Interdisciplinary Immersive Analytics at the Electronic Visualization Laboratory: Lessons Learned and Upcoming Challenges

G. Elisabeta Marai* Angus G. Forbes† Andrew Johnson‡

Electronic Visualization Laboratory
Department of Computer Science
University of Illinois at Chicago

Abstract

We describe the lessons learned from three recent Immersive Analytics projects which take place at the Electronic Visualization Lab. These successful collaborative projects use the CAVE2 immersive hybrid environment. All three projects visualize multifaceted scientific data and benefit from interdisciplinary collaborations with experts across application domains. We further outline the challenges and opportunities ahead for the field of Immersive Analytics.

Index Terms: K.6.1 [Immersive Analytics]: Virtual Reality—Interdisciplinary collaborations; K.7.m [Technology]: Displays—CAVE2

1 Introduction

The Electronic Visualization Lab (EVL) at the University of Illinois at Chicago has been at the forefront of Virtual Reality research since 1992, when it introduced the CAVE Automatic Environment, a projection-based virtual reality system. Since then, the lab has developed a wide range of immersive systems to be used for visualization and visual analysis tasks, such as PARIS [13], ImmersaDesk [7], and CAVE2 [9]. These technologies can be and are used for Immersive Analytics (IA) tasks — deriving insights from data by augmenting the humans' ability to analyze and make sense of the large and multifaceted datasets which are common across scientific disciplines.

In this paper, we describe the challenges and the lessons learned from recent IA projects which take place at EVL. These collaborative projects primarily make use of the CAVE2 immersive environment (though some also are portable to 3D desktop environments or can also run on head-mounted displays). We further summarize some of the opportunities ahead for the field of Immersive Analytics.

2 Background

CAVE2, unveiled in 2012, is a 74 megapixel 2D / 37 megapixel passive 3D hybrid reality environment designed and built based on our lessons from building the original projector-based CAVE in 1991, and the large high-resolution tiled LCD displays we designed and built in the 2000s. 36 computers drive 72 LCD panels (18 columns of 4) arranged in a 22' wide 320 degree circle. 14 Vicon tracking camera allow us to track 6 objects in the space (glasses and/or controllers) and a 20.2 surround system provides audio feedback. The goal of CAVE2 was to provide a space where groups of people had sufficient screen real estate and resolution to show multiple representations, from pdfs and tables of numbers, through movies and high-resolution images, to fully immersive 3D spaces, or all of those at the same time. We traded off increased immersion (no visuals on the floor) to gain the convenience of having the lights on in the room and rapidly moving tables and chairs in the space to try and create a modern War Room / Project Room [22] for collaborative work. As described below, The ENDURANCE team gave us an opportunity to first try out this model.

While the CAVE2 provides the hardware for an Immersive Analytics environment, we use SAGE (and now SAGE2) [20] and OmegaLib [8] as the software. OmegaLib, built on top of Equalizer, VTK, OpenSceneGraph, C++ and Python is the open source software we have developed to drive CAVE2 (and other devices) in fully immersive interactive mode. SAGE2, built primarily in Javascript and HTML5, allows us to run 36 interlinked web browsers in the CAVE2 as one single shared canvas where multiple users can interact simultaneously, adding content (pdfs, movies, images, javascript applications) to the walls, moving, resizing, and interacting with that content, and sharing their desktops. Users interacting with the immersive 3D world can use a tracked controller, while other members of the team are simultaneously interacting through their laptops or tablets.

Running both SAGE and OmegaLib simultaneously allows the users to choose how much of the CAVE2 screen real estate they want to use for VR immersion, and how much they want to use for sharing different types of related documents. At points it is important for the entire space to be immersive tracked 3D, at other times a mix of immersive 3D and other documents, and at other times no immersive 3D on the screens. One of the major lessons we learned in the mid 1990s with the original CAVE [6] was that it was at best extremely difficult to integrate multiple useful representations into the same virtual world. Some data representations fit naturally into that paradigm, and others are best left in 2D. Collaborators from different disciplines want to see their data in familiar ways so multiple representations can often be better than a single shared representation. The resolution of the original CAVE made that difficult, but the resolution of newer room sized displays can accommodate and encourage this type of integration.

