

Lessons Learned from Employing Multiple Perspectives In a Collaborative Virtual Environment for Visualizing Scientific Data

Kyoung S. Park

Electronic Visualization Laboratory
1120SEO, MC154, 851S. Morgan St.
University of Illinois at Chicago
Chicago, IL 60607, USA
1 312 996 3002

kpark@eecs.uic.edu

Abhinav Kapoor

Electronic Visualization Laboratory
1120SEO, MC154, 851S. Morgan St.
University of Illinois at Chicago
Chicago, IL 60607, USA
1 312 996 3002

akapoor@evl.uic.edu

Jason Leigh

Electronic Visualization Laboratory
1120SEO, MC154, 851S. Morgan St.
University of Illinois at Chicago
Chicago, IL 60607, USA
1 312 996 3002

spiff@evl.uic.edu

ABSTRACT

This paper explores the concept of multiple perspectives to enhance collaboration by allowing remote participants to tailor their views, user-interfaces and roles to their particular needs and expertise. It describes a preliminary design study conducted on users of a collaborative CAVE-based virtual reality tool for visualizing oceanographic data. Results will focus on the patterns of activity within this environment, in particular the manner in which participants transition between individual and group work during the course of a collaborative session.

Keywords

CSCW, awareness, multiple perspectives, subjective views.

1. INTRODUCTION

In the real world, individuals who are trying to solve a common problem gather (in workshops, for example) in the hopes that their combined experiences and expertise will contribute new perspectives and solutions to the problem. In most collaborative and immersive virtual environments (CVEs) to-date, the default assumption has been to display the collaborative world in the same way to all its participants. We believe that by employing multiple perspectives and in particular by encouraging role-

specialization, collaborators will be able to solve problems more effectively in CVEs. The following two examples will help motivate this concept:

- In a collaborative industrial design application, an engineer may be collaborating with an artist. They may be collaborating on the design of a new office chair. The artist can manipulate the design using tools that are specifically tailored for his/her expertise, such as interactive sculpting tools. Simultaneously the engineer can evaluate the impact of such design changes in terms of material stress and strain.
- In the context of scientific visualization we envision a potential application of multiple perspectives in the visualization of multi-dimensional data-sets. Here a large number of dimensions may be partitioned across multiple viewers to assist in reducing the overall complexity of the content being visualized. The challenge is in providing the necessary interface to support this collaboration while minimizing the confusion and additional cognitive load that may result from having to coordinate the activity of multiple viewers all simultaneously viewing disparate parts of the data-set.

This paper describes our application of multiple perspectives to collaborative virtual environments with a special focus on scientific visualization. CAVE6D (Figure 1) is a collaborative CAVE-based visualization tool for exploring multivariate oceanographic data sets. In this case the data being visualized is a simulation of tidal patterns in the Chesapeake Bay. A unique feature of CAVE6D, as compared to other CVE visualization tools, is its ability to allow participants to customize their views while working collaboratively. Hence even though the participants are all viewing the same data-set they are seeing the data from

all the users see and share the same information and interface. Often this sharing results in users operating the interface sequentially. WYSIWIS can however be relaxed so that individual users possess the ability to configure their shared users interface to best suit their working needs, either group needs or personal needs. These kinds of tools are termed *collaboration-aware*[9]. Gutwin et al [3] describe techniques such as *Action Indicators* and *Process Feedthrough*, to highlight actions made by collaborators, as a means to enhance group awareness during the use of tools that are collaboration-aware.

In VR, the use of multiple perspectives is fairly recent and sparse. Stoakley introduced a non-collaborative system called Worlds in Miniature (WIM) in which a user interacting with a virtual space would be presented with two simultaneous views of the scene – a first person view (egocentric view) and another miniature third person view (exocentric view) [18]. The exocentric view afforded gross manipulations and navigation of the scene whereas the egocentric view afforded fine manipulation of objects in the scene.

A similar implementation is found in a collaborative architectural layout system called CALVIN [10]. CALVIN introduced the notion of *Mortals* and *Deities*. Mortals were participants in the environment who would view and manipulate the world from an egocentric perspective and interface. Deities would modify and view the world from an exocentric perspective. CALVIN provided

decidedly different physical perspectives as well as through different filters.

