly spaced vectors it would make them farther away the faster the fluid is moving because this is how we see.

Again there is the possibility for added error. While the vector quantity will undoubtedly be rendered in an iconic fashion much like the second scalar in the scalar-scalar case, there is even more potential for error as this icon must indicate several values at once. Can a human glean this much information from a single image? In the 2-dimensional case they are expected to understand the

Figure 1: Traditional visualization of flow past a cylinder. On the left we display vorticity and velocity, on the right turbulent charge and turbulent current.

relationship among one scalar and the 2 components of the vector–a complex 3 variable relation. Does an artistic presentation serve to help, or does it only increase the error? Again, this question can only be answered by observing CFD researchers working with both data presentations.

1.3 Vector - Vector Correlations

This last class of CFD visualization is already pushing the limits of traditional visualization techniques. Most traditional tools allow only a single vector field to be viewed at one time, making two images a necessity. Even if one can get several vector fields into a single image the sheer number of arrows drawn from each sample point serves to make the visualization very difficult to use.

By simply selecting different icons for each data field, artistic methods allow one to present more vector data at once than a traditional method. But there still exists a potential for error. In this case CFD scientists will be presented with three visualizations: a multi-vector-valued artistic image, a multi-vector-valued traditional image (i.e. arrows), and several traditional visualizations covering all the variables for a single domain.

Three visualizations are being analyzed for this stage because there are two questions to answer. First, is switching to an iconic representation useful? And second, is it helpful to present all of the data in a single image? Answering these questions will give a better understanding of how to present CFD vector-valued data in a useful manner.

2 Methods of Testing

For each of the correlation-understanding cases discussed above a user study will be conducted. Example visualizations will be presented for each case and the fraction of correct correlations will be used to compare different techniques. Each type of correlation will be considered independently, results will not be pooled for general artistic vs. traditional conclusions.

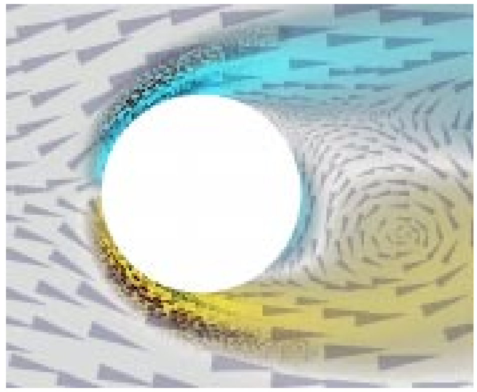

Figure 2: Artistic rendering of the above data. Triangles indicate velocity, background color indicates vorticity, particle color indicates turbulent charge and particle size/direction indicate turbulent current.

2.1 Test Subjects

The test subjects will be CFD researchers will at least one year of experience dealing with visualizations. This actually leans the test against artistic methods but there is no pool of subjects with significant CFD knowledge that has not worked with traditional visualization techniques. While useful human perception questions could be addressed by using a pool of untrained subjects, this research deals specifically with the CFD case. It is not, however, clear *a priori* whether familiarity with the data will increase or decrease the tendency to get confused by artistic visualizations. This is a question to be addressed by the research.

2.2 Test Visualizations

Since the scientists will be asked to describe correlations in the data it is essential that the presented data have correlations that can be fairly easily understood. Exceptionally complex cases will only cause everyone to get the wrong answer for all the visualizations. So the test cases will be fairly simple CFD problems. But the subjects will not be told which quantity is depicted by each color/icon/image to prevent prior CFD knowledge from influencing answers.

2.3 Sources of Error in Testing

It is hoped that this method of testing will reduce both confusion and non-visualization-based inference in the testing. It is less likely an experienced CFD scientist will be overwhelmed or totally confused by the visualizations than someone with no scientific visualization experience. But it is far more likely that someone with experience will make inferences using knowledge they have and did not get from the visualization–unless they do not know which variables they are looking at and therefore cannot mentally compute anything.

3 Research Plan

There are three phases to this research: visualization development, user testing, analysis of results.

3.1 Visualization Development

The first three weeks will be used to generate a good sample set of visualizations to use for use in the user study. A significant amount of CFD data is available to the authors so all of the effort for this portion of the time will go towards generating both traditional and artistic visualizations. The goal is to put together three to four of each type of visualization for each case to be examined. We have access to software for generating both traditional and artistic visualizations but, as with most scientific visualization software, it takes a fair amount of time to generate a useful picture. Plus, we must settle on the correct correlation for each visualization.

3.2 User Testing

Two weeks will then be spent presenting the sample set of visualizations to test subjects and recording the correlations they find. We plan to run these tests on ten CFD researchers in addition to ourselves.

3.3 Analysis of Results

The last week will be spent analyzing the data collected. This involves deciding which correlations were correct–while we will settle on a correct answer for each case during phase 1 there is still a small grey area. The actual calculations and data formatting will only take a few hours to complete. At this time we will also be able to draw conclusions about which cases benefit most from an artistic representation of data.

4 Summary

This research works to quantatively demonstrate that artistic methods are useful for CFD visualization. Further, we will compare the utility of artistically-inspired visualizations with the current state of the art to determine if researchers have reason to begin using artistic visualization techniques today.

CV for Jonathan Reiter

Personal

Name: Jonathan Reiter
E-Mail: jr@cfm.brown.edu

Education

2000 A.B. Computer Science, Brown University (expected)
2000 A.B. Mathematical Economics, Brown University (expected)

Employment

1996-98 Research Assistant, Thomas Jefferson University Hospital
1998-present Research Assistant, Brown University

Related Work

Interactive Visualization and Steering of Large-Scale CFD Computations, on-going research project at Brown University

CV for Joseph J. LaViola Jr.

Personal

Name: Joseph J. LaViola Jr.
Email: jjl@cs.brown.edu

Education

Ph.D, Computer Science, Brown University (Expected May 2002)
Sc.M, Applied Mathematics, Brown University, (Expected May 2001)
Sc.M, Computer Science, Brown University, October, 1999
B.S., Computer Science, Florida Atlantic Universty, 1996

Employment

Brown University Computer Graphics Lab, Researcher, May 1998 - Present
Fraunhofer Center for Research in Computer Graphics, Research Scientist, May 1997 - December 1997
UCS, Inc., Software Technician, Summer 1996
Internnational Business Machines, Computer Information Systems Dept., May 1995 - Novermber 1995

Honors

IBM Research Assistantship(1999)
IBM Cooperative Fellowship(1998)
Elected as Associate Member to Sigma Xi(1998)
FAU's Aaron Finerman Award(1996)
FAU's Faculty Award for Outstanding Undergraduate Achievement(1996)
Microsoft Senior Achievement Award(1995)

Relevant Experience

Active Stock and Options Trader, Investment Researcher for 12 years

Interactive Visualization of Surfaces in 4D

October 5, 1999

Participants
Prabhat (PI)
Daniel Acevedo Feliz (Co-PI)
Collaborators
Prof. Thomas Banchoff (Maths)

Abstract

A variety of computer graphics techniques have enabled the display of surfaces (which exist in higher dimensions) on a 2D screen. In the proposed work we are going to develop novel techniques to generate, navigate and interact with a surface in 4D, in an immersive environment. We shall try and investigate if visualization in the CAVE is any different from 'seeing' the surface on a computer screen and if it qualitatively helps the user to understand the surface better.

1 Introduction

The techniques of interactive 3-dimensional computer graphics can be extended so that users (particularly mathematicians) can examine surfaces that lie in 4-dimensional space. A significant body of literature exists in Topology and differential geometry which deals with the nature and properties of such surfaces. Unfortunately conventional computer graphics limits the user to view and interact with these fascinating shapes from 'outside' and try to make sense of them. The development of immersive environments presents a great opportunity to investigate these surfaces from a new perspective. But will the fact that we can interactively manipulate, navigate and generate a surface help us to get an intuition into what it really is ? We plan to look into this interesting question and find some answers (and perhaps more questions!) for ourselves.

The hypothesis that we hope to test is that by using the interaction techniques that we develop (tetrahedral control space and interactive generation) the user can visualize surfaces in 4 dimensions better than the existing methods.

2 Related Work

A lot of work has been done in the visualization of surfaces in 4D but it is largely confined to the conventional 2D screen[1][2]. There is a free software *geomview* which has some demos of generation of 4D objects[3]. Some work on visualization in immersive environments like the CAVE is being done in NCSA at the Univ. of Illinois- Urbana Champaign[4]. Our work will be different in the sense that it will provide the user with a control space to navigate the surface, an intuitive way to interact with the surface and the ability to generate it.

3 Work Plan

Our first job will be to import some VRML models which were made by David Cervone under Prof.Banchoff some time back. They will be ported over to the cave and "interacted" with using some of the standard techniques being used right now for a Scene Modeler application. This will be done in the first week.

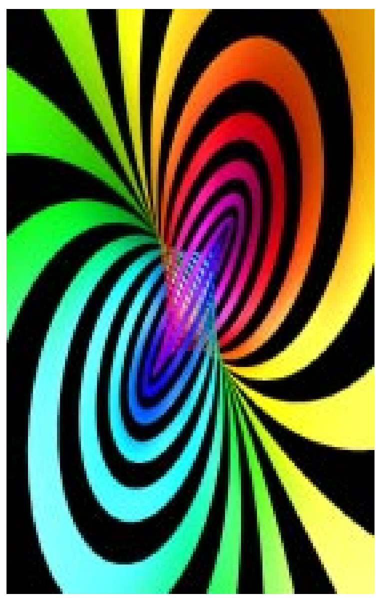

Fig.1 "In- And Outside the Torus" A VRML model made by David Cervone

In the next 2 weeks we will work on developing basic interaction techniques like rotating a surface in 4D, developing a "tetrahedral" control space (which will give us a choice of projection and viewpoint) and generating the surface interactively by varying the 4D parameters. In the fourth week we will implement some existing techniques like "ribboning", color coding and single-sided paint application to help the user visualize the surface better. In the fifth week we will conduct a user-study amongst the Graduate/Undergraduate students who are taking related courses to test if the visualization methodology has indeed helped them to understand the surfaces better. In the final week we will be working on the abstract and presentation.

4 Facilities

The whole purpose of the project is to visualize the surfaces in an *immersive* environment and see if the tools that we develop aid us in understanding 4 Dimensions in an intuitive manner. We will be working in the cave and building on *jot*'s interaction techniques. Prof. Banchoff has been working on visualizing surfaces using graphics on a standard computer monitor for a long time. He has many issues in the pipeline which we shall tackle to the extent that time permits.

5 Significance of Work

Unfortunately there is time enough only for developing a "toolkit" for visualizing surfaces. Once an effective interaction methodology is in place we can build on it for tackling bigger problems which arise from visualization of 4D datasets. We can then interactively 'fit' surfaces onto such datasets and try to see what class of surfaces the data corresponds to[5][6]. Besides a more intuitive understanding of 4D visualization could lead to development of interfaces which enable you to kinesthetically navigate through higher dimensional datasets.

References

[1] D.Banks, PhD Thesis, "Interacting with surfaces in Four Dimensions using Computer Graphics", 1993.

[2] A.J.Hanson, T.Muzner and G.Francis, "Interactive Methods for Visualizable Geometry", *Computer,* July 1994.

[3] `http://www.geom.umn.edu/software/geomview/geomview_toc.html`

[4] `http://new.math.uiuc.edu/galleryNSF/`

[5] T.Banchoff, "Visualizing Two-Dimensional Phenomena in Four-Dimensional Spaces : A Computer Graphics Approach", *Statistical Image Processing and Graphics,* 1986.

[6] Kocak, Huseyin, F.Bishopp, T.Banchoff and D.Laidlaw "Topology and Mechanics with Computer Graphics: Linear Hamiltonian Systems in Four Dimensions", *Advances in Applied Mathematics,* Vol. 7, 1986, 282-308.

Participant's CV:

Prabhat

NAME: Prabhat

BIRTHDATE: May 4th, 1977.

PLACE OF BIRTH: New Delhi, India.

POSITION: Graduate Student, Computer Science Dept., Brown University.

EDUCATION:
- B.Tech in Computer Science and Engineering from the Indian Institute of Technology, New Delhi (India), 1999.

RELEVANT COURSES:
- Computer Graphics
- Digital Image Processing
- Computer Vision
- Independent Study on Virtual Reality

RELEVANT PROJECTS:
- Virtual Walkthrough on an SGI Onyx RE-2
- Work on a Hyperboloidal Projection System for OmniGraphics
- Implementation of a Distributed VE

CAPABILITIES AND SKILLS:
- Computer Science and Engineering background.
- Programming experience in C, C++ and Java.
- Exposure to VR packages and programming in OpenInventor, OpenGL and IRIS Performer.

Daniel Acevedo Feliz

NAME: Daniel Acevedo Feliz

BIRTHDATE: October 15th, 1974.

PLACE OF BIRTH: A Coruña, Spain.

POSITION: Graduate Student, Computer Science Dept., Brown University.

EDUCATION:
- B.S in Civil Engineering in the University of A Coruña (Spain), 1997.

RESEARCH EXPERIENCE:
- Research Assistant in the Area of Continuous Media Mechanics and Structures Theory of the School of Civil Engineering of the University of A Coruña (Spain), since September 1995.
- Research Assistant in the Group of Visualization for Engineering and Urban Design of the School of Civil Engineering of the University of A Coruña (Spain), since September 1995 until August 1998.
- Research Assistant in the Graphics Group of the Computer Science Dept. of Brown University, since June 1999 until September 1999.

CAPABILITIES AND SKILLS:
- Strong Civil Engineering background including topography, hydrology, structures, and advanced numerical methods.
- Computer Science background and coding in C, C++, Java and Fortran.
- Virtual Reality background, as research assistant in the Graphics Group of the Computer Science Dept. of Brown University, specially oriented to the CAVE environment and its user interface programming. Experience with JOT (API used in this project) and with the different architectures that the project will use (SUN, SGI and IBM).

Interactive Visualization and Manipulation of Molecular Structure in a Virtual Reality Environment

PI: Pramod Paranthaman
Co-PI: Quan Gu

ABSTRACT

As a result of intense research in the fields of chemistry and molecular biology it has been found that the geometric and topological study of molecular structure is very relevant to the study of molecular behavior. The shape and topological structure of protein molecules play an important role in their function, for instance, the identification of structures such as cavities and channels may assist in the design of closer packed proteins. Many of the existing visualization techniques, when applied to large molecules may obscure important details of features within the molecule. Conventional display devices may not provide adequate depth information when compared with stereo displays available in virtual reality environments. Also, the use of 3D input devices such as pinch gloves may provide an easy to use and more intuitive interface, for viewing and manipulating molecular models, than 2D input devices such as mice.

This proposal describes a project to develop a molecular visualization tool using new and innovative techniques for the visualization and manipulation of molecules in an immersive display environment.

SPECIFIC AIMS

The overall aim of this project is to create a VR tool which chemists and molecular biologists can use to visualize and manipulate fairly complex organic molecules.

More specifically the program will be able to read standard NSF protein data bank files, which are freely available on various publicly accessible servers. It will support conventional molecular visualization techniques, such as the wire-frame model, ball and stick model and the space filling model (Figures 1, 2 and 3). We plan on exploring other representations, by making use of transparency and varying degrees of opacity in the space filling model. Time permitting, the project will also include molecular surface representations allowing the researcher to interactively switch between the Van der Waals, Connolly, and solvent accessible surface views (see background and relevance). The program will also support an easy to use gesture-based interface for viewing the models.

