

THEME ARTICLE: Special Issue on Immersive Analytics

Immersive Analytics Lessons from the Electronic Visualization Laboratory: a 25 Year Perspective

G.E. Marai

Electronic Visualization
Laboratory, University of
Illinois at Chicago

Jason Leigh

Laboratory for Advanced
Visualization and
Applications, University of
Hawaii at Manoa

Andrew Johnson

Electronic Visualization
Laboratory, University of
Illinois at Chicago

This article provides a 25 year-long perspective on Immersive Analytics through the lens of first-in-kind technological advancements introduced at the Electronic Visualization Laboratory, University of Illinois at Chicago, along with the challenges and lessons learned from multiple Immersive Analytics projects.

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 Virtual Environment, the first projection-based virtual reality system in the world. Since then, the lab has developed a wide range of immersive systems and technologies to be used for visualization and visual analysis tasks, such as PARIS, ImmerSaDesk, CAVE2 [1], and SAGE2 [2]. These technologies paved the way for the resurgence in virtual reality in the past few years, and lead to the introduction of Immersive Analytics—which investigates how new interaction and display technologies can be used to support analytical reasoning by immersing the users in data.

reasoning by immersing the users in data.

At the same time, many interaction and display technologies have intrinsic limitations in their applicability to Immersive Analytics. Immersive environments attempt to simulate interaction and perception as realistically as possible, primarily through technologies such as 1) stereo imagery delivered separately to each eye and 2) head and body tracking so that the environment responds immediately to the user's physical moves. The most common immersive technologies are head-mounted displays, or small rooms with rear-projection displays on walls, floor, and/or ceilings. These instantiations of immersive technology have been criticized in the visual analysis literature first and foremost for their limited resolution — the number of available pixels divided by the display area. Because resources are necessary for immersion, these displays “cannot show as many pixels as state-of-the-art desktop displays of the equivalent area” [3]. The number of pixels available is a critical and limited resource in visualization design, and thus visual analysts are rarely willing to trade resolution for immersion [3]. Immersion is most useful when a sense of presence is an important aspect of the intended analytical task.

The second major limitation of many immersive technologies is the typical lack of integration of visualization with the rest of an analyst's workflow. High-resolution immersive technologies are almost always special-purpose devices that are at a different physical location than the user's workspace. Using these high-resolution devices requires users to leave their office or laboratory and go to some other location, whether down the hallway, in a different building, or, occasionally, at a different geographical location. At these locations, analysts are typically required to stand rather than sit, which can be taxing over many hours. The most critical problem, however, is that in these environments analysts typically do not have access to their working environment of their own computer, including access to applications such as web browsing, email, text and spreadsheet editing, and data analysis packages; almost all such standard applications rely on input devices such as a mouse and keyboard. The lack of integration with a standard desktop environment, along with the ability to switch rapidly between visualization and other applications is a serious limitation in Immersive Analytics [3].

In this paper, we describe EVL's endeavors to address these and other limitations in Immersive Analytics, along with the challenges and the lessons learned from recent EVL projects. These collaborative projects primarily make use of the CAVE2 immersive environment (Figure 1), although some are also portable to 3D desktop environments or can run on head-mounted displays. We further summarize some of the opportunities ahead for the field of Immersive Analytics.



Figure 1. The CAVE2 is a 320-degree hybrid environment. After each of two rounds of data collection in Antarctica, the 8-person NASA-funded ENDURANCE team (not all pictured here) held a group meeting at EVL over several days. They used CAVE2, SAGE2, and OmegaLib to, among other things, work out the bathymetry of the bottom of a lake from the sonar collected by a robot. Through a combination of a 3D 1:1 scale representation of the sonar data on right-half of CAVE2, a 2D dynamically updated bathymetric map in VTK in the left-half, and 2D details about the missions on the walls, along with more detailed data on their laptops, the environment gave the team a productive analytical space to work in.

ENVIRONMENTS FOR IMMERSIVE ANALYTICS: EVOLUTION

Immersion: CAVE

From a historical perspective, when the first CAVE environment (Figure 2) was under development at EVL in 1991, its large projection screens were, even then, being used for collaborative debugging of the code that ran the CAVE itself. Once the CAVE was completed, scientists in a variety of disciplines began to make use of the environment, both individually and in groups, to get a new interactive 3D view of their data. The data was either stored locally or streamed from supercomputing resources. The environment gave analysts a place to discuss their data while being immersed within it, with the added advantage of being able to perceive a co-located collaborator's location, posture, and to some extent their facial expression. The CAVE's hand-held controller evolved quickly through multiple iterations to give analysts easy control over navigation and interaction in this unfamiliar space. By the mid 1990s, as networks improved and CAVE's and their related devices multiplied, users could share the same virtual space with remote users (depicted as avatars) and work with more common visualization libraries such as VTK, the visualization toolkit.