This paper will first present the concept of Multiple Perspectives in the context of previous related work. It will describe a design study to explore how users make use of multiple perspectives in a collaborative scientific visualization system. Finally it will highlight some of the observations, lessons learned, and ideas for future exploration.

2. RELATED WORK

We define Multiple Perspectives as a broad term to encompass multiple and possibly heterogeneous viewpoints, representations and roles, that can be adopted within both a collaborative and non-collaborative context. The following describes related research in the area.

A number of traditional Computer Supported Cooperative Work (CSCW) systems have implemented multiple perspectives in the past. These include systems such as Colab, Rendezvous, GroupKit, and Suite where shared objects can be viewed from different perspectives, such as through *public* views that are accessible to the entire group, or *private* views for personal use [2][11][12][13][17].

The use of multiple perspectives in early synchronous groupware products was motivated by the over-restrictiveness of *collaboration-transparent* WYSIWIS (What You See Is What I See) interfaces. Strict-WYSIWIS is a mode of operation whereby

Ms9825hosgeudipanssife0ddTuapperfacelikss i -111p articipa. Furthirrtie,D/F

suggests that the user's views should be allowed to converge and diverge depending on the application being used, the activities being performed, and the degree of cooperation desired.

Smith's approach to supporting subjective views is an access matrix that describes a shared object's range of possible representations [15]. The orthogonal axes of this matrix define two independent factors, appearances and modifiers. Appearances describe the major geometric changes of representation of an object, a room, or a space. Modifiers describe the set of operations that may be applied to an appearance. In later work Smith introduced the notion of the *re-coupled view* [16] as a way to overcome the difficulty of maintaining a common frame of reference and awareness when using purely subjective views.

A number of researchers have cited the potential benefits of multiple perspectives to enhance learning [4][8]. Larkin and Simon [8] suggest that differing representations induce distinct cognitive processes and hence provide a better fit to individual learning styles. Salzman, et al. [14] recently provided support for this idea in a VR learning environment to teach electric field concepts. In their study, high school students who interacted with both an exocentric view and an endocentric view of an electrical field exhibited a significantly higher level of mastery than students who only interacted with one of these visualizations.

Hennessy [4] has shown that the use of multiple perspectives can help constrain alternate and possibly incorrect interpretations of a piece of knowledge. For example the orthographic and isometric views in a CAD system contain degrees of freedom that are only disambiguated when the views are interpreted together. On the other-hand Ainsworth [1] indicates that multiple perspectives hinder rather than enhance learning because of the cognitive load imposed by having to simultaneously support the following tasks: 1. the learner must learn to understand each representation; 2. they must understand the relationship between the representation and the domain it is representing; 3. they must understand the relationship between each representation.

These contradictory findings would suggest that multiple perspectives help rather than hinder when the benefit of the multiple perspectives is offset by the increase in the cognitive load incurred in interpreting these perspectives. This cognitive load may be lessened if the proper tools are provided to coordinate the correspondences between the perspectives.

3. CAVE6D

CAVE6D [23] is a CAVE-based collaborative application for visualizing time-varying multivariate oceanographic data-sets. It allows multiple users to visualize, discuss, and interact with a data-set as well as with each other. Avatars that depict the users are equipped with long pointing-rays that can be used to point at features of interest in the data set. In this study we used a 15-day tidal simulation of the Chesapeake Bay. Table 1 lists the




Figure 2. CAVE6D's menu interface for the Chesapeake Bay.

Table 1. Descriptions of graphical object tools in the Chesapeake Bay data-set in CAVE6D

Topography	A solid or wire frame representation of the landmasses.
Contour Salinity	This visualizes salinity levels as contour lines. The magnitude of the salinity is depicted as a number next to the contour line.
Surface Horz Vec	This visualizes the velocity of the tide along the surface of the bay. This is interactively adjustable, allowing one to inspect tidal velocities at varying depths in the bay.
Bottom Horz Vec	Used in conjunction with Surface Horz Vec one can view the tidal velocities at both the deep and shallow regions of the bay.
Vert North South	This visualizes the tidal velocities cross sectional from North to South. This tool can be moved along the east/west axis.
Vert East West	This visualizes the tidal velocities cross-sectional from East to West. This tool can be moved along the north/south axis.
Tracer Iso	This visualizes salinity levels as iso-surfaces (3D contour diagrams). Red represents higher salinity levels and blue represents lower salinity levels.