3 Case Study: ENDURANCE

In July 2013 EVL hosted the ENDURANCE team in our CAVE2 Hybrid Reality Environment [9]. We have been working with the NASA-funded ENDURANCE team since 2007 to explore ice covered Lake Bonney in the McMurdo Dry Valleys of Antarctica. This work involved the team sending an Autonomous Underwater Vehicle under the ice in 2008 and 2009 to take sonar readings of the bottom and collect chemical data, as a precursor to NASA doing similar work on Jupiter's moon Europa. The ENDURANCE team had previously used EVL's large displays to plan the mission, looking at QuickBird Satellite imagery on our 100 megapixel wall, and then later validating the data in a multi-disciplinary meeting on our

*e-mail: gmarai@uic.edu
†e-mail: aforbes@uic.edu
‡e-mail: ajohnson@uic.edu

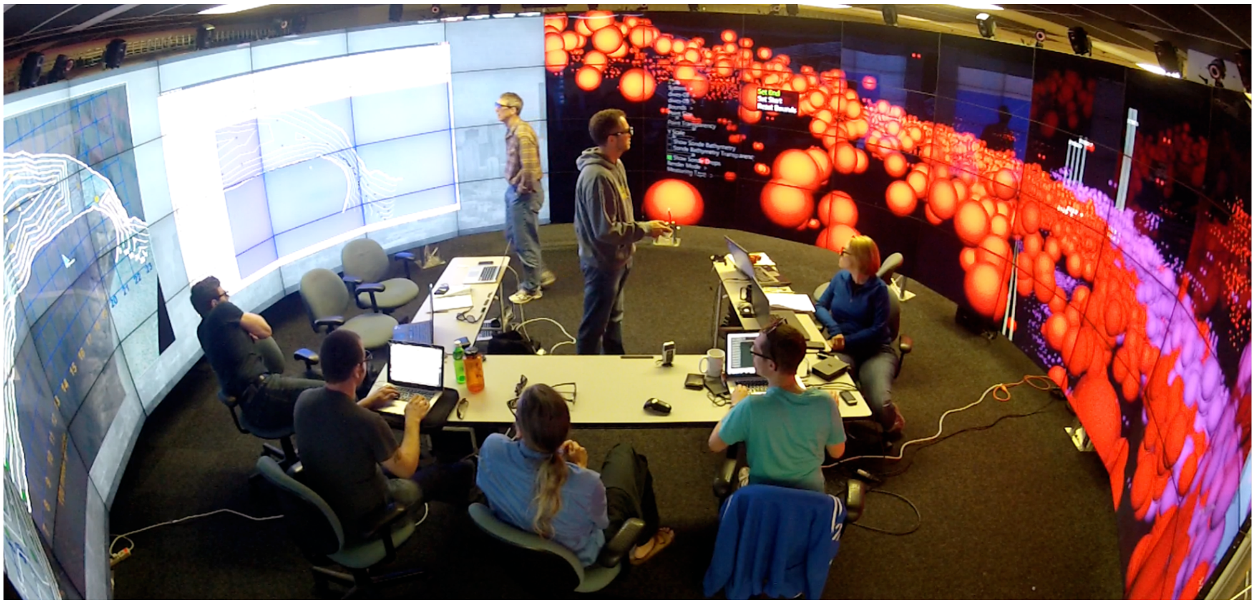

Figure 1: The ENDURANCE team examines lake data in CAVE2.

Cyber-Commons wall [15].

The ENDURANCE team spent two days working in CAVE2, allowing us to see how a multi-disciplinary team can work in an Immersive Analytics environment. During the meeting, team members sat at tables inside CAVE2 with their laptops. Different members of the team had different responsibilities and different expertise, and had brought their local data with them. The walls of CAVE2 were used for shared representations. Detailed data was kept private until is was needed and then uses could easily share their screen or drag and drop a relevant data file to the wall to add to the current conversation. The goal was to quickly answer questions about the data that had been collected and the processing that had been done on it.

One of the goals of the project was to create a detailed map of the bottom of the lake for the first time (See Fig. 1). This was particularly challenging as current sonar processing algorithms were not designed for this kind of environment and new algorithms needed to be tested. One way to test these was to 'dive' into the reconstruction of the lake. One of the team members has scuba dived in the real lake Bonney and wanted to swim through the lake at 1:1 scale to evaluate the sonar reconstruction and make changes to that reconstruction interactively. We were able to link the changes he made in the immersive 3D world to a VTK-based bathymetric representation of the lake shared on the other wall of CAVE2. The first person view was better for seeing the local area in detail, the bathymetric view gave the team a way to see what the overall contours looked like, and where they might be incorrect. He also had the ability to recolor the sonar points based on which dive they were collected on, and how far off axis they were so he could better judge the quality of the data. If he had a question about a particular dive and the actual sensor data he could ask someone in the room to look it up and show the results on another part of the screen. This created a very interactive session where different members could comment quickly and get answers quickly.

The large screen space also allowed subgroups to form when there was a particularly interesting question to answer. The subgroup could work on their own using their laptops and some of the shared CAVE2 screens while the rest of the team went on with their work using the rest of the space. At the end of the meeting one of the team members said that the team got more done in 2 days than in 6 months of email, Skype, and Google Hangout. He felt this was because we were altogether with our shared data and could quickly get answers which led to other questions that we could quickly get answers about. The space helped keep the team productive [19].