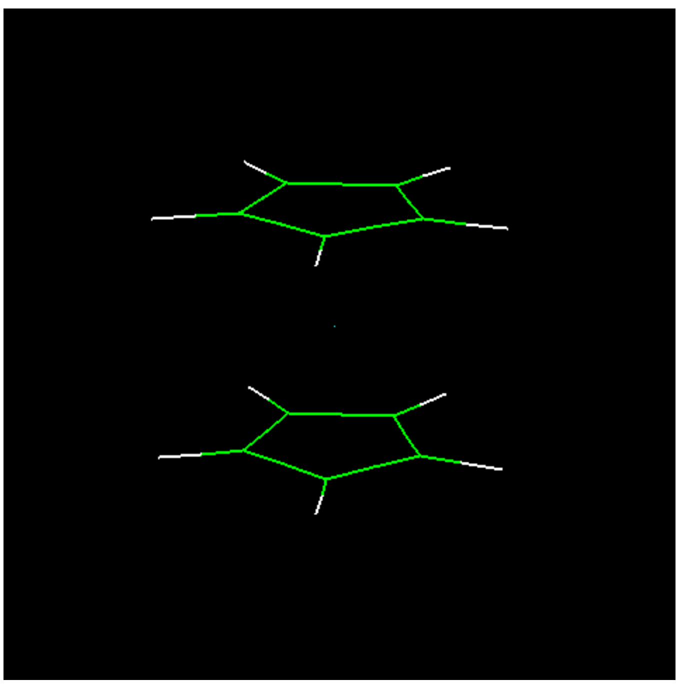

Figure 1

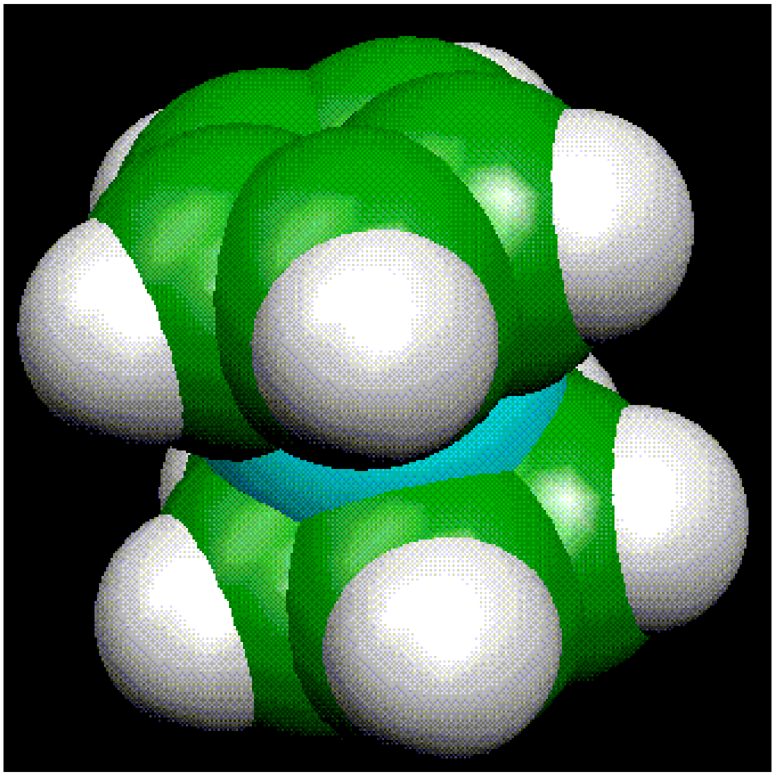

Figure 2

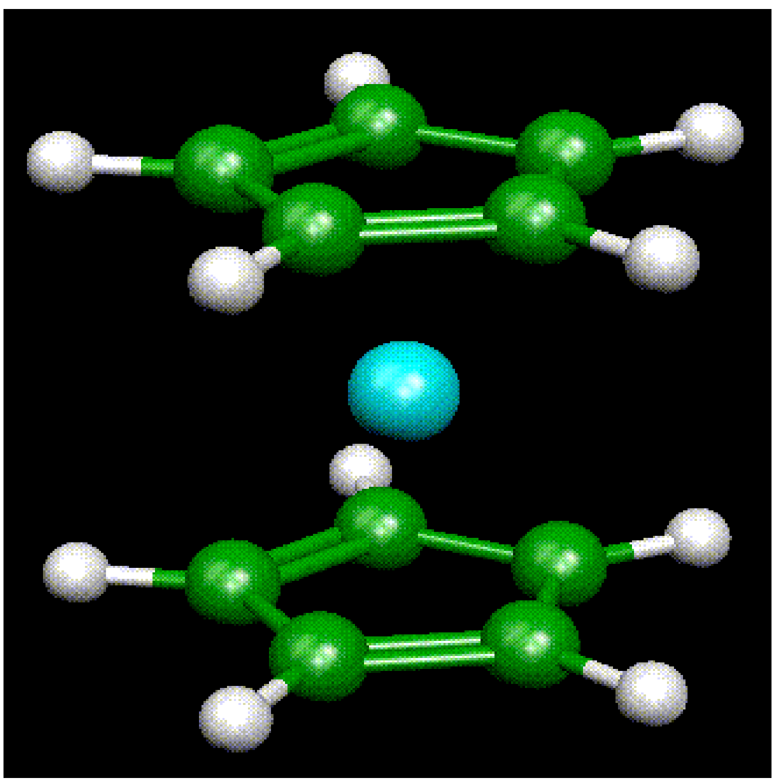

Figure 3

BACKGROUND AND SIGNIFICANCE

To better understand the functioning of complex proteins and enzymes it is often necessary to visualize their molecular models. Most models are obtained either experimentally, through X-ray crystallography and nuclear magnetic resonance (NMR) studies, or through theoretical modeling. The visual identification and inspection of cavities, pockets and channels in three dimensional molecular surfaces is useful in identifying binding sites for other molecules. This can be of great use in the field of drug design where molecules are designed to fit a specific receptor or active site.

As the number of atoms constituting a molecule becomes very large, interesting features of the molecule become obscured. This problem is usually overcome by clipping away parts of the molecule to reveal hidden features or making certain groups invisible and color coding others. For very large convoluted molecules the results may be confusing (Figure 4). Also effects such as depth cueing have little effect due to the high density of atoms and it becomes very difficult to judge the relative depth of atoms close to each other. It is fairly easy to get disoriented when using techniques such as clipping.

We propose that the use of a 3D virtual reality environment will eliminate or significantly reduce the effects of these problems. Making the researcher believe he is physically present within a molecule may give him a better sense of orientation. One reason for this is the wider field of view of VR environments such as the CAVE tm. The stereo display capabilities and the motion parallax effect provide better depth perception than the conventional CRT display.

The investigator will be able to manipulate the viewing parameters through the use of pinch gloves, we plan on developing an intuitive gesture based interface for this purpose. The program will support the visualization techniques mentioned earlier and, time permitting, will include the following surface visualizations:

Van der Waals surface - Van der Waals force is a weak attractive force acting between molecules. The surface corresponds to the molecular envelope containing atomic spheres of Van der Waals radii.

Connolly surface – The Connolly surface is composed of two kinds of surface patches. The first part is the Van der Waals surface of each atom which is accessible to a probe sphere of a given radius, generally 1.4 Angstrom units, corresponding to the contact surface. The second is the inward facing part of the probe sphere when it is in contact with more than one atom (Figure 5).

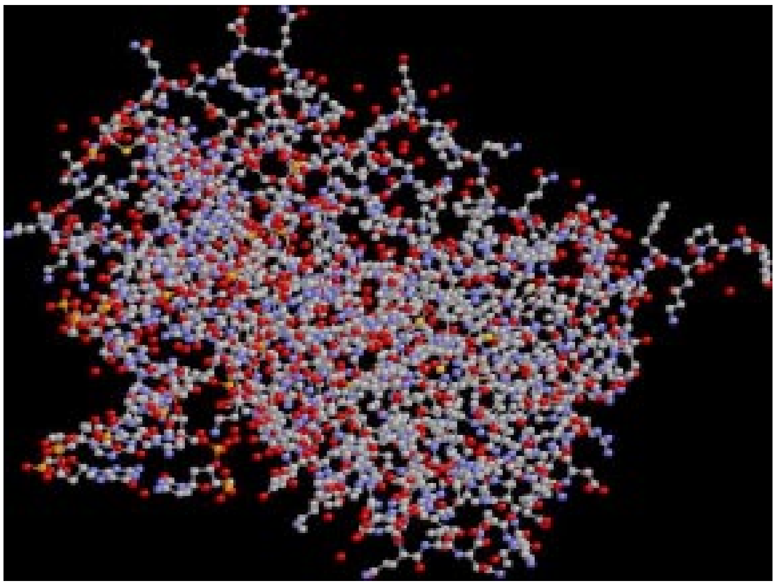

Figure 4

Specific properties can be mapped on each of the previously described surfaces. The simplest of which is color coding the surface as a function of the atom color (Figure 6), purple for iron, grey for carbon etc. Another possibility is to map the molecular electrostatic potential (MEP) on the Connolly surface. The MEP can be used to predict regions within the molecular substrate which are most reactive towards other molecules.

When visualizing larger molecules, we plan on giving the investigator the ability to selectively make surfaces and atoms partially transparent, so it becomes easier to view obscured features normally not visible.

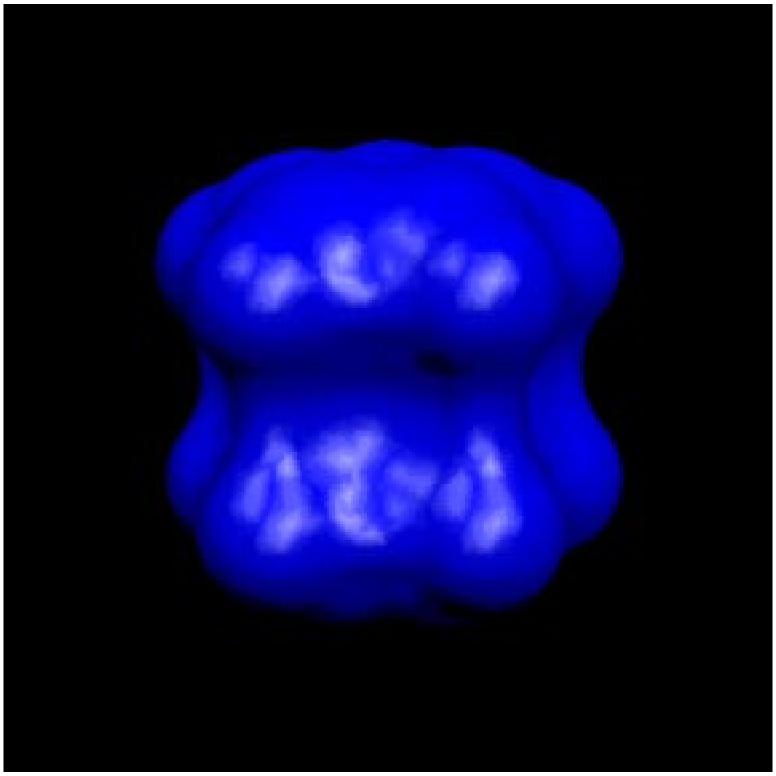

Figure 5

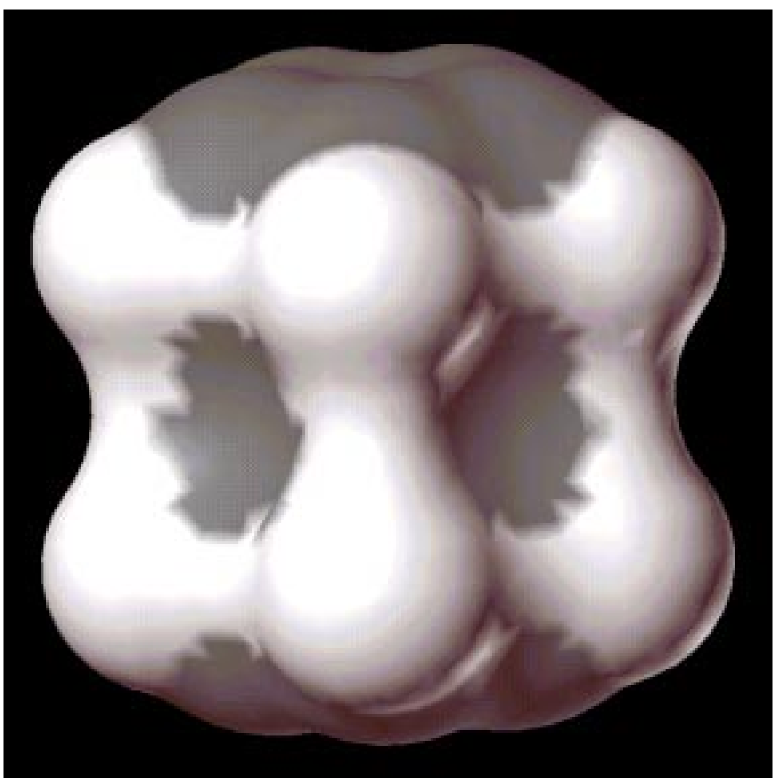

Figure 6

RELATED WORK

Much work has been done in the field of molecular modeling and visualization by various researchers. Most of the work focuses on conventional display devices and a few have extended their tools to work in VR environments. Very little work has been done on investigating alternative molecular visualization techniques geared specifically towards VR, and is currently an open area for research.

WORK PLAN

The project is divided into two main parts. The first part reads the protein data bank files and generates the required image database and will take approximately 1 week. The second part generates the appropriate view of the model based on the gestures from the pinch gloves. There are numerous public domain programs that read NSF protein data bank format files, adapting the source code to work with our project shouldn't take up much time. We plan on using World Toolkit, a commercially available virtual reality library, which we currently have at the Cave to do the actual rendering. By the end of the 4th week we plan on having a working program. In the last week we plan on conducting user trials with investigators from the Chemistry department, so we can incorporate the things we learn from their suggestions and comments into our project.

Pramod Paranthaman
Computer Science Department
Brown University

Education and Professional Experience:
Brown University
Master of Science, Expected May 2000
Concentration: Computer Graphics
Worked at the Virtual Environment Navigation Lab (VENLAB) in the Cognitive and Linguistic Sciences Department as a graphics programmer.
Relevant Courses: Interactive Computer Graphics, Introduction to
Computer Graphics, User Interface Design

American Megatrends Inc.(AMI), Madras, India
System Software Engineer, 1997-1998
Worked on the design and implementation of various large software projects in the BIOS division at AMI.

College of Engineering, Guindy, Anna University, Madras, India
Master of Computer Applications, 1997
Relevant Courses: Interactive Computer Graphics, Computer Aided Design, Software Engineering

Loyola College, University of Madras, India
Master of Science (Physics), 1994
Developed CrysMols, an interactive teaching aid for visualizing crystal and molecular structure for the Department of Chemistry.
Relevant Courses: Organic Chemistry, Quantum Mechanics

Relevant Computer Skills:
Extensive programming experience in C, C++, and Java.
Experience in WorldToolkit, Java3D and Open GL.

Quan Gu's CV

Chemistry Department

Brown University

Skills

¨ Expertise in chemistry at Ph.D level. Especially the understanding of complex molecules called polymers.

¨ Extensive knowledge of biology, physics and other natural sciences. Knowing how scientists from other field utilize chemistry.

¨ Some CS course and programming experiences. Understanding the CS terminologies, can write simple programs.

Experience

¨ Conducting chemistry research.

¨ Wrting research proposal, papers about chemistry.

¨ Using commericially available molecular visualization softwares to drawing chemical structures for publications.

¨ Taken CS4, CS252, CS176, CS127 courses in brown.

Visualization of Color Spaces

Principal Investigator: Tomer Moscovich

CO-Investigator: Daniel Keefe

Consultants: Barbara Meier, Timothy Miller

1 Abstract

Used properly, color can be a powerful tool for communication. However, it is, by nature, a complex interaction of the physical world with the human visual system making it difficult to understand. Using color in an optimal way requires a deeper understanding of how it works. An interactive visualization system to aid in exploration of the nature of color could provide much useful insight.