Figure 3. The instructor is observing on an ImmersaDesk while the two participants are collaborating in their respective CAVEs.

visualization tools (or parameters) available in CAVE6D. Collaborating participants can switch any of these viewing parameters on or off, *globally (propagating their changes to everyone)* or *locally (affecting only one's own local view)*.

Figure 2 shows the CAVE6D menu interface for enabling and disabling visualization parameters. The items on the left panel depict the available visualization tools. The corresponding items on the right panel allow users to toggle between global and local settings. In this illustration, the surface horizontal vector field is the only parameter that is being globally shared amongst all other users in CAVE6D session. The black text highlighting the vertical east west vectors indicates that the tool is interactively manipulable and affecting only one's view.

CAVE6D enforces no floor control in the manipulation of these parameters. Time on the other-hand is globally shared and hence participants view all time-varying data synchronously. The original designers of CAVE6D enforced this as a way to provide a common frame of reference for its collaborating users, to reduce confusion when coordinating between potentially divergent views.

4. THE DESIGN STUDY

The study consisted of three pairs of computer science students graduate students. To minimize individual differences we chose to limit the study to computer science students who have already considerable experience with virtual reality. Furthermore, as this study would require students to work in pairs we tried to pair students who already have an established working relationship with one another. Most were familiar with the concept of CVE, but none of them had prior experience with CAVE6D and the Chesapeake Bay data-set. In the future we intend to apply a refined version of this experiment to study real oceanographers or students of oceanography.

As shown in Figure 3, each pair of participants collaborated with each other in their respective CAVEs, and a third participant (the instructor) would observe the proceedings on an ImmersaDesk. All of them could speak to each other via a high quality audio connection using headphones and microphones. An assistant was assigned to each CAVE to correct any technical problems that arose. One evaluator recorded the users' answers to a number of queries. A video camera recorded one workstation's screen and all conversations.

The users were organized into three treatment groups: those who were allowed to use only *localized (private)* views; those who were allowed to use only *globally (fully) shared* views; and those who were allowed to use *either* view. The experiments ran for three days with each group experiencing one session per day for approximately one hour. The first day consisted of a training session to introduce the groups to CAVE6D. The remaining days consisted of guided search sessions where they were asked to search for specific trends/patterns in the data set. On the second day, groups were assigned to one of the three treatment conditions. On the third day, they were assigned to the remainder of the two treatment conditions. Each session was preceded by a pre-test and succeeded by a debriefing questionnaire.

4.1 The Training Session

Each group experienced a 1-hour collaborative training session in which each person in the pair was placed in their own respective CAVE. The training session was given remotely by the instructor on an ImmersaDesk. Each participant was able to see each other as avatars with long pointing rays emanating from their hands so that they could point at features in the visualization. In addition high quality zero-latency audio mediated their conversations.

The training regime would consist firstly of a description of the Chesapeake Bay geography identifying the location of individual landmarks such as Baltimore, Virginia, the Atlantic Ocean, etc. Then the instructor would run through each of the CAVE6D's visualization tools, describing their function in the Chesapeake data set. For example there were tools to display ocean tide vectors and tools to display iso-surface visualizations of salinity.

During the training session the images from each of the CAVEs were composed using a video switcher (so that each CAVE's images appeared in one half of the video screen) and recorded onto videotape accompanied by audio from all three participants.

Finally the instructor would familiarize the groups in the use of global and local views. Participants then practiced manipulating the visualization parameters until they felt comfortable with them. The training session was followed by a debriefing, in which the students gave us feedback on their experience.

4.2 The Guided Search Sessions

On the second day each group was asked to perform a set of guided search tasks over a period of one hour. The goal of the guided search tasks was to allow participants to focus on the single task of using the tools to verify the existence (or absence) of a specific trend in the data. The plan was to offer the more difficult, unguided free-form search opportunity on the third day, after they had gained some expertise in wielding the tools.

In the guided search task, the instructor would read one trend at a time and ask the participants to use the tools to determine whether the trend existed or not. A sample question would be "How often does the cycle of the tide repeat itself?" Some of the questions were intentionally ambiguous to encourage discourse and hopefully coordination.