The CAVE did have drawbacks—it was hard to stand for a long time, users suffered from arm fatigue when using a wand, the fan noise from the projectors and the active stereo glasses made it hard to work all day inside. Scientists wanted more resolution, especially to be able to display numbers alongside the visualizations, as well as brighter, more accurate colors. Last but not least, only one application could run at a time, forcing developers to try and fit all the relevant data into a single visual representation. Nevertheless, CAVE systems found success

in the collaborative design realm, and the spin-off lower-end GeoWall system found general success in the educational domain, in particular as part of geoscience labs and in museums.




Figure 2. EVL's projection-based, first-in-kind CAVE VR system was used in conjunction with the Visualization Toolkit (VTK) to interactively analyze geospatial data collected from the Chesapeake Bay. Note the wired glasses and controller, and the low-resolution of the environment. The CAVE size was 10 ft (or 3m) cubed.

Resolution and Collaboration: LambdaVision, SAGE and SAGE2

In response to some of the critical issues surrounding CAVE systems, in the early 2000s EVL's research shifted towards high-resolution, large tiled displays connected by high-speed networks, such as the 100-megapixel LambdaVision wall (2005) and the smaller 24-megapixel table format LambdaTable (2004). While these displays originally did not have stereo display or head and hand tracking capabilities, they soon proved to be capable of immersing users in their data through sheer size, resolution, and interactivity. With a size of 17 ft wide by 6 ft tall (5m by 1.98m), the LambdaVision wall would cover the field of view of several analysts standing in front of the wall, effectively giving the analysts a sense of immersion. These kinds of displays were successfully used to show extremely high-resolution images or visualizations (Figure 3), not only at EVL but also at partner institutions who were using EVL's software on their own walls. For example, analysts at the National Center for Microscopy and Imaging Research used this hardware and software to examine scans of mice brains down to the level of single neurons; at the Lamont-Doherty Earth Observatory to visualize kilometers of core-drilling data; and at the US Geological Survey to check the validity of new aerial photography of US cities after 9/11.

Yet our research partners, analysts across science and engineering, were increasingly asking to see multiple related datasets simultaneously, as opposed to single high-resolution images. To investigate these issues and try and satisfy this demand in a more general way, EVL created the SAGE software, to allow groups of people to simultaneously interact with multiple digital artifacts on a large shared high-resolution wall. We wanted to create a display and interaction framework to give collaborators easy access to enough physical display space and resolution to see the data in a representation familiar in their disciplines, and make sure that important features did not get lost — those features stayed on the screen until they were dealt with. Making more screen real estate available meant there was more room for individuals to externalize their own representations, leading to more productive conversations. Based on the lessons we learned from networking CAVE environments, there was no turn taking, no token passing, and no permissions to take control. Analysts were simply using social cues to avoid collisions. As in the CAVE, much of the visualization software was custom built, but

common image and movie formats could also be displayed through SAGE. Unlike the CAVE, these walls were high resolution, with brighter and more accurate colors, though with the tradeoff of border bezels on all LCD displays. Multiple co-located users could now individually interact with multiple related pieces of data (text, images, simulations) at the same time.

Our software was further evolving to support analyst workflow integration with immersive displays. Unlike in the CAVE, analysts typically sat in chairs at tables in front of walls like LambdaVision. This setting allowed analysts to use their laptops and phones as private displays, and the wall as a public display, and interact with the wall through their touchpads or mice. With the enhanced resolution, analysts would sometimes walk up to the wall, and have their conversations nearer to the visualizations. So, touch became an important modality when close to the wall, as well as air mice located near the wall. Last but not least, whereas developers had to write specific applications to run in the CAVE, analysts could use SAGE to run an application on their own laptop, and screen-share that content to the wall.