4 Case Study: Dark Sky

In July 2015 an interdisciplinary team of EVL researchers set out to develop a visual analysis tool for large-scale cosmological simulations of dark matter formation [12]. The data and required tasks were provided by the Dark Sky project[1] hosted by Stanford University; the project had been awarded a DOE INCITE computing allocation at the level of 80M cpu-hours. The largest simulations run from this project cover nearly 12 Gigaparsecs on a side (38 billion light-years across), and use 1.1 trillion particles to discretize the volume, totaling nearly half a Petabyte of output.

The tool developed by the team consists of an immersive linked multiview display which allows domain experts to interact with visual representations of spatial and nonspatial cosmology data (Fig. 2). There are three primary types of data involved in this project. The first is the raw particle data which is described by a position vector, velocity vector, and unique particle identifier. The second type of dataset is one Halo Catalog for each snapshot of time; each catalog groups sets of gravitationally bound particles together into coherent structures. Along with information about a given halo's position, shape, and size, the catalog contains a number of statistics derived from the particle distribution, such as angular momentum and relative concentration of the particles. The final dataset type links the individual halo catalogs, thereby creating a Merger Tree database. These merger tree datasets form a sparse graph that can then be analyzed to better understand how galaxies form and evolve through cosmic time.

We used the next generation CAVE2 immersive environment, and the D3 API and OmegaLib framework for virtual reality to display tree data, respectively 3D particles and halos. The D3 nonspatial views were projected into the immersive environment. Nonspatial data was represented as time-aligned merger trees, and through a pixel-based heatmap. Spatial data was represented through GPU-accelerated point clouds and geometric primitives. The user could

[1] http://darksky.slac.stanford.edu/scivis2015/

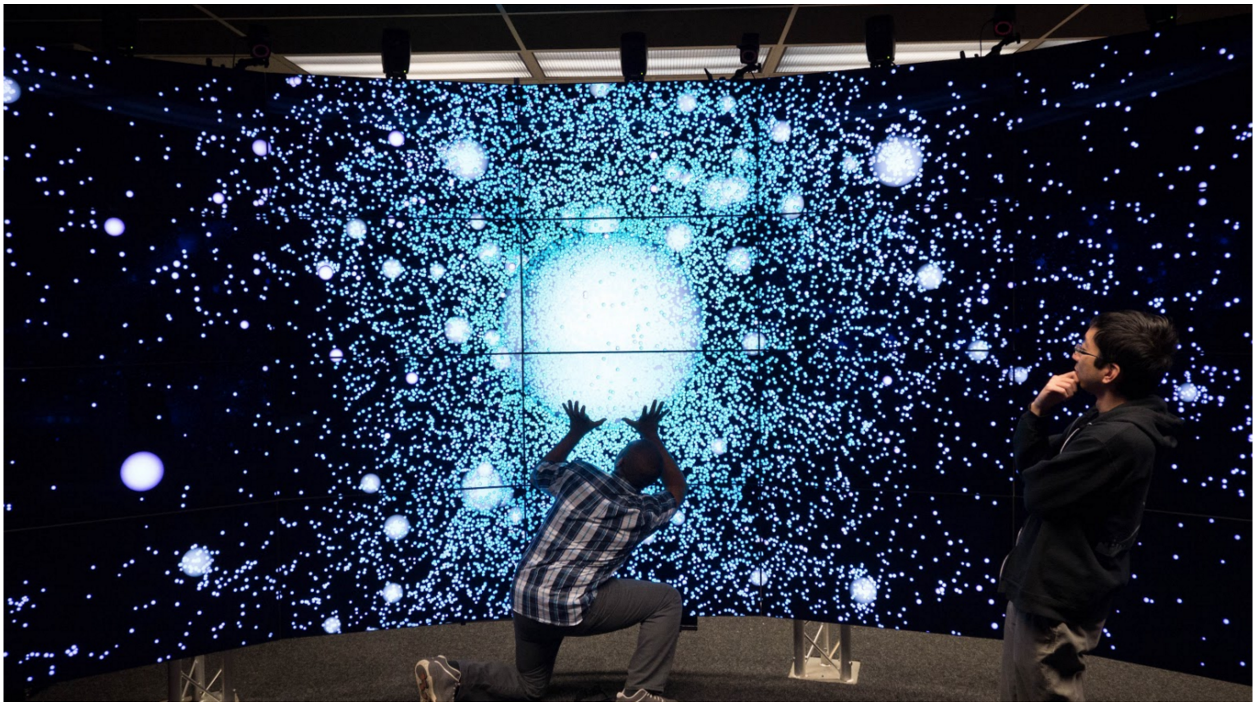

Figure 2: Young researchers examine as a subgroup large scale cosmology data in CAVE2

select a halo and visualize a 3D representation of the raw particles, as well as the halos at the particular time stamp. The interaction and a communication channel between D3 and OmegaLib allowed spatial and nonspatial views to be linked effectively [16]. We further implemented a 3D time lapse function, which can overlap several selected timesteps to show the flow and path of the halos and/or particles over time (Fig. reffig:timelapse). The time lapse creates a static 3D representation of the merger trees. The representation can also be animated to show the halo formations at each timestep. While the animation is playing, the user can freely move through the environment and zoom in on a desired halo formation.