When the pair had agreed on an answer for the trend, the evaluator would be called in to record it. The record would note the visualization tools that were activated by each participant and the tools that were considered the most important in revealing the trend. After recording the answer, the instructor would move on to the next question. If after a preset amount of time (10 minutes) they were not able to agree on an answer to a trend they were asked to move on to the next trend. Each group was given five trend questions. At the beginning of each task the instructor would

reset each participant's visualization tools to their original un-activated settings. The guided search task was followed by a debriefing, as on Day 1.

On Day 3, we permute the condition amongst groups so that, for example, formerly group with local views would be asked to search for new trends using global views and then either views. Because of time constraints we only allowed each guided search session to last 30 minutes rather than the full hour, thus we asked groups three questions for each session. In each case they were allowed no more than ten minutes for each trend even if they were not able to agree on an answer.

5. OBSERVATIONS

5.1 The Training Session

As described earlier, a remote instructor on an ImmersaDesk trained two students, in their respective CAVEs. They were instructed on basic Chesapeake Bay geography and how to operate CAVE6D's visualization tools. All participants in the environment were represented as avatars composed of a head, body, and single hand. The name of each user was displayed on the "jerseys" of each avatar. The head and hand's position and orientation were tracked with electromagnetic sensors. The body's orientation was locked to the front wall of the CAVE and hence altered when the user navigated through the space using a joystick. Drawn at the end of each avatar's hand, is a long pointing ray to allow the avatars to point at features in the visualization. Figure 4 illustrates two participants' different perspectives.

5.1.1 Occlusion caused by Avatars

Avatars are useful in a CVE because it allows a participant to indicate one's location relative to a large space, and convey natural gestures such as the nodding of one's head or waving of one's hand. The orientation of an avatar's head is very useful for determining where the avatar is looking. The pointer at the end of the avatar's hand is useful for pointing out interesting features to one's collaborators. In some cases however when two avatars are attempting to share the same view the presence of the avatars can occlude each other's view. Furthermore it is disconcerting to users when avatars interpenetrate one another. Collision detection can be used to help mitigate this problem. However Tromp et al [22] have found that collision detection in small enclosed spaces can impair navigation, especially for novices to VR.

5.1.2 Multiple Viewpoints within One Visualization

In previous work on desktop collaborative systems (such as text editors) [13] users are able to simultaneously navigate and edit different regions of a text document. The equivalent of this in CVEs, is that avatars are able to navigate independently through the virtual scene. Hence viewers are able to obtain multiple viewpoints from a single representation. However, unlike

traditional CSCW systems the avatars are able to rotate around the data and view it from any arbitrary orientation as in a real world workspace [21]. Hindmarsh et al. [5] describe such views as *Fragmented views*, i.e. features of the world are fragmented to different perspectives due to the narrow field of view on desktop CVEs, and have noted problems of losing group awareness.

5.1.3 The Need for a Shared Interface as well as a Shared Visualization

During the course of instruction it became clear that it was necessary, from time to time, to be able to share the view of the menu as well as the visualization. This would allow the instructor to confirm whether each student was using the correct visualization settings. Furthermore, in the context of the general use of CAVE6D, it would allow participants to be more aware of each other's actions. Since CAVE6D did not support this, when users were using localized views, the instructor had to spend a considerable amount of time confirming that each student had set his/her visualization parameters correctly.

Drawing from previous work by Gutwin [3], a solution might be to provide an over-the-shoulder view to allow an instructor to view the environment through each participant's eyes, and to provide a "radar" view to quickly identify which visualization parameters were being viewed by each participant. In addition *Action indicators* and *process feedthrough* could be used to enhance group awareness by signaling or illustrating when participants are performing actions.

5.1.4 Misinterpretation of Gestures due to a Mismatch between Coordinate Systems

Since CAVE6D is a fully immersive application each student is immersed in the data. When, for example, they cruise through the trenches in the Chesapeake Bay, their companions will see their avatar perform likewise. That is, the participants and their respective avatars navigate through the data in world coordinates. However the menu interface was presented to each CAVE in local coordinates. Hence, when activated, the menu would float in a constant location in the physical CAVE. This is the typical method of operation for most menu systems in VR. The problem however is that even though the participants are operating the menu in local coordinates, their avatars are represented to others in world coordinates. Hence the avatars appear to be giving arm gestures that may mislead the remote viewers into thinking that they were gesturing at something in the world coordinate frame. The solution might be to stop world coordinate updates of the avatar when menus are being used, and use action indication via audible prompts to provide feedback. Alternatively one could situate menus in the world coordinate frame. This may be impractical for large navigable spaces, as it would require users to navigate to the menu in order to use it.