The second generation SAGE2 middleware [4] further lowered the barrier to entry for using these kinds of high-resolution large displays, by building on top of web-based technology such as Javascript and HTML5, and by using a web browser as user client and web servers to run the displays. Any large display could be a SAGE2 display or part of a tiled SAGE2 display; any computer can run a SAGE2 display or be part of a cluster of SAGE2 displays. Analysts could connect to those displays through a web browser on their laptop or tablet. As well as making it easier to share content onto the display, and copy content off of the display onto their local machine, web technologies made it easier to leverage the science that has moved into portals and notebooks on the web.



Figure 3. Before the NASA-funded ENDURANCE team travelled to Antarctica, a 10-person team (not shown here) came to EVL and used SAGE and LambdaVision to examine related high-resolution imagery (including QuickBird photos) of the lake they were going to collect data from. The team successfully used this analysis to plan where they would set up camp and where they would need to melt through the ice so that the NASA robot would be able to enter the lake. Note the higher-resolution and less-intrusive technology of the environment, despite the visible display bezels. The LambdaVision wall was 17 ft wide by 6 ft tall (5m by 1.98m).

Embodied Hybridization: CAVE2

CAVE2, unveiled in 2012, was designed and built based on these 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. CAVE2 (Figure 1, Figure 4) is a 70-megapixel 2D / 35-megapixel passive 3D hybrid reality environment, that since being created at EVL has been adopted at multiple institutions across the world. Originally 36 computers, now 18, drive 72 LCD panels (18 columns of 4) arranged in a 320-degree circle with diameter 22 ft (6.7m). 14 Vicon tracking cameras allow us to track 20 objects (glasses, controllers, or anything with a rigid body marker attached to it) in the space, while 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 tables of numbers and pdf documents, 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 of rapidly moving tables and chairs in the space to try and create a modern Project Room [5] for collaborative work. In this work space, post-it notes and flip charts were replaced by digital artifacts on the walls. Unlike physical project rooms, this environment allowed analysts to save the state of the room, let another group use it, and then bring the room back to the previous configuration.

While the CAVE2 provides the hardware for an Immersive Analytics environment, we use SAGE (and now SAGE2) [4] and OmegaLib or Unity3D for the environment's 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 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 in OmegaLib or Unity3D can use a tracked controller, while other members of the team are simultaneously interacting through their laptops or tablets.

Running both SAGE2 and OmegaLib or Unity3D 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 times it is important for the entire space to be immersively 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. In our experience, 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 accommodates and encourages this type of integration. Furthermore, SAGE2 allows remote groups to share pointers, windows, or entire walls, and since SAGE2 is web-based it allows us to use commonly available web-based teleconferencing sites and place small webcams in the CAVE2 for remote collaborators to see and talk to each other about the shared content.



Figure 4. CAVE2 immersive exploration of dark matter formation in the universe. Note that while little information can be shown on a typical desktop display (one rectangle in the 320-deg. wall shown), the high-resolution tiled environment allows analysts to examine detailed information without losing the context of the larger dataset. Analysts (team of 7, not all shown here) also appeared to use kinesthesia to navigate through this space.

Into the Future: Cyber-Commons and Continuum

From an analytics perspective, we found that CAVE2 was very good for small group meetings (5-10 people), but that larger groups needed a larger physical space, especially when they needed to break up into smaller groups for more focused analysis work. Starting back in 2000, EVL began to research larger spaces for embodied collaborative Immersive Analytics—collaborative spaces enhanced with 3D capabilities and tracking. To this end, EVL linked our shared visualization environments with distance collaboration environments like Argonne National Laboratory's Access Grid project, which enabled multi-site multi-camera and multi-microphone group-to-group interaction. We set up this environment in EVL's large (40 ft by 20 ft, or 12.19m by 6m) meeting room space, to make it easier to hold distributed research meetings with colleagues. The projection screens for the multiple AccessGrid video windows were soon augmented with a variety of increasingly larger displays, some immersive, enabling us to begin holding in this room our visualization, user interaction, and virtual reality courses as well.

In 2009 a new tiled LCD wall display replaced the previous separate displays, and this second meeting space became officially the Cyber-Commons (Figure 5)—a digital distributed meeting place for discussions, research, and learning. An early version of the Cyber-Commons room incorporated a large floor-to-ceiling projection screen. The projection screen had both stereo immersion and head and hand-tracking using the classic CAVE technology. Cyber-Commons continued to evolve over the next decade through a variety of better tiled LCD displays, at the same time as SAGE was replaced by SAGE2. SAGE2 made it easy to save the state of a collaborative session and bring that state back when the collaboration resumed. Later display iterations used only non-tracked CAVE2-style panels, which had stereo immersion capabilities. Cameras and microphones were always available for remote connections. With a size of 20 ft wide by 6 ft tall (6m by 1.8m) and stereo display capabilities, although no hand or head-tracking, Cyber-Commons gave us a place for group analytics, where 30-40 people could meet, and where we could have tables and chairs ready to place in an appropriate configuration.