The entry point of this application is a pixel heatmap of the merger tree forest, i.e. the collection of all merger trees in the simulation data. From this overview of the tree data, the user can select a particular tree of interest. A second view shows a 2D representation of the selected tree, with time mapped on the horizontal axis. From here, the user can select a particular time step, or a time lapse interval, and immersively explore, in the third view, the corresponding 3D particle set and halos. The time lapse visualization can be used to show all the merger-trees in 3D space. Alternatively, the user may start directly with the immersive exploration of a particular timestep or time lapse interval and browse the corresponding 3D particle set and halos, then navigate to the merger tree data. Our application attains a reasonable rate of 60 fps.

To evaluate the usefulness of the overall application, we have demonstrated the tool to several groups of visitors, as well as to a senior domain expert from the Adler Planetarium, who has significant experience in immersive environments. The expert found the interaction and flow to be "very smooth". The expert remarked the density and distribution of particles inside halos "showed well the power of the halo". The expert further appreciated the ability to analyze the relationship between mass and size by turning the particles off, since "halos could be very compact and still have high mass". Overall, the domain expert was impressed with the application, which was found to be "very nice", and was keen to show it to colleagues at the Planetarium.

We observed that the large screen allowed visitor subgroups to analyze together the data when a particularly interesting observation was made—for example, the fact that some halos evolve in parallel, never merge, and dissipate. Navigation in the virtual environment came natural. We further noticed that users never lost track of the context of the data they were examining, despite the large scale of the data and their initial unfamiliarity with it; in fact several users were able to navigate towards an interesting area, and then do a precise 180 degree turn (possibly using muscle memory), and return to their previous location.

5 Case study: BRAINtrinsic

Magnetic resonance (MR) imaging techniques such as functional Magnetic Resonance Imaging (fMRI) and diffusion weighted imaging (DWI) enable neuroimagers to collect and derive data about how different brain regions connect from both a structural and a functional point of view [14]. Analogous to the concept of genome for genetic data, a brain connectome is a whole-brain comprehensive map of neural connections [21]. As neural connections exhibit complex patterns of function and structure, the field of brain connectomics has emerged in order to understand these imaging big data.

EVL researchers have collaborated with scientists from the UIC Department of Psychiatry and the UIC Department of Bioengineering over the last two years in order to investigate effective approaches toward representing multimodal neuroimaging data that take advantage of the opportunities that immersive systems provide (see Fig. 4). The brain connectome is typically mathematically represented using connectivity matrices that describe the interaction among different brain regions. EVL researchers developed interactive visual analytics systems that enable clinical neuroscientists to identify patterns in the neuroimaging data of a single individual [5, 10]. Specifically, our *BRAINtrinsic* system enables comparison tasks in which an individual patient's connectome can be compared to other patients or to averaged datasets (such as those that are available via from the Human Connectome Project data