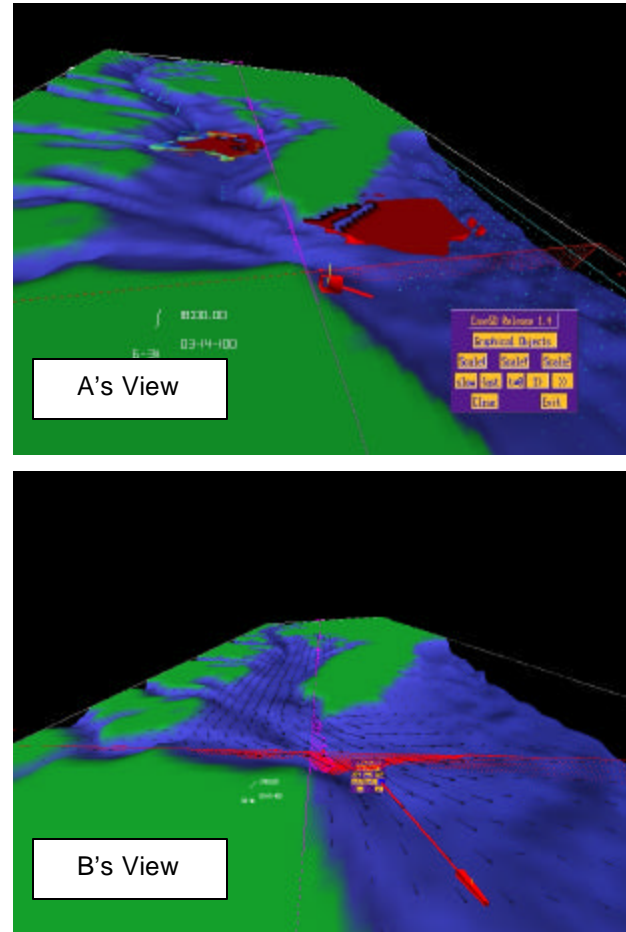


Figure 4. Tailored views for two different participants, depicting that Topography, Vert East West, Vert North South tools are globally shared between them. The Tracer Iso-surface and Bottom Horz Vec tools are visualized only in A's view only. Likewise, Surface Horz Vec tool in B's view only.

5.1.5 The Need for Sharing Intent as well as the State of the Interface

It was observed that simply showing the state of the interface was not sufficient since it did not offer the instructor any feedback on which particular interface item the student was attempting to operate. For example CAVE6D offers a tool for displaying a cross-sectional view of the ocean vectors at any given slice in either the North/South or East/West direction. When a slice is moved, there is no indication of which user was doing the manipulation, or which slice was being manipulated. And since the arm gesture used to move the slice was represented in local coordinates, the avatar appeared as though it were raising its arm for no apparent reason. This problem was worsened by the fact

that each avatar's pointing ray could not be toggled on and off hence further amplifying the unintended hand gestures.

5.2 The Guided Search Sessions

Overall we found that regardless of treatment condition, collaborators largely worked independently except to converge on their discoveries. The recurring pattern of activity was 1. problem interpretation, 2. agreement on parameters to use, 3. independent search for a trend including some adjustments to viewing parameters, 4. reporting of discoveries to their partners, and 5. negotiating a conclusion based on their combined discoveries. Negotiation may require revisiting the site where the trend was discovered to collectively confirm the discovery. This was especially true when conflicting discoveries were found.

5.2.1 Patterns of Discourse

The participants tended to say few things when they were searching intensely for their trends. The occasions for discourse were 1. to agree on the tools to use, 2. to disambiguate situations when CAVE6D was acting faultily, 3. to announce and debate over a discovery, or 4. to resolve conflicts. When localized views were used, groups spent more time talking to each other to synchronize their views. False trends tended to introduce an added degree of uncertainty and hence appeared to produce more conflict and discussion.

5.2.2 Usage Patterns for Localized and Global Views

As mentioned earlier, search patterns across groups tended to be largely the same – the participants would work mostly independently except to converge on their discoveries. They also tended to work independently when they needed to search over a large body of terrain. They would partition the terrain so that each would search for a trend over a different part of it. Hence the navigation provided multiple perspectives over a single perspective.