In 2017 the Cyber-Commons began another major evolution as part of the Continuum project (Figure 6), with a renovation to increase space, reduce HVAC noise, add multiple display walls and a variety of sensors to give the room more knowledge about the people inside, and in general, to turn the environment into an active assistant in the investigations taking place there. In conjunction with custom software, some of these sensors will

provide position, head and gaze tracking capabilities. Having multiple display-rich spaces available in EVL also allows for multiple simultaneous yet separate meetings and classes, each in the appropriate venue. We are currently running CAVE2, the Continuum, as well as a large tiled LCD wall in a second classroom space (since the Continuum is now booked up half of each week), and a smaller meeting room with three large 4K displays, all running SAGE2. We do not want our students or our local collaborators to think of these displays as a rare resource. We want students and collaborators to be able to pick the correct space for the task.

While EVL researches also a variety of head-mounted displays, our Immersive Analytics experience indicates that these devices are most useful when integrated into larger, high-resolution collaborative spaces. The use of integrated augmented and virtual reality could release analysts from being tethered to their laptops and phones towards using their own private data views within headsets, in the context of high-resolution data that is shared in a larger collaborative space.



Figure 5. After the first round of data collection in Antarctica, the 10 person ENDURANCE team, including earth scientists, geochemists, the builders of the robot that collected the data, and visualization experts, used EVL's Cyber-Commons and SAGE environment to initially assess the data. Multiple data representations were shared on the public wall, while the analysts maintained full access to more detailed data on their laptops. Note the reasonably seamless tiled displays, and the variety of interaction modalities used by analysts.



Figure 6. EVL's new Continuum room will combine: a large touch screen wall; a side-wall of 4K displays for extremely high-resolution visual analysis; a passive stereo wall, all driven by SAGE2; and a wall of whiteboards for quick collaborative brainstorming. The room is a quiet space with movable tables and chairs for quick reconfiguration, while a suite of cameras,

microphones, and sensors track the users so that the room can help with their investigations in an immersive and augmented reality space.

CASE STUDY: ENDURANCE AND COLLABORATIVE HYBRID REPRESENTATIONS

In July 2013 EVL hosted the NASA-funded ENDURANCE team in our CAVE2 Hybrid Reality Environment [1]. We have been working with the 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 (Figure 3), and then later validating the data in a multi-disciplinary meeting on our Cyber-Commons wall [7] (Figure 5).

During their third and last EVL meeting, the ENDURANCE team spent two days working in CAVE2 (Figure 1), 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; they had brought their local data with them. The walls of CAVE2 were used for shared representations. Detailed data was kept private until it was needed and then users could easily share their screen or drag and drop a relevant data file to the wall to add to the current conversation using SAGE2. 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. This was particularly challenging as current sonar processing algorithms were not designed for this kind of environment with an extreme layer of salinity, 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, while the bathymetric view gave the team a way to see what the overall contours looked like, and where they might be incorrect. The diver 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 the team was sitting all together 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 [8].

CASE STUDY: DARK SKY AND COLLABORATIVE EMBODIED NAVIGATION

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 (Figure 7) [9]. The data and required tasks were provided by the Dark Sky project hosted by Stanford University, using data from the San Diego Supercomputing Center. The team consisted of visualization and astronomy researchers.

One of the goals of this project was to model dark matter, a collisionless fluid, as a discretized set of particles that interact only gravitationally. Such a simulation requires a large number of particles — typically on the scale of 10K to 100K particles. Over the 14 billion years of evolution, these particles cluster into gravitationally bound structures that pull in baryonic matter that forms stars, galaxies, and clusters of galaxies. Developing visualizations for the structures formed by these particles through gravitational interaction and collapse, required first identifying the structures, developing appropriate representations of the components or the structures themselves, and then correlating these visualizations across time steps.