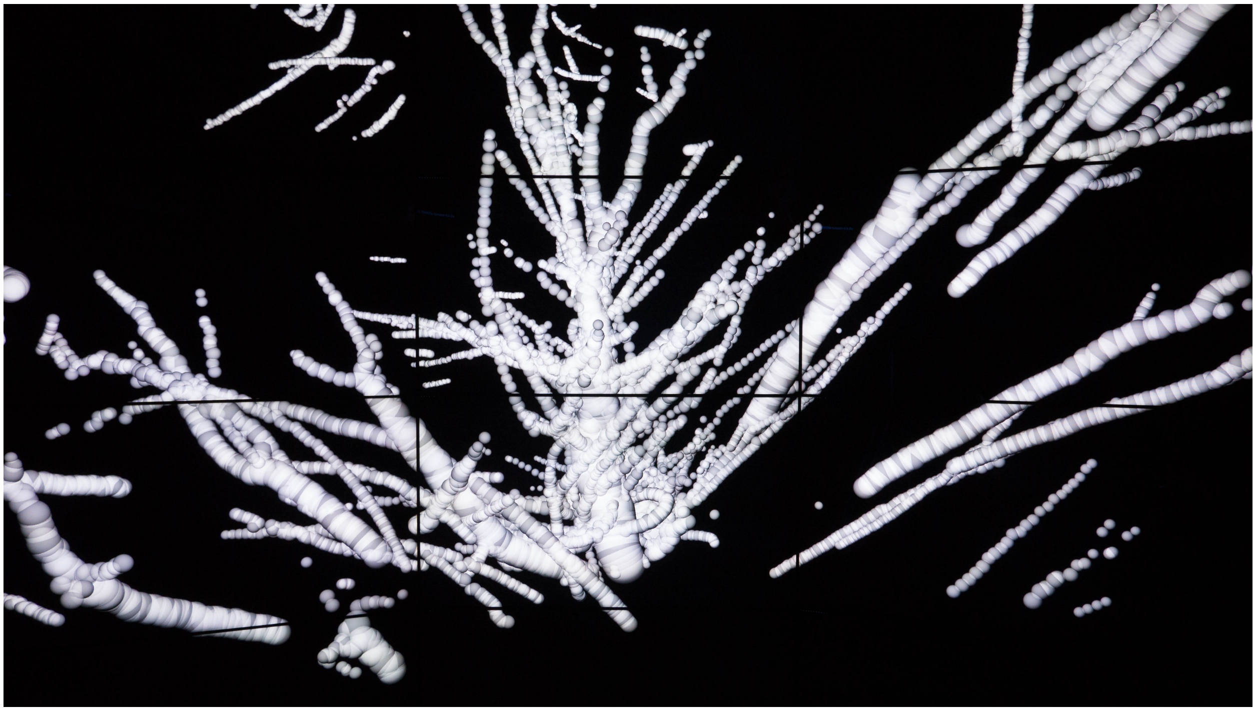

Figure 3: Timelapse visualization of halo merge trees. A group of visitors noted that some halos evolve in parallel and dissipate.

repository[2]).

Although the effectiveness of utilizing 3D for representing data has been debated [24], Alper et al. [1] have shown that in some situations visualizing 3D networks can outperform 2D visualizations. *BRAINtrinsic* introduces a dynamic and interactive VR-compatible visualization platform for connectome representation and exploration utilizing the intrinsic geometry of the connectome [4]. The intrinsic geometry represents the brain connectome after non-linear multidimensional data reduction techniques are applied to identify meaningful relationships between non-contiguous brain regions. The position of nodes in our visualization is based on the strength of the interaction that each brain region has with the rest of the brain, whether structural or functional. A user can easily switch between representations embedded in a neuroanatomical view and intrinsic geometry representations created via various dimensionality reduction pipelines.

Initial studies using *BRAINtrinsic* have helped researchers identify details about the importance of the "Rich-club property (where nodes with high nodal strengths to form tightly interconnected groups [23]) in the human connectome. Our Immersive Analytics application has also facilitated the exploration of differences between structural and functional resting state connectomes in patients with psychiatric disorders [25]. Fig. 5 shows a psychiatrist using the desktop version of *BRAINtrinsic* to explore resting state connectome data.

6 IA Challenges and Directions for Future Research

In addition to the three case studies reported above, we have also used CAVE2 to judge our university's Image of Research contest, have held meetings related to evaluating different user interfaces and reviewing data for an information system for nurses, and have held graduate classes in the space. Many of these lessons we have learned reinforce the lessons learned from the War Room research of the 1990s. However, with almost all of our information living its entire life in a digital form, these lessons emphasize that the way we access and interact with information has changed. We summarize below the challenges that became apparent through our work in IA.

6.1 Hardware Resolution

Immersive Analytics require enough screen real estate to show multiple representations simultaneously. They also require enough resolution to show context plus detail and for text to be easily readable. The environment further needs the ability to show 3D everywhere, for everyone in the room, even if 3D is not needed all of the time for analysis.

6.2 Interaction Support for IA Group Work

Immersive Analytics tasks require the ability to link representations together, to quickly annotate, and to brainstorm as a group, as though standing together at a whiteboard. Users should be able to quickly add new information to the conversation, to save the state of a session and bring it back the next day, or next year. They should further have the ability to quickly transition from controlling the immersive space to the 2D space, ideally using the same controller that knows what it is pointed at and acts accordingly [18].

6.3 Tracking Support for Group Work

Immersive Analytics require the ability for subgroups to form, work on a particular problem quickly (in 2D or immersive 3D), using part of the space, then bring their findings back to the group. The ability to track multiple users and multiple controllers is important, as is having enough tracked glasses and controls to support simultaneous work. The users should further have the ability to interact at same time and not have to take turns or swap controls.