When available, groups preferred to use a localized view to test out small hypotheses individually without disturbing the overall view. Then they would use a global view to present their findings to their partner. The use of localized views was particularly prevalent when users wanted to manipulate the visualizations—such as the depth of the tidal vectors in the bay.

These transitions between individual and coordinated activity are consistent with observations in the literature on workflow [17].

Groups that had to use only global views (similar to strict WYSIWIS systems) ran into frequent problems of users toggling the same viewing parameters at the same time and hence accidentally undoing the viewing option. In addition participants also complained about cases where their partner would inadvertently change the visualization without forewarning them.

Here again is where Gutwin's [3] suggestions for providing action and process feedthrough indicators would be helpful.

In strictly localized views participants engaged in lengthier discussions to ensure that their visualizations were consistent. So while a localized view typically favored independent work, the additional discourse needed to synchronize views required participants to increase coordination.

When groups were allowed to use either views at will, they tended to work independently preferring to use localized views. However, during the debriefing session, participants commented that overall they preferred a global view, even though their interaction history seemed to indicate more frequent use of localized views.

There were situations where participants would disagree on their conclusions. In such cases they attempted to prove the veracity of their discovery by asking their partner to temporarily switch to their viewpoint. This was clumsy in local-only interactions and was easier in global and either view interactions. However as there was no way to save previous viewing states they were hesitant to make those viewing changes. This suggests that in order to support multiple perspectives well there needs to be a means to capture these states for side-by-side comparison. Furthermore, the recording of these states will allow a faithful recreation of the visualization that is needed for future documentation of the discovery or for long-term asynchronous collaboration.

5.2.3 Usage Patterns for the Visualization Tools

The strategy adopted by all groups was to turn on all tools that seemed relevant to the task and then eliminate the unnecessary ones gradually. The most favored tools depended on the viewer's point of view and their comfort with using the tool. In some cases the most favored tool was the least efficient tool for answering the query but also resulted in the users providing the most precise answer. For example, in answering the query "What are the differences in the salinity levels in the sea and in the bay in general," one group used the Contour Salinity tool while another group used the Tracer Iso-surface tool (see Table 1.) As the Contour Salinity tool depicted precise salinity levels with numbers next to the contour lines, answers to the query took longer to determine but was more thoroughly determined. On the other-hand the Tracer Iso-surface revealed the trend very quickly at the expense of precision.

5.2.4 Individual and Group Differences

Individual differences between participants played a noticeable role in the nature of the collaboration and, in retrospect can be considered a weakness of this experiment— in particular since the subject pool was so small. The more dominant participant tended to take control of the collaboration. In addition, it appeared that more pro-active participants made greater use of localized views to explore alternative solutions. In future pairings of participants we will attempt to better match their personalities.

It appeared that more cooperative groups tended to work more effectively by dividing the task amongst themselves, integrating their thoughts and findings, and carefully re-examining disagreements. In one case it was noted that partners who did not adopt a cooperative problem solving strategy took significantly longer to converge on an answer for the trend. The question that resulted in this difficulty was “over what period does the tide cycle?” Groups that distributed this problem so that one participant viewed the system clock while the other viewed the patterns of the tide were able to solve the problem quickly. The group that did not adopt this strategy was not able to answer the query in the allotted time because both of the members were trying to watch the cycle of the tide and the clock at the same time.

5.2.5 Learning of the domain material occurred

As a side effect of this study it was found that participants were able to clearly and completely articulate each of the trends they found during the debriefing session. This implies that the guided search process may be useful as an instructional technique for new users of the system.

6. CONCLUSIONS

Overall we found that collaborators mostly worked independently, even in globally shared views. The recurring pattern of activity was 1. problem interpretation; 2. agreement on visualization tools to use; 3. independent search for a trend including some adjustments to viewing parameters; 4. reporting of discoveries to their partner, 5. negotiating a conclusion based on their combined discoveries. Negotiation may require revisiting the site where the trend was discovered to collectively confirm the discovery. This process often requires each collaborator to share their visualization parameters with the other participants. Hence some ability to record the settings of those tools and restore them could greatly assist the process. Participants will be able to take snapshots of events of interest, saving all the visualization parameters, and recall them to share with their collaborators.