Typically, dark matter structures are identified through a process known as halo finding. Halos may represent galaxies or clusters of galaxies. Through this finding process dark matter halos are identified either via local particle density estimation or through simple linking-length mechanisms. Within a galaxy cluster, smaller halos (substructures) may be identified which correspond to the location of galaxies. As these structures and substructures interact, merge, separate and grow, the domain scientists believe that the structure of the Universe grows and changes along. Visualizing the state of the simulated halos during the lifetime of the Universe can provide necessary inputs to understanding observations from next generation telescopes.

There were 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 abstract views were projected into the immersive environment SAGE2. Abstract data was represented as time-aligned merger trees, and through a pixel-based heat map. Spatial data was represented through GPU-accelerated point clouds and geometric primitives. An analyst could 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 abstract views to be linked effectively [10]. 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 (Figure 7 inset). 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, an analyst can freely move through the environment and zoom in on a desired halo formation.

We observed two senior domain experts from the Adler Planetarium, over repeated visits, performing in depth analyses of the halo structures, as well as several groups of visitors. The Adler experts had significant experience with immersive environments. The astronomy investigating team worked side by side with visualization researchers, and made frequent use of laptops, personal notebooks, desks and chairs brought inside the CAVE2. During the expert workflow, the lights were almost always on, to better facilitate access to the scientists' computer workflow. The visualization researchers used their own laptops to repeatedly recode parts of the visual encodings, according to suggestions from the domain experts.

We further noted the use of multiple “natural” interaction styles as the analysts worked at different distances within the CAVE2 environment: pointing at the 2D representations on the display wall, using tracking and controllers near the wall or from the middle of the CAVE2, and interacting with both the immersive and the 2D representations via laptops from the back of CAVE2. Navigation in the virtual environment came naturally. We noticed that the analysts 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, they were able to navigate towards an interesting area, and then do a precise 180 degree turn (possibly using kinesthesia, i.e., awareness of the position and movement of the parts of the body by means of sensory organs, i.e., proprioceptors, in the muscles and joints), and return to their previous location.

Overall, the domain experts were extremely impressed with the application, and were keen to show it to colleagues at the Adler Planetarium. The ability to examine details within context — without getting lost, despite the scale of the data—was particularly appreciated. Finally, the experts brainstormed for ways to port the application to the Adler museum, where OmegaLib was already installed and a 3D display is in use.

The Halo project is one of the most popular demonstrations in our lab, and has been shown to several groups of visitors, ranging from 6 to 45 people at one time. For these exploratory analyses, no visitors make use of a computer workflow, and the lights in the room are off. We observed that the large screen allows visitor subgroups to analyze together the data when a particularly interesting observation was made. The fact that some halos evolve in parallel, never merge, and dissipate is often noted by visitors. For this group of users, as with the expert group, we also noted kinesthesia and embodied navigation of the data — navigation in which the sensorimotor capacities, body and environment play an important role. These types of analyses are more about presentation of science using large immersive displays rather than trying to work out the “truth” in these spaces. Such demonstrations allow groups of visitors to investigate topics in science, as an outreach example of getting these types of experiences out to regular people, at scales where groups of users (as typical in museums) interact and talk to each other.



Figure 7. Collaboratively examining in the CAVE2 immersive environment the formation of dark matter halos. Shown is a single timestep view of several large halos which are starting to cluster. Each tile of the immersive display is the size of a regular desktop display, indicating how little information and context could fit on a single display. The inset in the top right shows a time lapse detail of 89 timesteps, which shows the paths taken over time by halos. Note how halos merge, get created, or disappear through time. With 3D glasses, analysts (7-person team, not all shown here) can see the depth of the halos and the paths they form, as well as where halo formations start and eventually end.

IMMERSIVE ANALYTICS CHALLENGES AND DIRECTIONS FOR FUTURE RESEARCH

In addition to the case studies reported here and elsewhere, we have also used these environments to judge our University's Image of Research contest for almost ten years, have held meetings related to evaluating different user interfaces and reviewing data for an information system for nurses, and have held graduate and undergraduate classes in this space. The CAVE2 form factor is uniformly appreciated for being able to represent all items, and for showing all the items at full resolution and, where appropriate, also at correct size (for example, winning images of research are printed on banners, and thus the ability to view the submissions at banner size is important). The spaces are routinely commended for leading to better discussions and to better decisions compared to single device or user spaces.