2 Introduction

Color is a very powerful tool for visualization, interaction, and communication. For many people good use of color may be essential. It is used by scientists to visualize data, designers of software, and of course artists. Color, however, is very complex. It is not simply a physical property of objects, but the interpretation of light by the human perception system. As the product of such a system it is difficult to understand. Poor use of color is not merely inefficient, it is often confusing. A better understanding and awareness of the subject could be usefull to many people.

The perception of color is initially governed by the response of three types of cone cells in the retina to the different spectral components in light entering the eye. Various combinations of brightness and frequencies produce the many colors we see. These responses are the result of evolved biochemical processes, and as such do not yield a clean and simple relation. Radically different spectral distributions can be perceived as the same color. Visualizing the mapping between cone space and perceived color space could prove very interesting and instructive.

3 Goal

The goal of this project is to create an interactive method for people to learn concepts about color perception and color spaces. Instead of being task or lesson driven, the visualization is intended to be an exploration from which a deep understanding of color can be acquired. It should relate the physical, biochemical, and mental manifestations of color to each-other and to standard color models. The user interaction should be transparent, and should require as short a companion text as possible since it is to be more than an illustration. However, since the purpose is to get more than an elementary understanding of color, some basic knowledge may be necessary for the visualization to be accessible.

4 Related Work

The Color Group in the Computer Science department of Brown university has done quite a bit of work on visualizing color spaces and methods of explaining color. Some of these methods have been aiming for museum-like exhibits that illustrate a single point, while others involve high-level visualization of how color is used in painting. This project is intended to fill the space in between.

The educational programs range from two dimensional applets illustrating concepts such as color mixing and warm and cool colors, to immersive environments for viewing how colors are arranged in a color space. They are intended to provide a novice color user with an understanding of basic color principles. Work is also in progress on helping people use expert pallettes.

The visualizer of color in paintings creates three dimensional histograms showing the distribution of colors in a painting. It provides insight into the relationships of the colors chosen by the artist that can be compared to the histograms of other paintings and related to known principles of color use.

5 Project Plan

5.1 Week One - Research and Design

The first week of this project will involve research and design of a flexible interactive visualization environment that will enable users to freely explore color spaces and their relationships. Mapping between various spaces and interpretations should be possible, as well as their relationship to the physical world.

5.2 Weeks Two and Three - Implementation and Testing

The core functionality of the program will be implemented. Once this is complete, it will be informally tested. Feedback and suggestions will be noted.

5.3 Weeks Four and Five - Completion

The fourth and some of the fifth week will go towards completing the project. Suggested adjustments and fixes will be made.

5.4 Week Six - Testing and Paper

The writing of the paper will begin as will user testing. Non-expert users will be given time to interact with the program. Afterwards they will be asked to answer questions on color perception and color spaces. The success of the project will be partially measured by the resulting answers. A greater degree of success will be judged if people already familiar with color concepts can gain a deeper understanding after interacting with the visualization software.

6 People and Resources

The team is composed of two main designers and implementors, and of two outside consultants. Myself and Daniel Keefe will be in charge of authoring reasonable specifications for the project and the actual implementation. As principle investigator, I will coordinate the communication with the outside consultants. Barbara Meier will provide us with feedback tempered by her artistic background and knowledge of color. Timothy Miller has much experience in implementing color visualization systems, as well as a wealth of color knowledge and could provide us with much useful advice and feedback.

7 Significance

Color is a difficult subject, and many could benefit from a deeper understanding. This project is an extension of the Brown Color Group's goal of making color more accessible to people. It can help both people who need to use color well, and investigators of the nature and use of color.

Internet Traffic Visualization:
Individual and Collective User Modeling

Until now, most information about traffic to one’s web site comes from text log−files, a cumbersome and unintuitive method that says little about the behavior of visitors within the web site. Here I propose a method for visualizing individual and collective user access patterns amongst a set of html hyperlinks.

PI: Matthew Hutson
(Cognitive Neuroscience)

co−PI: Benjamin White (Neuroscience)
co−PI: Danah Beard (Computer Science)
Consultant: Rosemary Simpson
(Computer Science/Hypertext)

Use of the internet has grown exponentially in the past several years. Lacking from this growth has been an intuitive tool for studying user access patterns. Servers frequently log access reports, which contain such information as which specific http address is accessed, the time of each access, the referring address (which address the user came to this address from), and which internet browser each user is using. A typical log file may look something like this:

```
foo.bar.edu − − [16/Nov/1995:18:50:04 −0800] \
    "GET /$$87612/sigmod_record/ HTTP/1.0" 200 1252
foo.bar.edu − − [16/Nov/1995:18:50:14 −0800] \
    "GET /$$87612/sigmod_record/issues.html HTTP/1.0" 200 653
foo.bar.edu − − [16/Nov/1995:18:50:23 −0800] \
    "GET /$$87612/sigmod_record/9−95/ HTTP/1.0" 200 3565
foo.bar.edu − − [16/Nov/1995:18:50:29 −0800] \
    "GET /$$87612/sigmod_record/issues.html HTTP/1.0" 200 653
```

If the web site is popular, accruing millions of hits a day, the log file will contain thousands of pages of text similar to this. A curious web designer may read the log file to glean information about who is visiting the site and how often, but the information is not arranged in a manner that allows much insight beyond this. The designer has no idea about the patterns of behavior that people exhibit once they arrive at the site beyond the overall frequency of each page access. In understanding how people use a site it would be helpful to see their movement from one page to another in a serial manner. With this information one can tell what a user is looking for, how he is looking for it, and whether he found it, and then redesign the web page to better accommodate future users.

Current techniques in web use analysis involve data mining algorithms that scour log files and produce statistical reports on basic information. Methods have even been developed that can extract access reports in a log file referring to an individual user and list these sequentially. There also exist client−side applications that can track a user's web browsing and construct static visual maps of the web pages traversed. I suggest combining current access−log analyzing techniques with a method for visually mapping hperlinked nodes. From the log file I would extract information regarding links between web pages, and sequence them by user. Applying these sequences to a map of nodes and edges representing a web site I would create an animation following a user through the site, hilighting the node and edges in order.

By overlapping such traversals accumulated from many users, one would have a static map displaying node and edge access frequency by color or border thickness. If one wanted to see where people most frequently link to or from at a specific node (page) in the map (site), one could click on the node and the connecting edges (links) would change color according to popularity. One could then click on the most popular linked−to node and see where people tend to go from there. From such interaction one would gain much greater understanding of how pathway preferences relate to a larger pattern of site exploration.

This collective map could then be transformed into an animation by placing a virtual user, represented by a dot, in the map, and watching her traverse the network according to a random or specific individual user access log. Traffic patterns may be made more apparent by recreating 5 or 20 or 100 user paths simultaneously and watching the black dots cruise around the map like cars from a helicopter.

There are a number of uses for such a holistic visualization of web site traffic patterns. Designers need to know how visitors navigate their web site. Do they prefer returning to a central node frequently, or do they like to wander? How useful is the index? Does it help them find what they need quickly or does it send them on a goose chase? Marketing wizards and e−commerce developers would also pay handsomely for such a visualization system so they could track exactly how customers behave within a commercial web site and learn where to place promotional information.

There are also many educational applications. Educational multimedia attempts to present information to a user in a fashion that conforms to natural human learning tendencies (vision, sound, etc.), and hypermedia accelerates this process by allowing users to choose their own paths through the material. An analysis of how people navigate through hypermedia pieces (including hypertext) would feedback to the design process and tell authors how to rearrange their presentations. Frequent jumps between nodes might illustrate a conceptual link in the users' minds not considered or emphasized by the author. Unvisited nodes may need better placement or suggest exclusion from the material.

In the near future such work within the context of the internet will lead to other technological advances. While the proposed system can only inform a human as to a web site's traffic patterns, this type of procedure will evolve into a system that will actually suggest changes, and eventually make the changes automatically. New forms of hypertextual markup languages have been suggested (Procedural Markup Language, or PML; Adaptive HTML, or A−HTML) that will dynamically generate web pages based upon accumulated user access patterns. The content and organization of these pages will be automatically customized for each visitor. Knowing how to program such automated systems is impossible without first understanding intuitively how visitors navigate different kinds of sites.

Information learned from internet traffic analysis will also tell us about how people traverse hyperlinked data architectures in general. User modeling will be an important cognitive science in the fields of hypermedia design, HCI, and data mining. Understanding user tendencies when navigating networked terrain will be crucial to presenting large amounts of information in an effective and efficient manner.

First Week: Form a conceptual model of the final product. What do we want the visualization to look like? What should it tell us?

Second and Third Weeks: Design data collection procedure. What information do we need, and how should it be organized?

Fourth and Fifth Weeks: Graphics programming and user interface. This will involve a parallel process of coding and subjective testing.

Sixth Week: Fine−tuning the visualization. Generating preliminary suggestions about the design of the analyzed web site(s). Presentation preparation.

Curriculum Vitae for Matthew Hutson, Principal Investigator

Name: Matthew Hutson
Birthdate: April 14, 1978
Place of Birth: USA
Current Status: Researcher, Brown University
Home Address: Box 3417, Brown University, Providence RI 02912
Telephone: 401.331.4859

Education:
− 2000 Sc.B. in Cognitive Neuroscience, Brown University, Providence RI

Research and/or Professional Experience:
− Consultant, Multimedia Lab, Brown University
January 1999−present
− Web Master, Interlock Media, Cambridge, MA,
June 1999−September 1999
− Internet Assistant, 911 Gallery, Boston, Mass,
June 1998−September 1998

Relevant courses at Brown:
− CS15, Object−Oriented Programming
− CS16, Data Structures and Algorithms
− EL19, Hypertext Fiction
− EL100, Advanced Hypertext Fiction
- EL111, Hypertext, Cyberspace, and Critical Theory
-

Curriculum Vitae for Benjamin White

Name: Benjamin White
Current Status: Researcher, Brown University
Telephone: 401.867.6656

Education:
− 2001 Sc.B. in Neuroscience, Brown University, Providence RI

Research and/or Professional Experience:
− Brown University Computer Science Dept. (1999−present)
− Software Engineer, White Oak Technologies, Inc. (1999−present)
− Software Engineer, Lighthouse Technology Solutions, Inc. (1997−1999)

Capabilities and skills:
− Strong coding experience (awk, C, C++, Java, Perl)
− Graphics knowledge

Curriculum Vitae for Danah Beard

Name: Danah Beard
Birthdate: November 24, 1977
Place of Birth: USA
Current Status: Researcher, Brown University
Home Address: Box 0556, Brown University, Providence RI 02912
Telephone: 401.861.0131
website: http://www.danah.org/

Education:
− 2000 A.B. in Computer Science, Brown University, Providence RI

Research and/or Professional Experience:
− Software Engineer, Macromedia Generator, Macromedia, San Francisco, CA,
  June 1998−September 1999
− Research Assistant, Sociable Media, MIT Media Lab, Cambridge, Mass,
  June 1999−present

Relevant courses at Brown:
− CS32, Software Engineering
− CS123, Introduction to Computer Graphics

Capabilities and skills:
− computer science background and strong coding background (C/C++/Java)
− computer graphics knowledge

Modeling the length of wrist ligaments (BUSPFA draft proposal)

Elisabeta Marai, gem@cs.brown.edu

1.0 Specific Aims.

The goal of the proposed research is modeling the length of the ligaments that connect wrist bones to forearm bones. We intend to visualize the way these lengths modify over time, as the wrist moves. We will look for differences between the results obtained for healthy people and the ones obtained for people with injured wrists. We will analyze the results and hopefully draw conclusions on the link between ligament elasticity and wrist mobility problems.

2.0 Significance.

Understanding how the wrist works is an important issue in today's medicine, as doctors try to diagnose and solve wrist mobility problems. However, this goal is still unachieved; though there were some attempts of building a model for the wrist, its complex anatomy (concentrated in such a small area) makes this task difficult. We have already a reasonable model of the wrist-bones, thanks to scientific visualization, but nothing more; our project may be a good step towards adding to the existing bone-model a model of the ligaments.

However, without loosing sight of the big picture, our research addresses a specific problem: proving the link between ligament length and wrist mobility constraints.

In certain cases, constraints in wrist mobility are known to be caused by malunion of previously fractured bones. This type of problems can be identified with the help of CT and X-ray investigation. Sometimes refracturing the bones in order to achieve better union is necessary, but in the end patients usually recover their mobility.

But in many other cases CT and X-ray do not show evidence of malunited bones.Though apparently the bones are in perfect shape, patients still display wrist mobility constraintsconstraints. The most likely explanation, suggested by doctors, is that the mobility problem is in these cases caused by the ligaments that link the bones, and not by the bones themselves. A plausible reason could be loss in ligament elasticity. Studying ligament elasticity through ligament length appears in this context as particularly important.

3.0 Related Work. Basic Approach.

Current approaches to the problem of modeling ligament elasticity are analyzing ligaments collected from fresh frozen cadavers and "growing" bioengineered ligaments.

Unfortunately, elasticity of “fixed”/ bioengineered ligaments is very different from the one specific to “in vivo” ligaments. Bioengineered ligaments achieve only 25% of the native strength ([NRB]). “Fixed” ligaments, even laser-treated in order to reduce laxness, achieve only 24% and 40% of the strength/ stiffness of the intact ligaments ([AAOS]).

So far there has been no attempt of scientific visualizing ligament length starting from data collected by noninvasive means. It is very likely that researchers still wait to have appropriate MRI data in order to move to modeling soft tissues in the wrist. We try to prove with our research that in fact CT and anatomy book data suffice in order to get a good approximation of ligament length. We see this as the start of a first wrist ligament model; this model could be, after all, later enriched with additional MRI data. It is likely that MRI data would provide more accurate information about location of anchor points for specific patients; however this will not bring major modifications to our model, since we try to model length of ligaments and not surface or shape of ligaments.

Our ligament length model will be based on an existing wrist-bone model and on anatomy book data (such as anchor points and location of ligaments). We intend to start by interpolating between control points located on the surface of the bones, and to later evaluate the lengths of the resulting 3d curves. The bone model was reconstructed from serial CT data collected from patients of Dr. JJ Crisco, Sinai Hospital, using a version of the marching cubes algorithm. A parameterized version of this model (using manifolds [Gri]) is available thanks to C. Grimm, Ph.D. An animation of the model showing wrist movements specific to healthy and injured persons is also available, thanks to D.Laidlaw, Ph.D., and to C.Demiralp.

4.0 Work-plan.

The estimated duration of our project is 6 weeks, with the following milestones:

1st week: Get acquainted to input data and tool kit (Open Inventor). Consider possible models for ligament length. Choose a model.

3rd week: First visualizations. Compute length.

4th week: Refine visualizations. Add animation.

5th week: Compare results obtained for injured/healthy wrists. Start writing paper.

6th week: Wrap-up and conclusions. Finish writing paper. Final presentation.

5.0 Collaboration and Facilities.

The PI counts on the support of D. Laidlaw, Ph.D., and JJ Crisco, M.D. Dr. Laidlaw will counsel the PI on choosing and implementing the ligament length model, and Dr. Crisco will provide medical information throughout the project. Dr. Crisco will also be a great help in evaluating and interpreting the results.

The only facility needed is a Sun Ultra10 machine running Solaris, which is already available.

References:

[NRB] McGowan, J.C., “Report on the 3rd International Conference on Cellular Engineering”, Naval Research Biophysics Newsletter No. 33, 1997

[AAOS] Rokito, S., Shields, C., Park, S., Felix, B., “Efficacy of Laser Treatment on a Pathologically-Induced Lax Ligament Model”, American Academy of Orthopaedic Surgeons-1997 Annual Meeting

[Gri] Grimm, C., “Modeling Surfaces of Arbitrary Topology using Manifolds”

Visualization Tools for Sovereign Risk Analysis

Principal Investigator: Marc Majzner

Co-Principal Investigator: Joseph J. LaViola Jr.