As mentioned earlier, search patterns across groups tended to be largely the same – the participants would work mostly independently except to converge on their discoveries. Users preferred to use localized views to test out small individual hypotheses without disturbing the overall view while they were searching, then used global views to present their findings to their partner.

Over-the-shoulder views, radar views, action indicators and process feedthrough mechanisms may be useful for improving awareness in these environments especially when collaborators are watching significantly diverging views.

This study has given us insight into how we might improve the CAVE6D interface as a whole, as well as an experiential understanding of how collaborators behave in CVE applications

that provide the ability to manipulate multiple perspectives. In future work we are interested in pursuing some of the following research issues:

- Investigate techniques for creating Action Indicators and Process Feedthrough in CVEs.
- Compare guided versus unguided search for trends. That is, in unguided search, we would allow collaborators to attempt to find discoveries on their own without any suggestion of whether a trend existed or not. Our hypothesis is that customized views will be used more frequently in unguided searches than in guided searches. In some ways unguided searches would more closely emulate the manner in which scientists engage a data set for the first time. On the other-hand scientists are usually equipped with a deep knowledge of the domain area to be able to know apriori what might be of interest in the data set. So essentially both guided and unguided search tasks could be considered valid emulations of scientific discovery.
- This study did not effectively explore how multiple perspectives can enhance interpretation of multivariate data with *very high dimensionality*. In such a case we anticipate the value of customized views to be more prevalent as the tailored perspectives are needed to reduce the clutter in the visualization. Furthermore we anticipate a need to also provide techniques to transition from one representation to another.
- It would be interesting to observe how the *number* of participants will affect collaborative visualization with multiple perspectives. This could be done in the original context of multiple CAVEs (participants are distributed) or with multiple participants in a single CAVE (participants are co-located) where each wall of the CAVE projects a separate visualization for each participant.

7. ACKNOWLEDGMENTS

We would like to thank all the students for serving as volunteers and Chris Scharver for serving as an instructor for this pilot study. In particular we would like to thank Cathy Lascara and Glen Wheelless at the Center for Coastal and Physical Oceanography for giving us comprehensive information on CAVE6D and the Chesapeake Bay data set. We would also like to thank Natt Mintrask and James Costigan for providing audio/visual support during data capture.

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible by major funding from the National Science Foundation (NSF), awards EIA-9802090, EIA-9871058, ANI-9980480, ANI-9730202, and ACI-9418068, as well as NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the

National Computational Science Alliance. EVL also receives major funding from the US Department of Energy (DOE), awards 99ER25388 and 99ER25405, as well as support from the DOE's Accelerated Strategic Computing Initiative (ASCI) Data and Visualization Corridor program. In addition, EVL receives funding from Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan.

The CAVE and ImmersaDesk are registered trademarks of the Board of Trustees of the University of Illinois. ImmersaDesk2, PARIS, and Wanda are trademarks of the Board of Trustees of the University of Illinois.