6.4 Quiet, Comfortable Environments

Immersive Analytics require quiet and cool rooms with enough light to see laptop keyboards or read paper notes. Analysts want to be able to bring in their lunch or beverage while working. It is

[2]http://www.humanconnectome.org/data/

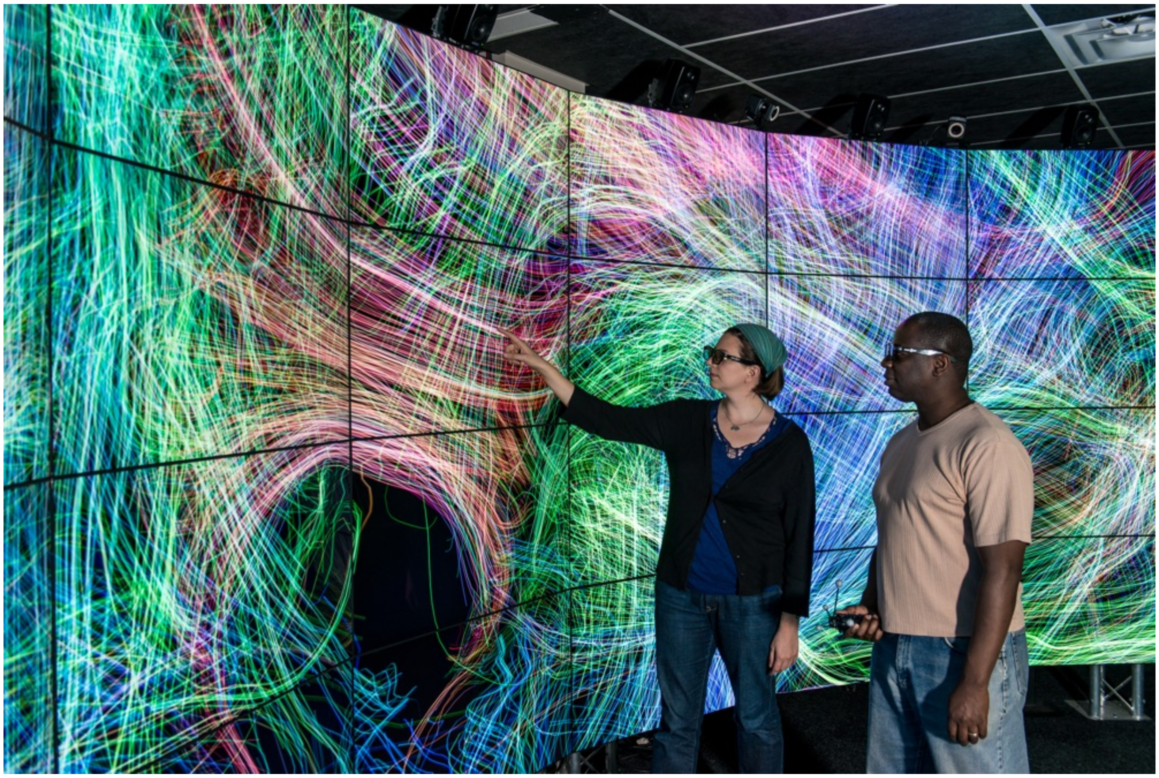

Figure 4: UIC graduate students explore a representation of connectome data in 3D using CAVE2.

further necessary to be able to quickly reconfigure tables and chairs, and the users should feel comfortable working in the space for 8+ hours straight.

6.5 Augmenting IA with Spatial Audio

In addition to the interaction and visualization techniques described above, the judicious use of spatial audio in virtual environments can function to capture a user's attention to aid navigation in VR systems [11]. Moreover, initial investigations into using 3D audio illustrates its potential for conveying additional channels of information during Immersive Analytics tasks [17]. EVL researchers are currently exploring the use of 3D user interfaces for designing effective virtual audio environments that augment the graphics capabilities of the CAVE2 [2, 3].

7 Conclusion

In conclusion, we presented in this work several lessons learned from three recent successful Immersive Analytics case studies performed on scientific data in the CAVE2 hybrid immersive environment. These case studies demonstrate that immersive environments can augment the humans' ability to analyze and make sense of large and multifaceted datasets. We believe that there is a bright future for visual analytics using VR technologies.

Acknowledgements

We gratefully acknowledge our wonderful collaborators across disciplines, who made this work possible. We further thank our students who worked with us in this project, in particular K.Reda, P.Hanula, K.Piekutowski, C.Uribe, A.Nishimoto, and K.Almryde. This work was partially supported by federal grant NSF CAREER IIS-1541277 and by a small Computing Research Association CREU grant.

References

[1] B. Alper, T. Höllerer, J. Kuchera-Morin, and A. G. Forbes. Stereoscopic highlighting: 2D graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2325–2333, November-December 2011.

[2] A. Çamci, P. Murray, and A. G. Forbes. Designing and controlling virtual sonic environments using a browser-based 3DUI. In *Proceedings of IEEE Symposium on 3D User Interfaces (3DUI) Poster Session*, Greenville, South Carolina, March 2016.