Abstract

The analysis of sovereign risk consists of the employment of statistical models to estimate the probability of default by a debtor nation. Statistical methodologies identify significant dependent variables that influence the probability of default. The large number of dependent variables makes formulating the probability of default difficult. This work investigates the use of an interactive application that assists the user in dynamic weighting and associating of these variables using statistical methodologies and empirical evidence. This tool is intended for use by sovereign debt professionals, policy makers and educated investors to rapidly manage risks posed by international investments.

Introduction

Daily, $1.3 trillion dollars of capital flow through the global financial markets. The direction of the capital flows is reevaluated continuously by professionals striving to secure the highest returns. Concurrently, policy makers monitor these markets to ensure that capital flows do not radically disrupt the societies to which the capital is loaned or rescinded. This constant analysis of the markets is extremely difficult, as 1998's global crisis revealed when it left in difficult positions both the most powerful policy maker (International Monetary Fund) and one of the world's most successful international investors (George Soros).

The difficulty of sovereign risk analysis lies in the large number of variables whose values continuously fluctuate. Moreover, the dependency between variables fluctuates as well, making the use of a static function near impossible. Individuals must continually monitor these variables and update their models to formulate estimates of sovereign risk.

Analysis is typically performed using spreadsheets and database programs alongside optional graphs. This method lends itself well to static models where dependencies do not fluctuate. In these instances, static functions may be derived which serve in formulating the probability of default. However, in the dynamic environment of sovereign risk analysis, functions must be continually estimated. This process should reconsider all of the original variables instead of only those used in the previous model. Additionally, empirical evidence is useful in formulating these new models. An additional problem involves weighing the accuracy of the variables' values since many of the values are provided by non-transparent agencies.

There is the need for a tool that provides a better environment for reevaluating these variables with regard to empirical evidence and the reliability of the values. This tool would not only increase productivity but would also allow better decision making increasing the efficiency at which capital is allocated.

Goal

This work will devise a tool to assist international investors and policy makers in formulating sovereign risk. It will allow for user interaction in weighting and associating variables. Users will be assisted by the availability of empirical evidence and visual cues on the accuracy of information. This project hypothesizes that this tool will allow for both productivity gains and better decision making.

Related Work

Related work in the public domain is sparse due to the economic advantage that such a tool would provide to its users. The sole tool in the public domain is a 3D visualization of risk management analysis for use by large banks in evaluating interest rates, currencies and equity indices. This tool, designed by Visible Decisions (VDI), a North American software company specializing in the delivery of business visualization solutions, deals with a similar problem or risk management. Although the VDI tool does visualize a couple of the variables that this proposal will, it is not complete enough to perform a sovereign risk analysis.

Research Plan

Database Creation: (2 weeks)

An initial collection of 42 variables tested by Taffler and Abassi (1984) will be initially integrated. Data collection of these variables will be difficult because the agencies releasing the numbers are just recently becoming transparent and because there are a lot of values when considering that each variable has as many values as there are countries. Moreover, there is no single available public source of this data. In the event that adequate current data cannot be found, historical numbers can be retrieved through well-known periodicals and journals. This data will be sufficient for demonstrating the hypothesis.

Application Development: (4 weeks)

A number of visualization techniques must be built into the framework to allow the user to view the data. (2 weeks) The application must also support a method to view historical data and to compare data between different time periods. (2 weeks) An extensible framework will be created for this with similar services being afforded to each component.

CV for Joseph J. LaViola Jr.

Personal

Name: Joseph J. LaViola Jr.
Email: jjl@cs.brown.edu

Education

Ph.D, Computer Science, Brown University (Expected May 2002)
Sc.M, Applied Mathematics, Brown University, (Expected May 2001)
Sc.M, Computer Science, Brown University, October, 1999
B.S., Computer Science, Florida Atlantic University, 1996

Employment

Brown University Computer Graphics Lab, Researcher, May 1998 - Present
Fraunhofer Center for Research in Computer Graphics, Research Scientist, May 1997 - December 1997
UCS, Inc., Software Technician, Summer 1996
International Business Machines, Computer Information Systems Dept., May 1995 - November 1995

Honors

IBM Research Assistantship(1999)
IBM Cooperative Fellowship(1998)
Elected as Associate Member to Sigma Xi(1998)
FAU's Aaron Finerman Award(1996)
FAU's Faculty Award for Outstanding Undergraduate Achievement(1996)
Microsoft Senior Achievement Award(1995)

Relevant Experience

Active Stock and Options Trader, Investment Researcher for 12 years

CV for Marc G. Majzner

Personal

Name: Marc G. Majzner
Email: mgm@cs.brown.edu

Education

Sc.M, Computer Science, Brown University (Expected May 2000)
B.S., Computer Science, Brown University, 1999

Employment

Brown University Graphics Group, Researcher, Sept. 98 - Current
Silicon Graphics, Engineer, Summer 1998 Brown University Graphics Group, Researcher, Sept. 97 - May 98
AT&T, Information Systems Manager, Summer 1997

Honors

National Honor Society President, 1994-1996

Relevant Experience

Active Stock and Options Trader
Heavy economics background

Waves in the Market: A Pixel-Based Approach to Visualizing Market Trends in Historical Stock Data

Principal Investigator: Joseph J. LaViola Jr.
Co-Principal Investigators: Marc G. Majzner, Paul Reitsma
Consultant: Jonathan Reiter

Abstract

Financial market indices such as the Dow Jones Industrial Average and the S&P 500 are often used as a tool for finding market trends and evaluating the performance of the overall stock market. However, these indices only provide a minimal amount of information and make it difficult to understand what is occurring on cross- and intra-sector levels. In the proposed work, we intend to investigate a pixel-based approach to viewing a large quantity of stocks which can be used for exploring market trends in historical stock data and provide a lower level representation for what occurs within and across market sectors.

1 Introduction

The principal method for viewing historical stock data, such as price and volume, has traditionally been with line and bar graphs (see Figure 1). These graphs are good at representing data for a small number of stocks. However, when investors wish to view more than a few individual stocks, these line graphs get cluttered and difficult to interpret. As an alternative, investors can examine market indices which group larger quantities of stocks together, such as stocks in an individual industry or sector, to form a single number which shows how they perform as a whole. Investors can then view these indices historically to identify market trends and patterns.

Figure 1: A typical chart for showing historical stock data. This chart shows stock volume and closing prices for the past year.

The major drawback with these stock indices is that in the process of providing a mechanism for investigation into the market trends of a group of stocks, important information on the individual stocks that make up the index is lost. In the case of market indices which encompass stocks within a single sector, information such as individual stocks that buck the market trend or which stocks are driving the index down are not represented. With market indices that encompass stocks across sectors, information such as how stocks from one sector influence stocks from another is also not shown. As a result, a method of viewing large quantities of historical stock information that has the power to show both high level holistic information and low-level individual stock patterns is needed.

1.1 Goals

The goal of this proposed work is to establish and evaluate a set of tools which allow investors to investigate a large quantity of stocks and how they relate to each other. By representing each stock in a grid as a small group of pixels (a stockcel), we can show volume, closing stock prices, and daily price spreads of hundreds of stocks by varying colors of the individual stockcels. Breaking up the grid into sectors allows investors to see how certain sectors effect others and how stocks within sectors influence each other through the course of time. Our hypothesis is that *viewing historical stock data using stockcels will provide investors with an easier and more intuitive method for exploring trends and price patterns in the stock market.*

2 Related Work

Unfortunately, the amount of related work in this area is unknown since most institutional and other investors do not want to reveal their methods for analyzing the stock market. However, one stock visualization tool that is publicly available and is a basis for this work is SmartMoney's "Map of the Market"[1]. This tool, shown in Figure 2, shows over 600 stocks from the NYSE and Nasdaq market sites. Each stock is represented by a rectangle in the grid where the rectangle's size reflects market capitalization. The colors represent a stocks percentage increase or decrease for a unit of time. Users can tell which stock is which by simply putting the mouse cursor over a given rectangle.

Although "Map of the Market" provides related functionality to the tools we are describing in this proposal, there are a number of drawbacks to the system. First, their tool set is static in that it does not allow the user to moving dynamically through a given time interval to view stock prices on a daily, weekly or monthly basis. Second, it only provides information for the past two years. Third, it does not show other information such as volume or price spread. Our proposed work will essentially extend this application to be more robust and more powerful.

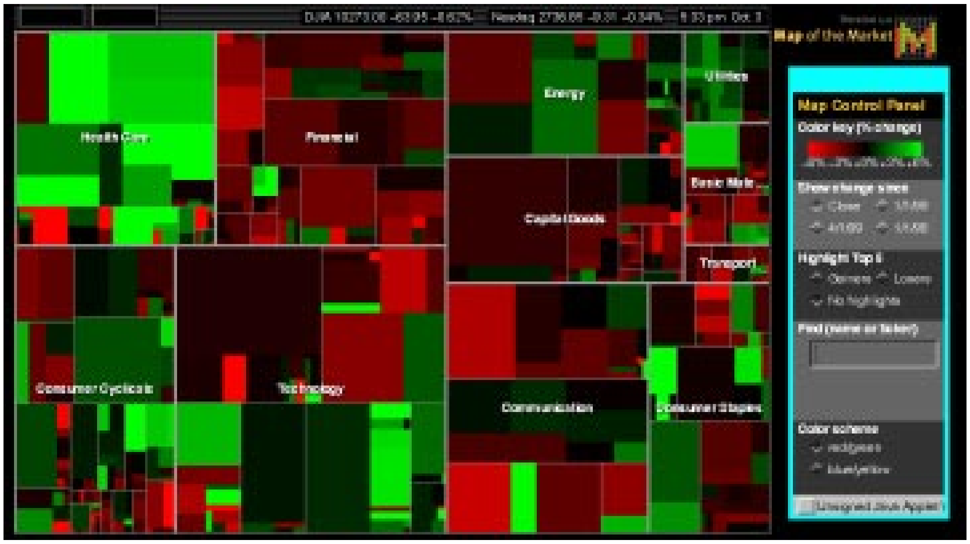

Figure 2: A snapshot of SmartMoney's "Map of the Market". Each stock is represented by a square and the size of the square reflects market capitalization. Over 600 stocks are represented.

3 Research Plan

This project is divided into three main parts: stock database creation, visualization application development, and evaluation.

3.1 Stock Database Creation

In the first two weeks of the project, the research team will collect historical stock data using the Yahoo historical quote website[2]. This website allows for simple downloading of text files which contain daily records for the high, low, close, and volume for a stock. The database will contain the last four years of data for anywhere between 500 to 3000 stocks depending on time constraints. Stocks will be taken primarily from the NYSE and Nasdaq market sites and will represent a total of 11 sectors which include: consumer cyclicals, heath care, technology, communication, financial, energy, capital goods, utilities, basic materials, transport, and consumer staples. Once the data has been downloaded, we will create data structures to hold the information that will be incorporated into a database which will serve as a back end to the visualization front end.

3.2 Visualization Application Development

In weeks three thru five, the visualization front end will be developed. The front end will contain two components; the stockcel grid and a graphical user interface. Each stockcel will be a certain number of pixels. This number will be determined based on the number of stocks in the database. The graphical user interface will provide the user with options for viewing closing stock prices, volume data, or price spreads within a given time interval in the past four years, and the user will be able to look at the data on a daily, weekly, or monthly scale. A slider will allow the user to view the data dynamically in the time interval. As the data moves, varying colors will represent changes in price, spread, or volume. Other features will be added time permitting.

3.3 Evaluation

In the last week of the project, we will evaluate the application by having a number of investors examine the system and answer a questionnaire as to the utility of the system, thus testing our hypothesis.

References

[1] Wattenberg, Martin. Visualizing the Stock Market, In *Late-Breaking Results, CHI'99*, ACM Press, 188-189, 1999. Available at http://www.smartmoney.com/marketmap/

[2] Yahoo Historical Quotes. http://chart.yahoo.com/d, 1999.

CV for Joseph J. LaViola Jr.

Personal

Name: Joseph J. LaViola Jr.
Email: jjl@cs.brown.edu

Education

Ph.D, Computer Science, Brown University (Expected May 2002)
Sc.M, Applied Mathematics, Brown University, (Expected May 2001)
Sc.M, Computer Science, Brown University, October, 1999
B.S., Computer Science, Florida Atlantic University, 1996

Employment

Brown University Computer Graphics Lab, Researcher, May 1998 - Present
Fraunhofer Center for Research in Computer Graphics, Research Scientist, May 1997 - December 1997
UCS, Inc., Software Technician, Summer 1996
International Business Machines, Computer Information Systems Dept., May 1995 - November 1995

Honors

IBM Research Assistantship(1999)
IBM Cooperative Fellowship(1998)
Elected as Associate Member to Sigma Xi(1998)
FAU's Aaron Finerman Award(1996)
FAU's Faculty Award for Outstanding Undergraduate Achievement(1996)
Microsoft Senior Achievement Award(1995)

Relevant Experience

Active Stock and Options Trader, Investment Researcher for 12 years

CV for Marc G. Majzner

Personal

Name: Marc G. Majzner
Email: mgm@cs.brown.edu

Education

Sc.M, Computer Science, Brown University (Expected May 2000)
B.S., Computer Science, Brown University, 1999

Employment

Brown University Graphics Group, Researcher, Sept. 98 - Current
Silicon Graphics, Engineer, Summer 1998 Brown University Graphics Group, Researcher, Sept. 97 - May 98
AT&T, Information Systems Manager, Summer 1997

Honors

National Honor Society President, 1994-1996

Relevant Experience

Active Stock and Options Trader
Heavy economics background

CV for Paul Reitsma

Personal

Name: Paul Reitsma
Email: psar@cs.brown.edu

Education

- BSc, Combined Honors in Computer Science and Mathematics, University of British Columbia, with Honors, 1998

Relevant Experience

- Research Assistant, Electronic Games for Education in Math and Science (E-GEMS), University of British Columbia, January to April 1997
  - Web-based educational game development for grades 5-7
- Research Assistant, TRIUMF National Laboratory, January to April 1995
  - Assisting Neutrino Group with software simulations

Honors

- Andries van Dam Fellowship

Skills

- Strong computer science background and programming experience
- Strong mathematical background
- Good technical writing background
- Basic grounding in most physical sciences

CV for Jonathan Reiter

Personal

Name: Jonathan Reiter
E-Mail: jr@cfm.brown.edu

Education

2000 A.B. Mathematical Economics, Brown University
2000 A.B. Computer Science, Brown University

Employment

1996-98 Research Assistant, Thomas Jefferson University Hospital
1998-present Research Assistant, Brown University

Related Work

Testing Monte Carlo Option Pricing Methods in the 90s Market, Brown University Department of Economics

Visualization of 3D Protein image By Emphasizing the uncertain factors

pi: Quan Gu co-pi: Pramod Paranthaman

Consultant: Dongbai Guo(computer graphics)

Abstracts

Visualization of the chemical structures three dimensionally can help scientists from different disciplines to communicate ideas more efficiently. But the current visualization techniques fail to illustrate uncertain factors like conformation changes, hydrogen bonding migration, etc. To illustrate these factors in a static picture, strategies borrowed from sports photographic techniques, oriental paintings will be utilized. To interactively visualize the motion of the chemical structures, animation techniques will be utilized.

1. Aims

To visualize the chemical structures in 3D as real as possible is an ongoing pursuit of many scientific software companies and scientists. Ignorance of the uncertain factors is the critical shortcoming of the current softwares. Novel ideas are proposed here to emphasize these factors in both static picture and interactive visualization.

Though not all the proposed ideas will be fully explored due to time limitation, the results from this research should provide the insight into the technic trends of the next generation of 3D chemical structure visualization software. Some detailed study of the visualization techniques burrowed from other disciplines might can be applied directly to other scientific visualization instances.