8. REFERENCES

- [1] Ainsworth, S. E., Wood, D. J. and Bibby, P. A., "Evaluating principles for multi-representational learning environments," *Proceedings of 7th EARLI conference*, Athens, 1997.
- [2] Dewan, P. "A tour of the suite user interface software," *Proceedings of 3rd ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 57–65, Oct 1990.
- [3] Gutwin, C. and Greenberg, S. "Design for individuals, design for groups: tradeoffs between power and workspace awareness," *Proceedings of the ACM CSCW'98 Conference on Computer Supported Cooperative Work*, pages 207–216. ACM, ACM Press, 1998.
- [4] Hennessy, S., Twigger, D., Driver, R., O'Shea, T., O'Malley, C., Byard, M., Draper, S., Hartley, R., Mohamed, R. and Scanlon, E. "Design of a computer-augmented curriculum for mechanics," *International Journal of Science Education*, 17(1):75–92, 1995.
- [5] Hindmarsh, J., Fraser, M., Heath, C., Benford, S. and Greenhalgh, C. "Fragmented Interaction: Establishing mutual orientation in virtual environments," *Proceedings of the ACM CSCW'98 Conference on Computer Supported Cooperative Work*, pages 217–226. ACM, ACM Press, 1998.
- [6] Johnson, A., Moher, T., Ohlsson, S. and Gillingham, M. "The round earth project: Deep learning in a collaborative virtual world," *Proceedings of IEEE Virtual Reality*, pages 164–171. IEEE, March 1999.
- [7] Kaput, J. J. "Linking representations in the symbol systems of algebra," In S. Wagner and C. Kieran, editors, *Research Issues in the Learning and Teaching of Algebra*. Hillsdale, NJ, 1989.
- [8] Larkin, J. H. and Simon, H. A. "Why a diagram is (sometimes) worth ten thousand words," *Cognitive Science*, 11:65–99, 1987.
- [9] Lauwers, J. C. and Lantz, K. A. "Collaboration awareness in support of collaboration transparency: Requirements for the next generation of shared window systems," *Proceedings of the ACM CHI'90 Conference on Human Factors in Computing Systems*. ACM, 1990.
- [10] Leigh, J., Johnson, A. E., Vasilakis, C. A. and De-Fanti, T. A. "Multi-perspective collaborative design in persistent networked virtual environments," *Proceedings of IEEE Virtual Reality Annual International Symposium '96*, pages 253–260, Apr. 1996.
- [11] Patterson, J. F., Hill, R. D., Rohall, S. L. and Meeks, W. S. "Rendezvous: An architecture for synchronous multi-user applications," *Proceedings of the ACM CSCW'90 Conference on Computer Supported Cooperative Work*, pages 317–328. ACM, 1990.
- [12] Roseman, M. and Greenberg, S. "Groupkit, a groupware toolkit for building real-time conferencing applications," *Proceedings of the ACM CSCW'92 Conference on Computer Supported Cooperative Work*, pages 43–50. ACM, ACM Press, 1992.
- [13] Roseman, M. and Greenberg, S. "Building real-time groupware with groupkit, a groupware toolkit," *ACM Transactions on Computer-Human Interaction*, 3(1):66–106, Mar 1996.
- [14] Salzman, M., Dede, C., Loftin, B., and Ash, K. "VR's frames of reference: A visualization technique for mastering abstract information spaces," *Proceedings of 3rd International Conference on Learning Sciences*, pages 249–255, Charlottesville, VA, 1998. Association for the Advancement of Computers in Education.
- [15] Smith, G. and Mariani, J. "Using subjective views to enhance 3D applications," *ACM Symposium on Virtual Reality Software and Technology*, pages 139–146, Lausanne, Switzerland, 1997. ACM, ACM Press.
- [16] Smith, G. and O'Brien, J. "Re-coupling tailored user interfaces," *Proceedings of the ACM CSCW'98 Conference on Computer Supported Cooperative Work*, pages 237–246. ACM, ACM Press, 1998.
- [17] Stefik, M., Foster, G., Bobrow, D. G., Kahn, K., Lanning, S. and Suchman, L. "Beyond the chalkboard: Computer support for collaboration and problem solving in meetings," *Communications of the ACM*, 30(1):32–47, Jan 1987.
- [18] Stoakley, R., Conway, M., and Pausch, R. "Virtual reality on a WIM: Interactive worlds in miniature," *Proceedings of SIGCHI*, pages 265–272. ACM, ACM Press, May 1995.
- [19] Snowdon, D., Greenhalgh, C., and Benford, S. "What You See is Not What I See: Subjectivity in Virtual Environments," *Framework for Immersive Virtual Environments (FIVE'95)*, December 1995.

- [20] Snowdon, D., J-Aro. K.M. "A Subjective Virtual Environment for collaborative information visualization," *Virtual Reality Universe'97*, April 2-5, 1997.
- [21] Tang, J. C. "Findings from observational studies of collaborative work," *International J. Man-Machine Studies*, 34(2):143-160, 1991.
- [22] Tromp, J., Bullock, A., Steed, A., Sadagic, A., Slater, M. and Frecon, E. "Small group behavior experiments in the coven project," *IEEE Computer Graphics & Applications*, 18(6), Nov/Dec 1998.
- [23] Wheless, G.H., Lascara, C.M., Cox, D., Patterson, R., Levy, S., Johnson, A., Leigh, J., Kapoor, A. "The Use of Collaborative Virtual Environments in the Mine Countermeasures Mission," *Proceedings of SPIE's 13th annual international symposium on aerospace / defense sensing, simulation, and controls*, April 5-9, 1999.