[3] A. Çamcı, Z. Özcan, and D. Pehlevan. Interactive virtual soundscapes: A research report. In *Proceedings of the 41st International Computer Music Conference*, pages 163–169, 2015.

[4] G. Conte, A. Ye, K. Almryde, O. Ajilore, A. Leow, and A. G. Forbes. Intrinsic geometry visualization for the interactive analysis of brain connectivity patterns. In *Visualization and Data Analysis (VDA)*, Proceedings of IS&T Electronic Imaging, pages 481–1–8. San Francisco, California, February 2016.

[5] G. Conte, A. Ye, A. G. Forbes, O. Ajilore, and A. Leow. BRAINtrinsic: A virtual reality-compatible tool for exploring intrinsic topologies of the human brain connectome. In Y. Guo, K. Friston, A. Faisal, S. Hill, and H. Peng, editors, *Brain Informatics and Health*, volume 9250 of *Lecture Notes in Artificial Intelligence*, chapter 7, pages 67–76. Springer, 2015.

[6] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '93, pages 135–142, New York, NY, USA, 1993. ACM.

[7] M. Czernuszenko, D. Pape, D. Sandin, T. DeFanti, G. L. Dawe, and M. D. Brown. The ImmersaDesk and Infinity Wall projection-based virtual reality displays. *ACM SIGGRAPH Computer Graphics*, 31(2):46–49, 1997.

[8] A. Febretti, A. Nishimoto, V. Mateevitsi, L. Renambot, A. Johnson, and J. Leigh. Omegalib: A multi-view application framework for hybrid reality display environments. In *Proceedings of IEEE Virtual*

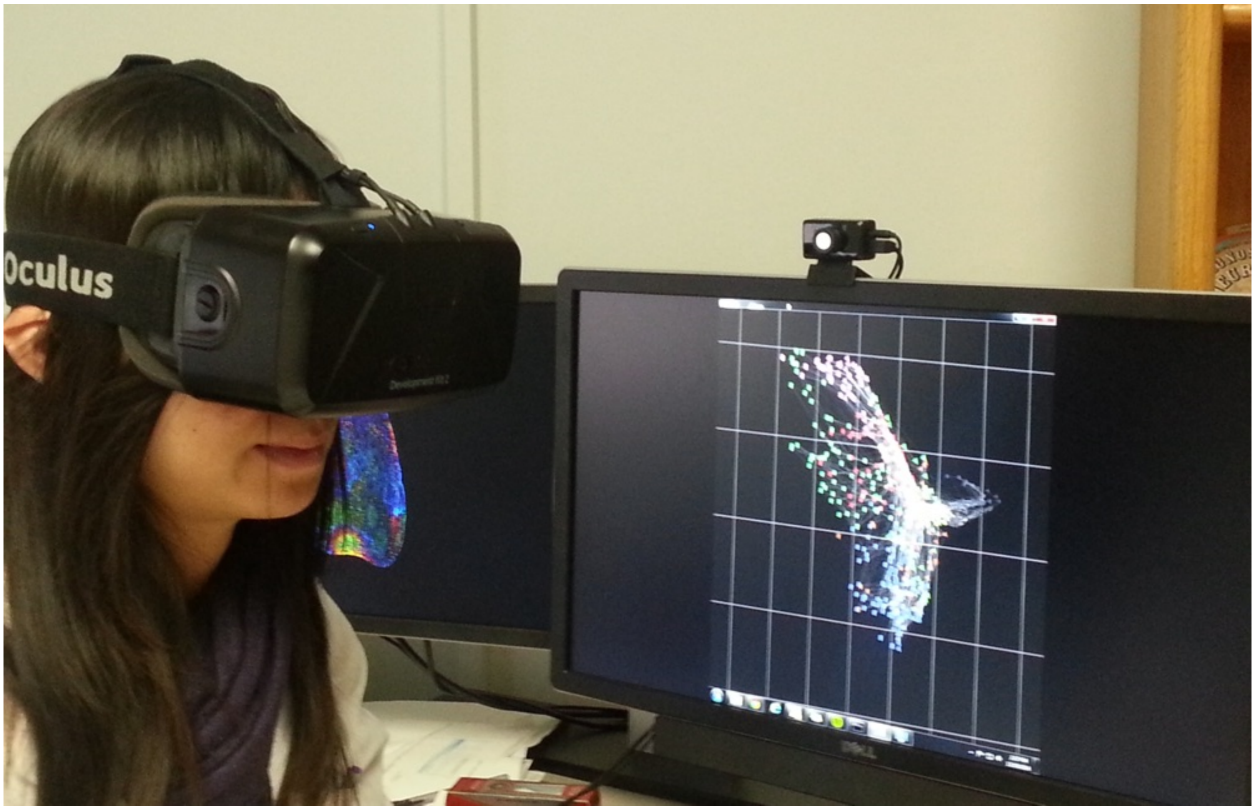

Figure 5: A researcher interactively explores the intrinsic geometry of the human brain connectome in 3D using the BRAINtrinsic application created at EVL.