2. Significance

To imagine the molecular structures by brain requires in depth understanding of the chemistry and high IQ. in understanding the 3D objects. For most of the scientists and students, see the 3D picture or interactive visualization of the molecular structure is the only way to get an insight of the molecule.

Unfortunately, currently available visualization softwares failed to show the uncertain factors in the molecules. Which often mislead scientist from multiple disciplines like chemical engineering, molecular biology, drug designing, molecular modeling, etc. to wrong conclusions.

Protein, also called polypeptide which is a very complex molecule made up of amino acid units. Enzymes are the proteins function as bio-catalyst. Conclusions drawn from the 3D visualization without considering the conformation changes or the variation of the hydrogen bonding are often wrong, which means the research on HIV , cancer be lead to wrong directions. So use protein as the visualization subject is interesting and necessary.

3. Comparison with the Related Works

The currently available 3D molecular visualization softwares are quanta, cache, chem3D, etc. All these softwares fail to show the (1) molecular rotation related variations (2) hydrogen bonding variations of the molecules. Therefore the resultant 3D static images are conceptually wrong. One goal of the proposed research is to give live to these “dead” images, so they can “fly” on the paper.

To realize the above goal, strategies burrowed from (1) available computer visualization techniques like “Conveying the 3D Shape of Smoothly Curving Transparent Surfaces via Texture” (2) photographic techniques like multiple exposure and distortion utilized in sports photos to show motion (3) oriental painting techniques like using contrast and simple contour to suggest the motion of water will be applied.

Historically, almost all the computer molecular visualization technologies are the implementation of the pre-exist molecular visualization methods. For example, the stick and ball model is derived from molecular model; the ribbon representation is created by artist first. Knowing what the molecules should looks like is more important than how to do the programming in computer molecular visualization development. So the main focus of the research will be on the effect (how people will think about the created pictures), rather than how to actually create a picture.

The other disadvantage of the current available software is that only one variable like one bond angle can be changed at one time. So the demonstration like the protein de-naturalization can't be explored by them. Another goal of my proposal is to allow the user to change multiple values of a molecule at one time. The realization of this seems to be relatively easy; write a temp file for the changes first, then modify the data for the molecule accordingly. A kind of animation style visualization can be achieve by this part of the work. It will allow the scientist to "play" with the molecule much more efficiently and intuitively by computer than the actual models for the first time.

4. Work Plan

Apparently, this is a proposal with huge ambition. I'm sure that I CAN'T finish everything I proposed here on time, but the result will be fruitful after all.

In the first two weeks, the main focus will be on the comparison of different visualization methods. Commercial available 3D molecular visualization softwares like cache, chem3D, quanta will be used to generate 3D images to begin with, and image edition softwares like photoshop will be used to create the effect.

In the third week, discussion will be carried to determine the future direction of the project. One possible direction is to spend the rest of the time to write a paper about our insight of the future for molecular visualization. The other possibility is to actually implement at least one of our ideas.

In the week 4-5, either coding is going on by using SGI stations or paper writing is going on with intensive interaction with experts in the field.

In the week 6, results will be wrapped up. Need beers for celebration.

Quan Gu

Chemistry Department

Brown University

Skills

- Expertises in chemistry at Ph.D level. Especially the understanding of complex molecules called polymer.
- Extensive knowledge of biology, physics and other natural sciences. Knowing how scientists from other field utilize chemistry.
- Some CS course and programming experiences. Understanding the CS terminologies, can write simple programs.

Experiences

- Conducting chemistry research.
- Wrting research proposal, papers about chemistry.
- Using commericially available molecular visualization softwares to drawing chemical structures for publications.
- Taken CS4, CS252, CS176, CS127 courses in brown.

Pramod Paranthaman

Computer Science Department

Brown University

Education and Professional Experience:

Brown University, Master of Science. Expected May 2000

Concentration: Computer Graphics

Worked at the Virtual Environment Navigation Lab (VENLAB) in the Cognitive and Linguistic Sciences Department as a graphics programmer.

Relevant Courses: Interactive Computer Graphics, Introduction to Computer Graphics, User Interface Design

American Megatrends Inc.(AMI), Madras, India

System Software Engineer, 1997-1998

Worked on the design and implementation of various large software projects in the BIOS division at AMI.

College of Engineering, Guindy, Anna University, Madras, India

Master of Computer Applications, 1997

Relevant Courses: Interactive Computer Graphics, Computer Aided Design, Software Engineering

Loyola College, University of Madras, India

Master of Science (Physics), 1994

Developed CrysMols, an interactive teaching aid for visualizing crystal and molecular structure for the Department of Chemistry.

Relevant Courses: Organic Chemistry, Quantum Mechanics

Relevant Computer Skills:

Extensive programming experience in C, C++, and Java.

Experience in WorldToolkit, Java3D and Open GL.

Dongbai Guo

Education:

Brown University, Providence, Rhode Island, USA

Ph.D. in Engineering, expected January, 2000

M. Sc. in Computer Science, May, 1999

M. Sc. in Engineering, May, 1998

Current Research:

Angiogram and intravascular Ultrasound Video Processing (Ph.D. thesis)

Research Experience:

Engineering Division, Brown University, Providence, Rhode Island (January, 1996 ~ present)

Medical video segmentation and visualization

3D vessel reconstruction from EKG gated single plane angiograms recorded in vivo

Ultrasound video segmentation

Time-variant 3D volume reconstruction and visualization

Bell-labs, Lucent Technology, Murray Hill, New Jersey (May 1999 ~ September, 1999)

Real-time facial animation through low bandwidth network

Real-time realistic text-driven facial animation

Facial model compression with implicit polynomials

Graphics Lab, Computer Science Department, Brown University (September, 1997 ~ May, 1999)

Automatic image mosaic: Robust mosaic from images without any adjacency information

Shanghai Jiaotong University Artificial Intelligence Lab, China (September, 1994 ~ July, 1995)

Chinese language processing: Automatic keyword extraction from articles

Technical Experience:

Programmer: Brown University, Providence, RI, NSF Center for Graphics and Visualization

Java applet for 3D interactive illustration of Phong illumination on web, summer 1997

Consultant: Brown University, Engineering Division, Prince lab computer cluster,

UNIX and Windows NT system, 1996, 1997, 1998

3D Vessel Reconstruction from EKG Gated Single Plane Angiograms

Principle Investigator Dongbai Guo,
Co Principle Investigator Tim Rowley, Song Zhang
Consultant (tentative) David Williams

Abstract—**We propose a system to construct time-variant 3D angiogram from EKG gated single plane angiograms. Our extension to the existing vessel reconstruction methods allow 3D construction from clinically widely used single plane angiogram. The frame correspondence required by the 3D reconstruction process is achieved by EKG signal input recorded along with the single plane angiogram video.**

Keywords— **EKG Gated Angiogram, 3D vessel reconstruction, in vivo.**

I. Introduction

Angiogram has been used widely clinically for diagnosing acute heart diseases such as atherosclerosis. X-ray camera images the contrast agent introduced into a blood vessel by a catheter. The radio opaqueness of the vessel image provide diagnostic cues to a radiologist. However, the use of single plane angiogram alone proved to be unreliable due to natural ambiguity arises from the single projection view of a 3D elongated structure. More views or additional imaging methods are required in order to accurately image the deceased vessel. Although through extensive training, an angiographer can learn to mentally reconstruct the three dimensional relationships of the coronary arterial shape, it is often more desirable to reconstruct and visualize the 3D geometry.

Previous work in vessel detection and reconstruction is limited to in vitro bi-plane angiograms. We propose a new method that uses EKG gated angiogram to reconstruct 3D vessel structure from two different single plane angiogram views. A prelimilary of sketch of the system is shown in figure 3.

II. Protocol

An X-ray camera is first calibrated with standard camera calibration procedure. EKG signal is then recorded along with angiogram images for 2 heart cycles. The camera is then rotated to a position (90 angle) to repeat the capturing process for two more cycles.

III. Camera Calibration

To link the projected views of the angiograms together and perform 3D vessel reconstruction, one must calibration the X-ray camera first to establish the mapping between 2D views and 3D reconstruction.

Many camera calibration methods with varies degree of automaticity has been introduced over the past year. Our choice of camera calibration is based on the following criterion: a) code availability, b) accuracy. Two methods fit this category: 1) Tsai's [1] linear perspective camera calibration model, implemented by Willson [1]. 2) An improvement of Tsai model [2], [3], developed and implemented by Heikkila and Silven. We expect both systems to be sufficient for our purpose. They should require a small amount of time to set up. Having two methods not only provides us backup but also allow comparisons impossible otherwise.

IV. Frame Correspondence from EKG Gated Angiograms

For biplane angiogram, two views of the same 3D structure are captured virtually at the same time, from which one may construct the 3D structure for a specific time. At any time, we have a pair of images which correspond to two views of the same 3D structure.

On the other hand, single plane angiogram can only capture one view at a time. The capture of the second view involves rotation of the imaging device which has significant lag that can not be ignored, thus addition assumptions are necessary to solve the computer vision problem. Based on these assumptions, we may find the corresponding pairs from two sequences of angiograms.

In this case, we assume there is little overall motion of patient during the imaging process (typically takes less than 1 min) and the motion of the cardiac blood vessel is cyclic. In other word, one observe the same repeated motion during the entire capture process. The period of this motion, is fortunately, accessible from EKG. Most angiogram imaging system allow capturing of the EKG signal, thus we may stamp every frame of angiogram video with EKG feed and thus determine the phase of heart beat for each frame. Two angiogram image sequences can then be paired up with this feed.

Note that the global rigid motion can be removed through point correspondence and camera calibration process, thus we may relax the first assumption to small rigid motion.

V. Image Feature Correspondence, Graphic User Interface

Once images are paired up, we may start to reconstruction the 3D geometry. The first step of reconstruction is

D. Guo is with the Brown University, Engineering Division, Providence, RI02912, USA. E-mail: dbg@cs.brown.edu

[1]http://www.cs.cmu.edu/afs/cs.cmu.edu/user/rgw/www/TsaiCode.html

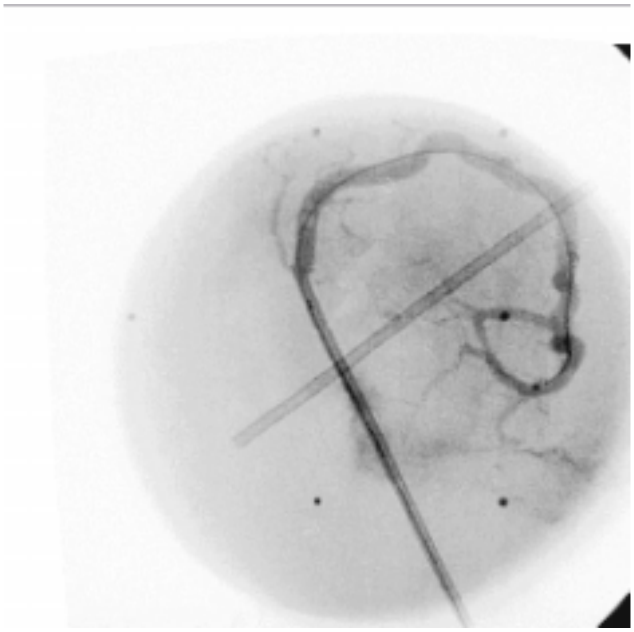

Fig. 1. One projected view of a blood vessel

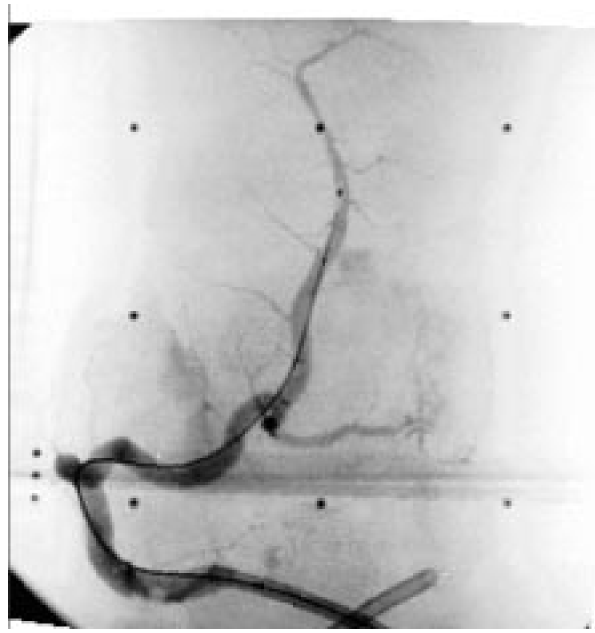

Fig. 2. A second projected view of the same blood vessel as figure 1. Notice the difference in overall image contrast and brightness and local variation in image intensity

to build point correspondences from two views of the same geometry. Such correspondences may be established automatically for simple cases by using computer vision techniques, how it is more difficult for angiogram images (see figure 2,3). First of all, even for biplane angiogram, two projections of the same 3-D vessel segment may not have the same image intensity, therefore common image intensity based correspondence matching algorithm will not work in this case. Secondly, two corresponding projections are taken at different times, the amount of the manually injected radio-opaque contrast agent for the two imaging sequences generally varies from time to time. In addition, the phase of the injection can not the controlled either. Therefore, the result images can not be fused by simply solving the correspondence problem, a manual matching process is necessary to build the image feature correspondence.

Our goal is to build, on top of existing interactive user interface (Jot, developed by Computer graphics lab of Brown University), an interactive drawing tool to match up two projections of the same vessel. This can be achieve with relatively small effort due to the large amount of existing library. Anatomical landmarks such as branches and crossing can be used to align two corresponding frames. In addition, we also mark out roughly the region of interest for the blood vessel segment that we want to reconstruct.

VI. Vessel center line detection

The 3D geometry reconstruction is based on 2D vessel center line detection, which generate parametric lines from each view of the 3D blood vessel. This is an image space operation.

We used a central detection method [4] that is based on the shortest weighted path (as in graph theory). A region of interest is first marked in the GUI as suggested in the previous section. We then generate a graph to describe the connectivity of the region, where each pixel corresponds to a vertex of the graph and an edge join two vertices indicates that the two pixels are adjacent. Edges are weighted by their lengths and average image intensities. We used standard graph theory algorithm to search for the shortest path joins the start and the end of vessel segment that we want to build 3D representation.

Other methods, such as ridge detection, modified snake, exist for vessel central line extraction. The above method has the advantage of easy to implement and relatively fast to compute, especially we have already marked the region of interest before-head thus dramatically decreases the computation cost.

VII. Vessel center line 3D reconstruction

The detected vessel center line in each view plus feature point correspondence between the views allow the reconstruction of 3D vessel center line. We use linear interpolation to interpolate correspondence a line segment bounded by two marked feature points. Since the correspondence is not exact, they does not exist a 3D point whose projections match the inputs. Therefore an optimization step is necessary to choose the best path. Existing methods including snake and nonlinear optimization, they are straight forward to implement.

The center line provide a 3D skeleton, we add additional information by computing the cross sectional shape from each views, so that to build a tubular geometry resample that of a real blood vessel.

VIII. Vessel Cross-sectional Geometry

To generate the cross sectional shape from the two projected views, one needs to estimate the scale factor of the perspective projection. This is described in the camera calibration section. In addition, the detected vessel geometry

also depends on the amount of the contrast agent injected in the vessel and its dissipation speed. It is very difficult to construct the right geometry; however, one may get a good approximation by normalize the vessel size with respected to the proximal end of the vessel, where there is in general very little chance of arterial decease and the cross sectional shape is very often circular. Clinically, the exact geometry of the vessel is not of utter most importance, since the abnormal narrowing and tapering what we want to look out for as they are signs for serious deceases. Therefore we may simplify the cross-sectional shape to ellipse and generate a quick reconstruction of geometry.

IX. Curve Coordinate System and Triangulation

The cross sectional curves may be joined together and triangulated to generate triangular model that is suitable for graphics hardware. To do this, we need to build the correct curve-linear coordinates. This problem can be solved by using Frenet formula as suggested by Leban *et al.* [5]. The code for this coordinate has been partially implemented in Jot. The reconstructed cross-sectional geometry and the curve-linear coordinates define the 3D vessel shape unambiguously. We may then triangulate this shape into triangular strips for visualization.

X. Evaluation

Our consult, Dr Williams, who is an experienced cardiologist, will provide suggestion for GUI and will evaluate the accuracy of the reconstruction. In addition, we will also test the accuracy of 3D reconstruction with radial opaque calibration object in the shape

XI. Software packages

We will use the following software packages:

- **Jot, a 3D C++ software package for interactive graphics.**
- **Matlab, a commericial package for matrix computation**
- **Camera calibration package, in Matlab and C.**
- **Leda, data structure and graph theory package.**

The principle investigator has experience with all the above packages. All of then has been installed/downloaded to computer science department here at Brown. Fulltime researchers in Brown will also provide technical help regarding to questions related to Jot, which is not as well documented as others.

XII. Work plan

The amount of work requires three programers with advantaged degree, with background in computer graphics and software engineering. The major tasks are described belong: a) The data acquisition and camera calibration can be done within 2 week. b) GUI tool may also be finished within two weeks by a programmer who is familiar with Jot architecture. c) Vessel centerline detection will take up to 2 weeks of programming. d) Vessel centerline reconstruction will take 2 weeks to accomplish. e) Vessel triangulation may be done within two weeks. f) Vessel cross-sectional geometric may require up to 2 weeks of programming, depends on the simplification assumption one will make. g) Software merging will take up to 1 week of time. h) software design will talk 1 weeks. i) project write up and presentation preparation will take 1 week. j) reference lookup will take 0.5 weeks. k) evaluation and adjustment will take 1 week.

Since this is a research project, we will not merge all the code into a single software due to the difference of the underlying modules.

The detail of the time schedule is shown in the table I. Note that *s*1, *s*2, *s*3 indicate three streams, not students.

XIII. People and Resources

Dongbai Guo, the principle investigator, has been working on medical image processing for three years. He has extensive experience in image processing and computer vision and will coordinate the research. Tim Rowley is a fourth year Ph.D candidate of computer science department and has extensive experience in computer graphics and software engineering. Song Zhang is a second year Ph.D candidate of computer science department and has more than one year of research experience in scientific visualization. All three investigators will work fulltime on this project. Guo will mainly work on image processing and computer vision modules of the program, Rowley will work on graphics and interaction module and Zhang will work on visualization module.

Our outside consultant, Dr. David Williams will provide angiogram data. He will also assign technicians for helping with camera calibration and image capturing.

References

[1] R.Y. Tsai, "A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses," *IEEE Journal of Robotics and Auomation*, vol. 3(4), pp. 323–344, 1987.

[2] J. Heikkila and Silven O., "Calibration procedure for short focal length off-the-shelf ccd cameras," in *Proc of The 13th International Conference on Pattern Recognition*, 1996, pp. 166–170.

[3] J. Heikkila and Silven O., "A four-step camera calibration procedure with implicit image correction," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1997, pp. 1106–1112.

[4] Wahle A, Oswald H., and Fleck E., "3-d heart-vessel reconstruction from biplane angiograms," *Computer Graphics and Applications - Applications in Surgery and Therapy*, vol. 16(1), pp. 65–73, 1996.

[5] Leban M. and *et al.*, "Angus: A new approach to three-dimension reconstruction of coronary vessels," 1995, vol. 22, pp. 325–328.

Fig. 3. 3D reconstruction of vessel geometry from two views

TABLE I
TIME TABLE.

	0.5 week	2 weeks	2 weeks	0.5 week	1 week
s1	reference	camera calibr.	cross section	soft. merge	
s2	acquisition	GUI	triangulate	evaluation	visualize
s3	design	vessel extract	vessel constr.		write up

Extending a Painting Metaphor for Visualizing Flow Fields

Participants:

Daniel Keefe (PI)
David Laidlaw
Song Zhang
Paul Reitsma
Marco DaSilva

Consultants:

Michael Kirby

Abstract:

This work builds upon previous research by David Laidlaw by proposing a more realistic model of the “brushstrokes” used to create his flow field visualizations. By capturing more of the expressive power of a painter’s brushstroke, will be able to improve upon previous visualization techniques and explore possibilities for future techniques, which take advantage of artistic concepts to produce simple, intuitive, and expressive images.

Introduction:

In the field of fluid dynamics, multi-valued datasets are common. For two-dimensional (2D) flow fields, each point in the field often has several data values associated with it. These data often include properties, such as velocity, which are determined directly from experimentation or computational analysis, as well as quantities, such as the rate of strain tensor, turbulent charge, and turbulent current, which are derived values. This vast amount of data poses an intriguing problem for data visualization experts. Researchers in fluid dynamics like to be able to see all of the data on a particular 2D flow field (up to eight data values for each point in the field) at once. This ability enables them to better understand the flow they are studying. It also helps them to draw conclusions about the relationships between the different data values within the field. The problem for visualization experts is that, unfortunately, useful visualizations of multi-valued datasets are very difficult to create. Traditional visualization strategies are, often, not appropriate for this type of data, since there is simply, too much data for each point to fit it all in one composite image. Visualization experts need to find a way to take this complex data and transform it into a simple, expressive image, which accurately captures the meaning of the data, but presents it in a simplified, expressive manner.

There is a discipline, which has been successfully solving this problem for thousands of years. This discipline is art. The power that artists have to simply express complex scenes or ideas is largely unrivaled. In the 2D case, painters like Van Gogh and Monet are able, through their varied use of brushstroke, color, and composition, to express even the most complex variations in light. It is amazing that we have no trouble comprehending the multitude of ideas that they have projected, with the aid of only a paintbrush, onto their canvases. This is the power of art.

Laidlaw [1] recognized that this power could be applied to the problem of visualizing mult-valued data sets. Specifically, he applied several concepts from painting to create impressive 2D visualizations of flow fields. This proposed project builds upon Laidlaw's initial work. The hypothesis we hope to test is that by defining a more realistic approximation for the "brushstrokes" used in the original visualizations (essentially by staying closer to the painting concepts this work is based upon), we will be able to create more compelling, expressive, and valuable visualizations. A new, more realistic brushstroke model will potentially make understanding the visualizations created much more intuitive and provide the ability to encode even more data for each point in the flow field.

Related Work:

Laidlaw [1,2] recognized the promise that art holds in solving the problem of visualizing multi-valued data from flows. His work focused on modeling the painting techniques of layering and varying "brushstroke." However, the brushstrokes used in this initial work were actually just simple shapes: triangles and ellipses. The dimensions and orientation of these “strokes” were varied to create a visualization of flow data. While this technique shows promise, it tends to handcuff the expressive power of a brushstroke that Van Gogh and Monet demonstrate in their work.

Curtis [3] developed an extensive and impressive simulation of watercolor painting. This work was not applied to any scientific problem, however. Others [4] have done similar work simulating oil painting techniques. This work tends to focus on rendering images, which have been scanned into a computer, in ways which simulate oil painting. Very little work has been done to take advantage of artistic power of expression in scientific applications of computers.

Work Plan:

The first step in this proposed project is the most important. It is research. Since this work focuses on improving an existing visualization by taking advantage of painting techniques, it is imperative that we pick the most promising techniques to model. This will involve going to museums to look at painting, actually doing some oil painting, talking to artists, and reading. Initial work suggests that brushstroke characteristics such as varying length, varying width along the stroke, color, paint streaking, and fluidity of the paint applied may all be useful characteristics to model. Data could be encoded in all of these parameters.

The second step of this work is to actually implement a computer-based model of a brushstroke. Here, we will take the defining characteristics of a stroke that we feel are most promising from a visualization standpoint and create a computer model of a stroke using these characteristics as inputs.

The third step is to make the stroke model work with some data. Note that this step inherently involves refining the model created in step two. There are two interesting different directions we could take at this point. The first is to apply the stroke model to data which was already visualized using Laidlaw’s method. A comparative study could be done between the two methods to determine the strengths and weaknesses of each. The second, and perhaps more intriguing direction to take, from a research standpoint, is to apply the stroke model to 3D data. Depending on the specifics of the model, the long, flowing strokes that we envision modeling with this project may lend themselves to a 3D application. The development of a 3D visualization technique for this type of data would be a significant research advance.

At key junctures throughout the project, a component of the work will be the analysis of the technique by flow experts. Mike Kirby will be providing expertise on flow visualization interpretation for this project.

We expect the first step of this project (research) to take approximate one and a half weeks. Some members of the team may be able to proceed with the second step of the project during this time, so we anticipate its completion about a week after the first step is complete. The next three weeks will be spent connecting the stroke model to some real data and refining the model as necessary. Again, the way that this is carried out may depend on whether the brush model lends itself to a 3D application. There will be approximately one week left in the six week timeframe to prepare a final abstract and presentation of the project.

Collaborators / Facilities:

A team of several people will be working on this project. As the principle investigator, I will manage this group and have several other responsibilities regarding the implementation of the project. David Laidlaw will be providing expert advise on the visualization of multi-value datasets as well as research directions for future work related to this 6-week project. Song Zhang will be providing our team with programming environment help and will serve as a resource for the *jot* programming toolkit. Paul Reitsma will focus on the extension of the project to 3D. Marco DeSilva will focus on the actual implementation of the brush model specification. Fiinally, Mike Kirby is a flow expert who will be available to the project for analysis of the visualization techniques. Several artists will also be involved in the project as it progresses.

Facilities available to this group at Brown University's Computer Science Department will be more than sufficient for this project. The project may make use of jot, a 3D programming toolkit created at Brown University, which will enable results from this project to be viewed in Brown's CAVE.

Significance of Work:

This work is significant because it addresses a problem which is present in many different fields. The inability to accurately visualize multi-valued datasets is a problem in fluid dynamics, medicine, and many other scientific areas. This problem will only increase as technology improves and the ability to collect even more data increases.

Summary:

This project proposes an extension to work completed by David Laidlaw on visualizing multi-value datasets using techniques from painting. Painting, and art in general, provide us with many powerful concepts which we can use to improve our computer visualization techniques. In this project, we will develop a new, more realistic computer model of an oil-painter's brushstroke. Through the parameters which characterize this brushstroke, we will encode data describing flow fields. We will then analyze this painterly technique to determine whether it indeed shows promise for either a 2D or 3D scientific visualization of flow fields or other multi-valued datasets.

References Cited:

1. David Laidlaw, David Kremers, Eric T. Ahrens and Matthew J. Avalos. Visually Representing Multi-valued Scientific Data Using Concepts from Oil Painting. *SIGGRAPH '98 Visual Proceedings* (Sketch #249).
2. R. Michael Kirby, H. Marmans, and D. Laidlaw. Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting. In *Visualization '99 Proceedings*.
3. Cassidy Curtis, Sean Anderson, Joshua Seims, Kurt Fleisher, and David Salesin. Computer-Generated Watercolor. In *SIGGRAPH '97 Proceedings* pages 421-430.
4. Aaron Hertzmann. Painterly Rendering with Curved Brush Strokes of Multiple Sizes. In *SIGGRAPH '98 Proceedings* pages 453-460.

An Analysis of Depth Cues in Three Dimensional Virtual Spaces

Principal Investigator: Danah Beard
Co-Principal Investigators: Leslie Welch, Avi Walsky, Matt Hutson, Benjamin White
Consultant: Bill Warren

Abstract

By effectively analyzing how various visual cues affect an individual's awareness of depth in a virtual system, we can better develop virtual systems that meet the needs of the user. In addition, we can advance the current base of knowledge regarding vision and the potential sex differences that may exist in cue prioritizing. This research proposes looking into the visual mechanisms that affect depth perception by using three dimensional computer systems that have controlled cues available to the user to determine depth.

Introduction

Detecting depth from three dimensional virtual systems is dependent on the amount of cues that the program makes available to you. Most systems use optic flow – objects getting larger as you get closer. Some use texture cues to give a better understanding of the objects' properties and distance. Very rarely, binocular disparity cues are given. This is done through stereoscopic glasses that are frequently used in virtual reality rooms. One cue that humans have in real space is rarely used – shape from shading.

Shape from shading is a cue that is dependent on the lighting in the system and is heavily integrated with optic flow and binocular disparity. This cue allows the individual to get information about the object by detecting the subtle changes in reflected light dependent on the eye's given position in relation to the object and the relationship between the object and the light source. By having access to two different shape from shading cues at any given point (stereoscopic vision) and by constantly moving relative to any object, the brain uses this cue to get depth information.

Since shape from shading is heavily coupled with both optic flow and binocular vision, limited amounts of research have been done on just shape from shading alone. Also, while research has been done decoupling stereoscopic cues from shape from shading cues, optic flow has never been measured (Cornilleau-Peres, 1993).

Goals

The primary goal of this project is to determine what significance shape from shading has in determining depth information. In particular, we want to know how this affects individuals, particularly along sex lines. In this manner, we may be able to determine whether or not this is an essential cue for individuals to accurately detect depth.

In order to do this, we will develop static and interactive systems to be used in virtual environments which will allow us to do user studies to determine what type of cues are necessary for individuals to have acceptable human-computer interaction in a three-dimensional environment. In addition, we will provide examples of environments which are contradictory to anything that could appear in real space. This will also allow us to determine which cues are necessary for depth detection and how important it is that these cues are available.

Significance of Work

Currently available three-dimensional systems acknowledge that there are a variety of cues used to develop a clear understanding of the depth of an environment. This does not mean that they adequately use them. Most systems leave out shape-from-shading as an unnecessary, extraneous cue. If, in fact, this cue does more than help the optic flow of the system by becoming a primary cue for some people, it will then become essential that developers of 3D systems actively use this cue.

Based on previous research, there is a strong possibility that individual differences will fall along sex lines. This would prove significant because it may begin to explain some of the difficulties that women appear to have using currently available 3D systems which lack this cue.

The implications of this work are broad reaching. Conclusions developed during this study will dramatically affect research in cognitive science, psychology, neuroscience and endocrinology. In addition, the results may affect the perceptual cues used in developing virtual reality systems, games and other 3D applications. By developing a system that accurately activates the user's visual system, a human-computer interface can be developed that is effective and useful.

Related Work

This proposal is based upon research that I completed in Amsterdam. While in Amsterdam, I determined that visual perception differences do develop as an individual undergoes sex hormone transformation. This information was gathered through self-acknowledgement. It is well known that the retina is bathed in sex hormones (Lanthier, 1988) but little is known of how the hormones affect the vision. Since depth perception is one thing that appeared to change within the transsexuals, it is only appropriate to determine what differences, if any, exist.

Coupled with the awareness that women have a more difficult time performing depth tasks in video games (Cassell, 1999), it appears that cues may be lacking in virtual spaces that are not lacking in real spaces.

Work Plan

By the end of the first week, a set of images and interactive programs will be available. These are the rough drafts of what will be used during the human experiment.

The second week will be dedicated to analyzing the test data for experiments. This will be done by running the test on random friends and ourselves, to determine if the tests are too easy or too difficult. The participants in this round of work will be individuals with cogntive science background who are aware of what we are attempting to measure. This interaction will allow us to develop a good balance of what is effective for the tests.

During the third week, we will modify the images and prepare for a final round. By this time, all participants will be scheduled for testing during the fourth week.

The fourth week will be spent running tests with participants. At this point, it appears as though 20 participants are appropriate for the first round (10 female, 10 male).

The fifth and sixth weeks will be spent analyzing the data in the context of the readings and creating an abstract of a paper that accurately describes what has been learnt. In addition, I will prepare a presentation of the work that has been done.