Reality (VR), pages 9–14, 2014.

[9] A. Febretti, A. Nishimoto, T. Thigpen, J. Talandis, L. Long, J. Pirtle, T. Peterka, A. Verlo, M. Brown, D. Plepys, D. Sandin, L. Renambot, A. Johnson, and J. Leigh. CAVE2: A hybrid reality environment for immersive simulation and information analysis. In *The Engineering Reality of Virtual Reality*, volume 8649 of *Proceedings of IS&T/SPIE Electronic Imaging*. San Francisco, California, February 2013.

[10] A. G. Forbes, J. Villegas, K. Almryde, and E. Plante. A stereoscopic system for viewing the temporal evolution of brain activity clusters in response to linguistic stimuli. In A. J. Woods, N. S. Holliman, and G. E. Favalora, editors, *Stereoscopic Displays and Applications XXV*, volume 9011 of *Proceedings of SPIE-IS&T Electronic Imaging*, pages 90110I–1–7. San Francisco, California, February 2014.

[11] F. Grani, S. Serafin, F. Argelaguet, V. Gouranton, M. Badawi, R. Gaugne, and A. Lécuyer. Audio-visual attractors for capturing attention to the screens when walking in CAVE systems. In *IEEE VR Workshop on Sonic Interaction in Virtual Environments (SIVE)*, pages 75–76, 2014.

[12] P. Hanula, K. Piekutowski, C. Uribe, A. Nishimoto, K. Almryde, J. Aguilera, and G. E. Marai. Cavern halos: Exploring spatial and nonspatial cosmological data in an immersive virtual environment. In *Virtual and Augmented Reality, 3D, and Stereoscopic Systems Conference Poster Compendium, Electronic Imaging'16*, 2016.

[13] A. Johnson, D. Sandin, G. Dawe, T. DeFanti, D. Pape, Z. Qiu, S. Thongrong, and D. Plepys. Developing the PARIS: Using the CAVE to prototype a new VR display. In *Proceedings of IPT 2000: Immersive Projection Technology Workshop*, pages 19–20, 2000.

[14] D. K. Jones. *Diffusion MRI: Theory, methods, and applications*. Oxford University Press, 2010.

[15] J. Leigh and M. D. Brown. Cyber-commons: Merging real and virtual worlds. *Communications of the ACM*, 51(1):82–85, Jan. 2008.

[16] G. E. Marai. Visual scaffolding in integrated spatial and nonspatial analysis. In E. Bertini and J. C. Roberts, editors, *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*, pages 13–17, 2015.

[17] K. McMullen. The potentials for spatial audio to convey information in virtual environments. In *IEEE VR Workshop on Sonic Interaction in Virtual Environments (SIVE)*, pages 31–34, 2014.

[18] K. Reda, A. E. Johnson, J. Leigh, and M. E. Papka. Evaluating user behavior and strategy during visual exploration. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pages 41–45. ACM, 2014.

[19] K. Reda, A. E. Johnson, M. E. Papka, and J. Leigh. Effects of display size and resolution on user behavior and insight acquisition in visual exploration. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2759–2768. ACM, 2015.

[20] L. Renambot, T. Marrinan, J. Aurisano, A. Nishimoto, V. Mateevitsi, K. Bharadwaj, L. Long, A. Johnson, M. Brown, and J. Leigh. SAGE2: A collaboration portal for scalable resolution displays. *Future Generation Computer Systems*, 54:296–305, 2016.

[21] O. Sporns. The human connectome: A complex network. *Annals of the New York Academy of Sciences*, 1224(1):109–125, 2011.

[22] S. Teasley, L. Covi, M. S. Krishnan, and J. S. Olson. How does radical collocation help a team succeed? In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 339–346. ACM, 2000.

[23] M. P. van den Heuvel and O. Sporns. Rich-club organization of the human connectome. *The Journal of Neuroscience*, 31(44):15775–15786, 2011.

[24] C. Ware. *Information visualization: Perception for design*. Elsevier, 2012.

[25] A. Q. Ye, O. A. Ajilore, G. Conte, J. GadElkarim, G. Thomas-Ramos, L. Zhan, S. Yang, A. Kumar, R. Magin, A. G. Forbes, and A. D. Leow. The intrinsic geometry of the human brain connectome. *Brain Informatics*, 2(4):197–210, December 2015.