Collaborators

This work will be completed by multiple participants. As the principal investigator, my primary role will be to organize the development and systematically manage this project. In addition to this role, I will be actively working on the implementation of the system as well as guaranteeing that the experiment is methodologically in tact.

Leslie Welch and Bill Warren will lend their experimental expertise and knowledge in assuring that any experiment that is run is well prepared and without methodological defects. In addition, they will continue to act as mentors to improve my background in cognition. By having access to both of their expertise, it is more probable that this project will appeal to both cognitive scientists as well as psychologists.

Benjamin White will be active in preparing the software that creates the necessary images and animations for the experiments. His background and expertise in graphics programming will be quite useful.

Avi Walsky and Matt Hutson will lend experimental knowledge through their expertise in neuroscience. In addition, their knowledge of the previous work and experimental difficulties will become quite useful.

Facilities

For this project, I will need to use at least 20 human subjects. These subjects will be acquired from the Psychology department. Brown University's Psychology 1 class requires students to participate in at least 3 psychology-related studies and the Professor running this class has accepted this proposal as one of the research projects which Psych1 students may do for course credit.

The appropriate forms for using human subjects have been submitted.

Brown University has created two new virtual reality laboratories which will be of great use for this project. In addition to doing tests on monitors, using the Virtual Environment Navigation Lab (VENLab) or the Cave system will be of great use to complete this project. These labs allow for immersive environments, which will enable me to research this material in a manner that is as similar to the real space as possible.

In addition to these items, I will need to have regular access to high-powered UNIX machines. This should be available by using the Sunlab within the computer science department.

Summary

Being aware of how various cues affect visual performance is key to developing software and hardware systems that can effectively perform depending on the user's capabilities and strengths. Without this awareness, virtual reality will not succeed in becoming a commonly used technology. In addition to the technological effects, having the ability to understand this section of visual system, particularly with regard to the potential sex-based disparity could have strong implications for the cognitive science, psychology, neuroscience and endocrinology worlds.

This research also allows for direct collaboration between cognitive scientists and computer scientists in order to develop systems that are more useful for the general public while advancing our understanding of the brain and visual system.

References Cited

Beard, Danah (1999). Sex-based vision through the eyes of transsexuals undergoing hormone treatment. School for International Training, unpublished.

Cassell, Justine (ed.) (1999). *From Barbie to Mortal Kombat: Gender and Computer Games*. MIT Press: Cambridge, MA.

Cornilleau-Peres, V (1993). Stereo-motion cooperation and the use of motion disparity in the visual perception of 3-D structure. *Percept Psychophys*, 54(2), 223-39.

Lanthier, Andre and Vasant Patwardhan (1988). Vasant In Vetro steroid metabolism by rat retina. *Brain Research*, 463, 403-406.

Danah Beard, Principal Investigator

Birthdate: November 24, 1977
Place of Birth: USA
Current Status: Researcher, Brown University
Home Address: Box 0556, Brown University, Providence RI 02912
Telephone: 401.861.0131
Email: danah@acm.org
Website: http://www.danah.org/

Education:
- 2000 A.B. in Computer Science, Brown University, Providence RI
- 2000 A.B. in Computational Gender Theory, Brown University, Providence RI

Research and/or Professional Experience:
- Software Engineer, Macromedia Generator, Macromedia, San Francisco, CA, June 1998-September 1999
- Research Assistant, Sociable Media, MIT Media Lab, Cambridge, Mass, June 1999-present
- Research Assistant, Welch Vision Lab, Brown University, Providence, RI, September 1999-present

Related papers:
- Beard, Danah (1999). Sex-based vision through the eyes of transsexuals undergoing hormone treatment. School for International Training, unpublished.

Relevant courses at Brown:
- Computer Science 32, Software Engineering
- Computer Science 123, Introduction to Computer Graphics
- Psychology 185, Motion Perception

Capabilities and skills:
- Computer science background and strong coding background (C/C++/Java)
- Computer graphics knowledge
- Directed knowledge of vision- motion perception and depth perception

Avi Walsky, co-Personal Investigator

Birthdate: December 27, 1977
Place of Birth: USA
Current Status: Researcher/Student, Brown University
Home Address: Box 5701, Brown University, Providence RI 02912
Telephone: 401.831.8986
Email: awalsky@cs.brown.edu

Education:
- 2001 Sc.B.in Visual Science, Brown University, Providence RI

Research and/or Professional Experience:
- Developed interactive web pages(1999) - A Good Thing Inc.

Relevant courses at Brown:
- Computer Science 123, Introduction to Computer Graphics
- Neuroscience 1, The Brain: An Introduction to Neuroscience
- Neuroscience 101, Neural Systems
- Neuroscience 102, Principals of Neurobiology
- Cognitive Science 102 - Neural Modelling Laboratory
- RISD FormZ - 3D graphics

Capabilities and skills:
- Strong neuroscience background
- Computer science background moderate C++/Java programming
- Computer graphics knowledge (both as user and programmer)
- Strong artistic background in 1,2 and 3 point perspective as well as other artistic techniques in creating the illusion of depth

Benjamin White, co-Personal Investigator

Telephone: 401.867.6656
Email: bwhite@cs.brown.edu

Education:
- 2001 Sc.B. in Neuroscience, Brown University

Employment/Research Experience:
- Brown University Computer Science Department, 1999-present
- Software Engineer, White Oak Technologies, Inc., 1999-present
- Software Engineer, Lighthouse Technology Solutions, Inc., 1997-present
- Lab Technician, Holy Cross Hospital, 1996

Relevant Skills:
- Strong coding experience (awk, C, C++, Java, Perl)
- Graphics knowledge
- Good knowledge of visual perception

Collaboration Statement from Leslie Welch

There are many cues to depth in the real world but human observers are not uniformly sensitive to all of them. When designing a 3D virtual space, it can be important to know which cues help users perceive depth and which can safely be ignored. Danah's research will investigate two cues to depth: shading and motion parallax. She will examine how sensitive human observers are to each depth cue individually, together and in conflict with each other. Research on other depth cues has shown quite striking individual differences in relative sensitivity and this is an additional aspect that Danah will be studying.

I will happily contribute my knowledge of visual perception and experimental design to the project. This is an interesting examination of how much of the real world needs to be translated into virtual environments.

Leslie Welch, Department of Psychology

An Analysis of Depth Cues in Three Dimensional Virtual Spaces

Principal Investigator: Danah Beard

Issues under question

- There are multiple cues to understand depth in 3D virtual systems (optic flow, shading..)
- People may rely on these cues differently; this dependence might be sex aligned
- Virtual reality systems use limited cues to provide depth information which could be limiting for individuals who rely on those

Hypothesis and Effects

- Women are *probably* more likely to depend on shape-from-shading cues while men are more likely to depend on optic flow

- This *may* explain why sex hormones are present on the retina and why sex-based performance differences exist in virtual 3D systems

Process of Project Completion

- Create a system to test individual’s cue preferences through conflicting cues
- Run user studies with male and female Brown students
- Analyze collected data in conjunction with available literature

Project Milestones

1: Develop test stimulus

2: Analyze stimulus and run test experiments

3: Evaluate and reorganize stimulus

4: Run user study with 20 individuals

5: Analyze data in reference to previous work

6: Construct paper, presentation

Participants and Roles

PI: Danah Beard - organize and implement

Co-PI: Leslie Welch - experiment knowledge

Co-PI: Matt Hutson - cogsci/neuro details

Co-PI: Avi Walsky - neuroscience details

Co-PI: Ben White - programmer

Consultant: Bill Warren - critique / advise experiment

ENCODING INFORMATION OF HIERARCHICAL DATA STRUCTURES

PI: Benjamin White
Co−PI: Danah Beard
Consultant: David Hornug

Abstract

Because it is so ubiquitous in daily life, the assimilation of data arranged in a hierarchical structure is an important problem. Various forms of graphical data presentation exist that take various approaches to presenting the data to the user. In this proposal we aim to develop a better technique to dealing with data trees.

Introduction

Human beings are often required to assimilate information which is presented to them in hierarchy form. A classic tradeoff in visually representing a large mass of data already arranged in a tree form is specificity for structure. That is, a given application may concentrate on making available a relatively large amount of data that is specific to a relatively small subset of the data tree. Alternatively, it may concentrate on conveying the more macroscopic characteristics of the entire tree at the expense of specific information regarding a much smaller subset of the data. (See Appendix A for specific examples of both methods.)

A very large body of work has been done regarding methods for visualizing tree data that focus on displaying specific data about a small subset of the data. Examples include MS−Windows Explorer[1], Macintosh Finder[2], ZTree[3], and TPEX Grinder[4]. Each of these tools attempts to pack in the greatest amount of information specific to a very small group of nodes: name, size, last−modified time, methods contained, etc. Because each of these fields takes up valuable screen real estate, there is a physical limitation to the number of fields that can be successfully included before there is no room for actual nodes. For most standard−size displays this translates into a listing of 15−40 nodes of the tree and 2−6 descriptor fields.

There is less data regarding research in the area of visually displaying the macroscopic characteristics of data. The earliest serious research yielded techniques such as the bifocal display[5], the fisheye lens[6], and the Perspective Wall. More recently Xerox PARC pioneered the cone tree[7] which was further developed by Tversky, et. al. at Brown University and by Carriere, et. al. at the University of Waterloo. Each of the above systems aims to represent a large amount of data in a generally but only vaguely informative manner. While most of these techniques allow the user to zoom in on a relatively small subset of the data tree, they primarily concerned with a more macroscopic view of the data tree. Later iterations of the cone tree have attempted to encode specific attributes of the nodes using various coloring and geometric schemes but they have not met with much success as these techniques are limited in the number of visual cues they can simultaneously provide without confusing the user.

Aims

The goal of the proposed project is to develop a visualization scheme that allows the efficient assimilation of both macroscopic and specific data pertinent to medium−sized data trees. Such a system would be immediately practical in a number of applications: file system browsing, web−directory browsing, source code browsing, etc.

We are proposing a technique that acts as a hybrid of the two. Fundamentally, it will be a macroscopic−information viewer. However, it will not only provide a visual system that efficiently conveys information about the macroscopic characteristics of the data tree but will allow a user to scale in and view relatively large amounts of detail of specific attributes of a small group of individual nodes. *We hypothesize that such a system will allow a user to efficiently assimilate information about both the macroscopic characteristics of the larger tree and specific information about a very small group of nodes, more so than using either the exclusively specific−information or macroscopic−information oriented tools.*

The proposed system is a 2D model that employs judicious use of perspective and other visual cues to add a sense of depth. One the most stark signs of immaturity in design is the failure of macroscopic information viewers to make good use of shading, hue changes, luminescence, etc. These are cues which humans manage daily and are well−suited to dealing with and will consequently be quite useful to the user when they are employed to encode information. Professor David Hornug will provide the necessary color expertise to devise an appropriately complex and useful strategy involving changes in color, hue, luminescence, tone, shading, and texture. Further, we will experiment with sound cues as they are often used to imply or connote motion.

This proposed technique differs in several significant ways from previous work on tree hierarchy visualization. First, no macroscopic−visualization scheme has ever attempted to incorporate a panel containing specific information. This is primarily because those systems were concerned with node scalibility and did not want to devote the valuable screen real estate necessary. Second, no previous system using macroscopic techniques has developed such an intricate encoding scheme that is capable of encoding a large amount of data using the aforementioned visual and perhaps audible cues. Third, prior efforts have only made use of optic flow in navigation of the data tree and did not attempt to incorporate other motion cues such as shading and texture.

Work Plan

The three phases to this projects are: information extraction, macroscopic display, and specific−information panel.

Information extraction refers to the preparation of data and conversion into an internal format that easily lends itself to being depicted as a tree. For example, it would not do to have the program crawl the file system or source code each time the perspective is changed; the program would first crawl the target system once, hold that information internally, and use it as needed. Initially, conversion functions would be made for tree data that seems promising. Hooks would also be provided so that later conversion functions may be plugged in to use to tool on other hierarchical data structures.

Macroscopic display refers to a set of functions that display the nodes to a user in the manner described above and which allow a user to change points of view within the graphical node structure.

Specific−information panel refers to a set of functions that would display detailed information about a small subset of the currently−viewable nodes.

We expect the first phase to take 1−1.5 weeks, the second phase to take 2−2.5 weeks, and the third to take 1−1.5 weeks. The remaining time (approximately 1 week) would be used to conduct user studies on the success of the model relative to a standard specific−information oriented tool and a macroscopic−orientation oriented tool. Final presentation and demonstration will be prepared during the final 2 weeks.

Future Applications

The proposed technique will be useful in many browsing contexts: file systems, the World Wide Web, source code trees, documentation trees, etc. Each of these applications would be directly enhanced by a better visualization scheme.

Tangentially but importantly, the general field of data visualization would benefit from the research done regarding the coordination of traditional optic flow cues with less−used shading and texture motion cues.

C.V. for Benjamin White

Personal

Name: Benjamin White
Email: bwhite@cs.brown.edu
Telephone: 401−867−6656

Education

ScB Neuroscience, Brown University (Expected May 2001)

Employment / Research Experience

Brown University Computer Science Dept. (1999−present)
Software Engineer, White Oak Technologies, Inc. (1999−present)
Software Engineer, Lighthouse Technology Solutions, Inc. (1997−1999)
Lab Technician, Holy Cross Hospital (1996)

Relevant Skills

Strong coding experience (awk, C, C++, Java, Perl)
Graphics knowledge
Good knowledge of visual perception

C.V. for Danah Beard

Name: Danah Beard
Home Address: Box 0556, Brown University, Providence RI 02912
Telephone: 401.861.0131

Education:

− 2000 A.B. in Computer Science, Brown University, Providence RI
− 2000 A.B. in Computational Gender Theory, Brown University, Providence RI

Research and/or Professional Experience:

− Software Engineer, Macromedia Generator, Macromedia, San Francisco, CA, June 1998−September 1999
− Research Assistant, Sociable Media, MIT Media Lab, Cambridge, Mass, June 1999−present
− Research Assistant, Welch Vision Lab, Brown University, Providence, RI, September 1999−present

Related papers:

− Beard, Danah (1999). Sex−based vision through the eyes of transsexuals undergoing hormone treatment. School for International Training, unpublished.

Capabilities and skills:

− computer science background and strong coding background (C/C++/Java)
− directed knowledge of vision− motion perception and depth perception

APPENDIX A

[1] **MS−Windows Explorer**

[2] **Macintosh Finder**

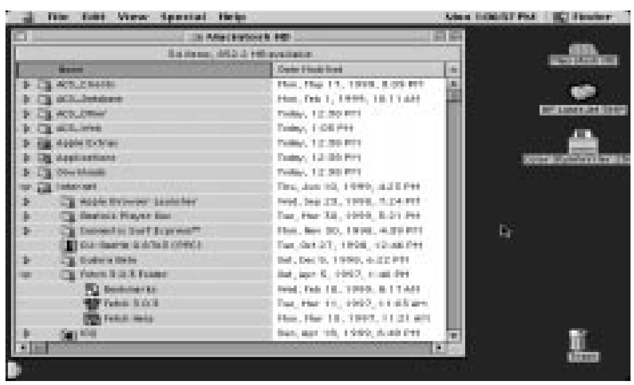

[3] **ZTree**

[4] **Grinder**

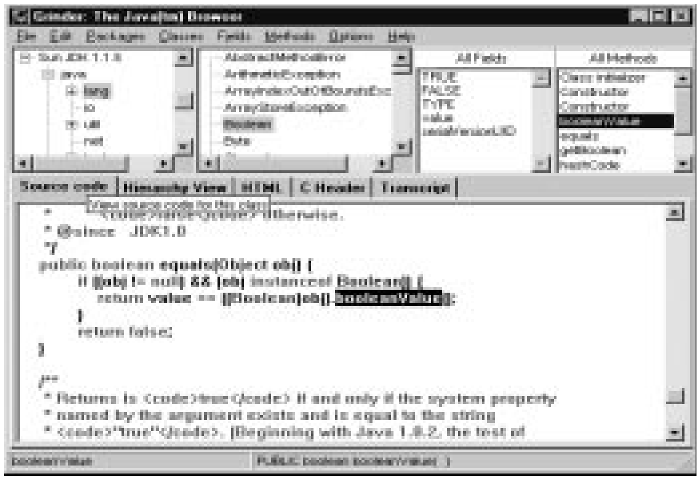

[5] **Bifocal display**

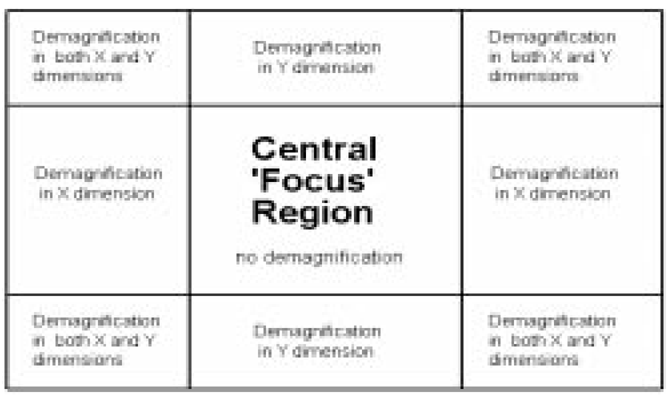

[6] **Fisheye lens**

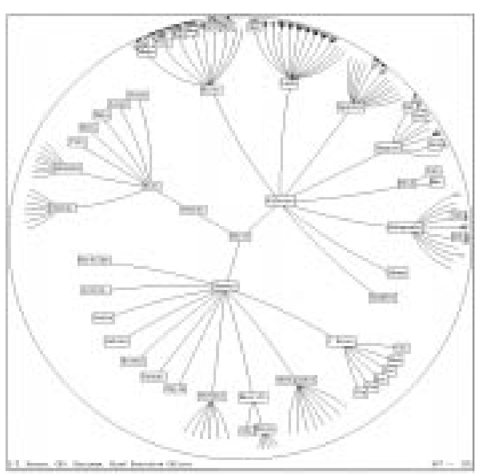

[7] **Cone tree**

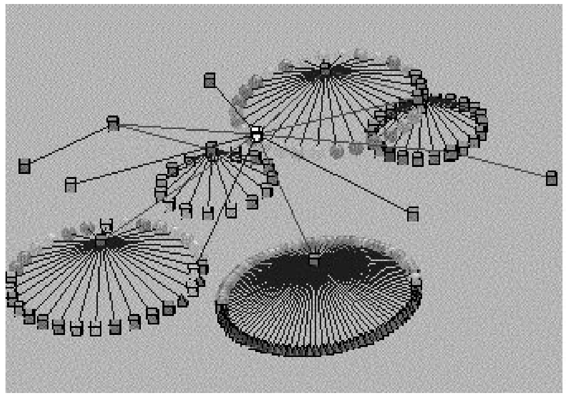

Avi Walsky
CS295-5
October 5, 1999

Enhanced Line Perception Using Neural Vision Models
Subset of:
Constructing Three Dimensional Objects from Two Dimensional Perspective Images

PI: Avi Walsky
Co-PI: Asohan Amarasingham
Consultant: Professor James Anderson (Cognitive and Linguistic Sciences)

Abstract

Enhanced line and edge perception in two-dimensional (2D) perspective images is the first step towards extracting three-dimensional (3D) objects from stereoscopic flat images. Enhanced edge detection alone has the effect of focussing and extruding 2D representations of 3D data. Its implementation provides an extra depth cue and for visualization of complex 3D images, which aid the human eye and mind in understanding and representing such complex data. Applying this technique to both architectural photographs and generated 3D renderings is a crucial step toward computer reproduction of real space objects from flat space images.

Aims

To enable computers to extract a 3D image from a single or stereoscopic set of images is an ongoing battle in the world of scientific computing. This project proposes to use human neural models to achieve such a result. Though the end result of creating such a system with 3D representational abilities does not fit within the 5-week time scale of this project, the primary creation and study of human eye edge enhancement models does.

The results of creating a system that can filter an image by enhancing edge and line separations is applicable to all scientific data representations as well as various computer art programs. This project will direct this filter toward its use in archeology and architecture as preliminary step in creating 3D models of real space objects/structures. Since the findings of such experiments could have great effects in enhancing pre-generated images from scientific data, this research will also apply the filter to various medical and scientific displays including MRI, ultrasound and CAT images.

Significance of Work:

On a large scale, the completion of the full 3D-imaging project would have drastic impacts in the fields of archaeology and architecture. In these areas it is necessary to catalogue and construct models of the structures and findings in the field. Cataloging archeological structures is extremely time and cost consuming as well as possibly detrimental to the objects being catalogued. Creating a program that can extract representations of such objects from multiple photographs would increase efficiency of archeological projects as well as decreasing their overall costs in manpower drastically.

Such a filter as proposed for this 5 week project has some applications in the field above, but most of its applications will be in enhancement of low-resolution images. Such images coming from ultrasound and even complex images obtained by MRI and CAT scans. Though archaeological and architectural images will be used for the learning period of the neural network design, such medical images will be applied this filter as a measure of the results. If possible, low resolution astronomical images and sonar scans of the ocean floor will also be used as a measure of the possible outcomes of such an enhancement network. The NASA and ocean scan images may prove too difficult to acquire, but their application is a strong intention of this preliminary computer vision model.

Avi Walsky
CS295-5
October 5, 1999

As a unintended but nonetheless welcomed outcome to such a neural vision model, a better understanding of how humans perceive tree dimensions may be formulated. Since it is the intent of this project to mimic human retinal and lower level neural filtering of images, a completed project will allow us to asses the “known” structure of the these biological filters by comparing human documented results to the computer models results.

Related Work

Personal experience of the Primary PI is limited to classes in human neuroscience and cognitive science with special emphasis on a class taught by the consultant, James Anderson, on neural modeling. Works mainly taken from the text of that class provide preliminary models for biological eye computations. The most important of such involves a model of the Limulus polyphemus eye and an accurate model of its lateral inhibition algorithms to represent photoreceptor and ganglion cell interactivity. The computer graphics skills of this PI will also be utilized to produce a result that not only works as it should, but displays something of meaningful, visual representation.

The CoPI’s background and paper on synaptic weights will be specifically utilized when creating the network with its moch-neurons and synapses. Both PI’s will utilize their programming skills in C++ to write the code for the network.

Work Plan

The first two weeks will be devoted to the creation of and debugging of the neural network that will be the actual filter. The next 2 weeks will be spent teaching the network, as the idea behind such an organization is that it learns how to see by itself using multiple images and simple data sets. The remaining time will be spent applying the filter to the proposed data sets above, some simple cases of which will be used in the training period. This last segment of the study is just an extension of the learning period with the implementation of more complex and novel stimuli.

The actual design of the network will be a Hebbian network with the specifics to be unfolded as different problems arise (as they always do). It will be programmed in C++ by both the PI’s.

Avi Walsky
CS295-5
October 5, 1999

Avi Walsky,Principal Investigator

Birthdate: December 27, 1977
Place of Birth: USA
Current Status: Researcher/Student, Brown University
Home Address: Box 5701, Brown University, Providence RI 02912
Telephone: 401.831.8986

Education:
- 2001 Sc.B.in Visual Science, Brown University, Providence RI
**this is the current name for my independant study (it may change)
Research and/or Professional Experience:
-Developed interactive web pages(1999) - A Good Thing Inc.

Relevant courses at Brown:
- Computer Science 123, Introduction to Computer Graphics
- Neuroscience 1, The Brain: An Introduction to Neuroscience
- Neuroscience 101, Neural Systems
- Neuroscience 102, Principals of Neurobiology
- Cognitive Science 102 - Neural Modelling Labratory
- RISD FormZ - 3D graphics

Capabilities and skills:
- strong neuroscience background
- computer science background moderate C++/Java
programming
- computer graphics knowledge (both as user and programmer)
- strong artistic background in 1,2 and 3 point perspective as well as other
artistic techniques in creating the illusion of depth

Avi Walsky
CS295-5
October 5, 1999

Asohan Amarasingham, Co-Principal Investigator

Birthdate: June 3, 1975
Place of Birth: New York, USA
Current Status: Graduate Student, Brown University
Division of Applied Mathematics
Department of Cognitive & Linguistic Sciences

Education:

B.A., 1997, Mathematics & Cognitive Science
University of Virginia

M.Sc., 1999, Cognitive Science
Brown University

M.Sc., 1999, Applied Mathematics
Brown University

Ph.D., expected 2003, Applied Mathematics
Brown University

Relevant Publications:

Amarasingham, A. & Levy, W.B. (1998) "Predicting the Synaptic Weight Distribution in A Self-Organizing Sequence Prediction Model" Neural Computation 10(1):25-57.

Capabilities and skills:
- training in mathematics, cognitive science, neuroscience
- computer programming (Matlab, C, C++, Fortran)

Avi Walsky
CS295-5
October 5, 1999

Visualizing Diffusion Tensor MR Image Using Streamtubes

October 5, 1999

Principle Investigator: Song Zhang
Co Principle Investigators: Daniel Acevedo, Paul Reitsma
Consultants: David Laidlaw, Eric Ahren

1 Abstract

MR measurements of water diffusion in organs and tissues having an orderly, oriented structure, such as skeletal, cardiac, and uterine muscle, portions of the kidney, the lens, and white matter, exhibit anisotropy(i.e., a dependence of the diffusivity on direction), which could be used to visualize the structure of the tissue. Previous work have been done on visualizing anisotropy and its direction in slices. We presents a method of visualizing diffusion tensor field in 3D using streamtubes.

2 Introduction

Within biological systems water molecules undergo continuous stochastic Brownian motion. In different tissues the rate of this diffusive motion can vary by several orders of magnitude - faster in liquids like cerebro-spinal fluid, slower in tissues like muscle. In some tissues the rate is anisotropic, or faster in some directions compared to others. Magnetic resonance imaging(MRI) can acquire images with intensity values sensitive to the diffusion rate of water. A quantitative image of the diffusion rate can be calculated from a set of such MR images. From a 2D slice or 3D volume image of the anisotropy we can infer the underlying tissue structure and better understand the anatomy of the nervous system, neuro-degenerative diseases, and neural development. Diffusion-rate images calculated from MRI measurements are second-order tensor fields. At each point in the MR image, there are totally 7 variables, 6 from the symmetric matrix representing the diffusion tensor, one from the intensity.Those many variables at one point pose a challenging problem in scientific visualization. Much of the previous work reduce the problem by visualizing a slice of the MR image, or visualizing only some of the information contained at every point. In one approach, Ellipsoids are used to present the tensor value at any spatial location. The ellipsoid is a natural icon, since its shape echoes the structure of the diffusion process. But this method is limited in visualizing 2D slices of the MRI data, in 3D, ellipsoids would block each other and won't provide a useful visualization.

3 Related Work

Since ellipsoid is a natural icon of the diffusion tensor, it's been used to generate pictures for a slice of MRI data. But when extended to 3D, ellipsoids probably will generate visual cluttering.
Also, an artistic approach is used to present a slice of the MRI data, different kinds of strokes is used to present different layers of information, as artists do in their paintings.

4 Methods

4.1 Streamtubes

To visualize the mri image volumetrically, we extend the method of ellipsoids to 3D, visualizing the diffusion rate with streamtubes. A streamtube

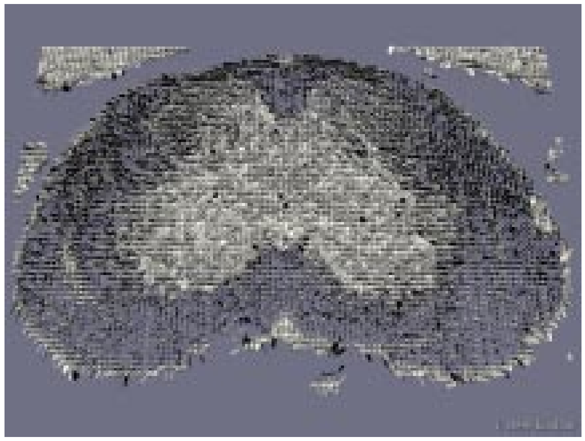

Figure 1: Visualizing slice of mouse spinal cord using ellipsoids

is a opaque tube that goes through 3D space. Our goal is to carefully map the variables in the diffusion tensor field to the visual features of the streamtubes. The visual features of the streamtubes include the direction of the streamtube, the cross section shape, the color, the texture ... The variables in one voxel includes six variables in the symmetric matrix of the tensor and the intensity. The tensor matrix could be represented by its three eigenvectors and three eigenvalues, ellipsoid is a natural icon for it in that it has a simple shape and exactly the same degrees of freedom as the tensor. Any point on the streamtube could represent the primary eigenvector(by the direction of the streamtube), the other two eigenvectors and their eigenvalues(by the cross-section ecllipse on that point). The primary eigenvalue could be mapped to the color on that point.

4.2 Programming tools

We choose jot as our programming environment because jot provide the support for manipulating mesh, interacting, and easy porting to stereo.

4.3 Porting to stereo

The resulting pictures should be streamtubes in 3D space, with support for user interaction, which could change the angle and position of the camera through which the user observes the streamtubes. Using stereo(workbench, cave) hopefully could greatly improve the human understanding of the visual features of the streamtubes, thus improving the understanding of the underlying mri data.

5 Validation

Real world mri data is hard to get and difficult to interprete, so making phantom data of some know geometric models could help us verify our method. We could also make 2D images of the data using the ellipsoid representation. By examing a sequence of slices through Z axis, we could get a idea of the 3D images. This method is tedious and indirect, but could be used to verify our method.

6 Timeline

1st week: Writing proposals
2nd, 3rd week: Experiment with mapping schemes from MRI data to visual features of the streamtube. Demo the pictures generated by the mapping schemes.
4th, 5th week: Porting the work to stereo. Demo the pictures shown in cave.
6th, 7th week: Validation, prepare for the demo.

CV of Song Zhang

Box 1910, Department of Computer Science
Brown University
Providence, RI 02912
sz@cs.brown.edu
http://www.cs.brown.edu/sz/
Education:

- Brown University, PhD Candidate.
- Nankai University, P.R.China, 1996, BS in Computer Science.

Research Experience:

- Research assistant, Brown Graphics Group, Sept. 1998-present. Research in Diffusion tensor field visualization.
- Research assitant, Nankai University, Sept. 1996-May. 1998. Developing a commercial OCR(Optical Character Recognition) product, Typ eReader.

Relavant Courses and course projects in Brown:

- cs123, a recursive raytracer.
- cs224, photorealistic rendering using photonmap, mini-sketch in Java3D, nonphotorealistic rendering, photorealistic rendering using metropolis method

Relavant skills:

- Experience in diffusion tensor visualization.
- Programming skill in C++, Java
- Knowledge in Graphics

Curriculum Vitae of Paul Reitsma

Box 1910, Department of Computer Science
Brown University
Providence, RI 02912
psar@cs.brown.edu
http://www.cs.brown.edu/∼ psar/

Education

- BSc, Combined Honours in Computer Science and Mathematics, University of British Columbia, with Honours, 1998

Relevant Experience

- Research Assistant, Electronic Games for Education in Math and Science (E -GEMS), University of British Columbia, January to April 1997
  - Web-based educational game development for grades 5-7
- Research Assistant, TRIUMF National Laboratory, January to April 1995
  - Assisting Neutrino Group with software simulations

Honours

- Andries van Dam Fellowship

Skills

- Strong computer science background and programming experience
- Strong mathematical background
- Good technical writing background
- Basic grounding in most physical sciences