We note that while the recent environments described here address the high-resolution Immersive Analytics problem, they do not solve the second, “traveling to another room” problem. However, these environments take concrete steps towards better integrating that other room with the visualization and analysis workflow, making it less isolated from regular tools, and making it more appealing as a collaborative meeting space. From a historical perspective, in the times of the original CAVE systems, almost all analysts who were using computers for visualization and analysis needed also to use workstations — for their raw power and storage. In contrast, today many analysts use a laptop, which in conjunction with the analyst’s smartphone, and often a notebook, makes it easier for analysts to take “their office” along to these more powerful shared Immersive Analytics spaces. As demonstrated by the two case studies, many analysts consider this effort well spent.

Many of the lessons we have learned reinforce the lessons learned from the War Room research of the 1990s, as well as recent research in collaborative non-immersive spaces. For example, Meade et al. [11] also sought to bring a team of analysts together in a display-enhanced room for collaborative work, like our two case studies; although beyond that similarity, our case studies involve multi-disciplinary analysis work as a team, and require 3D and immersion for the sonar / bathymetry work, respectively for the halo spatial relationships. 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 Immersive Analytics. Some of these challenges were first summarized in the IEEE VR Workshop on Immersive Analytics [12].

Cost-Aware Immersion and Hardware Resolution

Our experience indicates that Immersive Analytics requires 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, in order to be attractive when compared to desktop displays. 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. Unfortunately, at the current time, there are few inexpensive solutions that are bright, high-resolution, 3D, and ideally touch-enabled. Investigators who do not have easy access to CAVE2 systems or immersive wall displays need to prioritize based on their needs and usage patterns and may have to rely on several different displays with different uses in the same room.

Shared Spaces for Flexible Group Sizes

While original CAVE projects featured one to a few domain science analysts, the trend over the years has been towards projects for larger and larger groups, many featuring complementary expertise. Due to their cost and relative rarity, high-resolution Immersive Analytics spaces are increasingly being used by multiple groups under either a time-share or a space-share model. Under space-share models, modern spaces should accommodate seamlessly large group meetings (30-50 people) and multiple smaller group meetings (5-10 people). 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. Under time-share models, 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, and be able to copy information from the public displays back to their personal displays.

Interaction Support for Group Work

Our experience indicates that Immersive Analytics needs to further consider the informational needs of users. Some analyses are about presentation of science or of analysis results through carefully curated and reduced data to lay audiences. Other analyses are about interactive exploration of as much data and as many data representations as possible, for expert user consideration. Users 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. Our experience further shows that it is important to support multiple forms of interaction simultaneously, e.g. embodied navigation, touch at a wall, controller or tablet-based interaction in the space, and laptop/mouse interaction when seated in front of your laptop.

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. Our experience further points to the benefits of embodied navigation in the context of large spatial datasets.

Quiet, Comfortable, Integrated Environments

Our experience is furthermore that Immersive Analytics requires quiet and cool rooms with enough light to see laptop keyboards or read paper notes. Our experience confirms that integration with the analysts' typical web-based desktop flow is essential. There should be sufficient power available for laptops. We note that analysts furthermore want to be able to bring in their lunch or beverage while working. It is 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.

CONCLUSION

We summarized in this work our experience, over more than 25 years, of designing collaborative environments for Immersive Analytics. We believe that resolution and immersion are possible simultaneously in spaces that use large displays. We further believe support for integrated analytics workflows is essential, as well as support for collaborative work. As shown in several lessons learned from recent successful Immersive Analytics case studies performed on scientific data in our hybrid immersive environments, it is further important that these environments are quiet and comfortable. Our experience demonstrates 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.

ACKNOWLEDGMENTS

We gratefully acknowledge our collaborators across disciplines, who made this work possible. We further thank A. Forbes and our students who worked with us in these projects, in particular K.Redda, P.Hanula, K.Piekutowski, C.Uribe, A.Nishimoto, and K.Almryde. This work was partially supported by federal grants NSF CAREER IIS-1541277, CNS-1625941, and by a small Computing Research Association CREU grant. We gratefully acknowledge the Electronic Visualization Laboratory faculty, students, staff, collaborators and visitors for their help.

REFERENCES

1. 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.
2. Reda, K., Johnson, A., Papka, M., Leigh, J., Modeling and evaluating user behavior in exploratory visual analysis In *Information Visualization*, Vol. 15.4, pages 325-339.
3. T. Munzner. *Visualization analysis and design*. AK Peters/CRC Press, 2014.
4. 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

5. 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.
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. J. Leigh and M. D. Brown. Cyber-commons: Merging real and virtual worlds. *Communications of the ACM*, 51(1):82–85, Jan. 2008.
8. 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.
9. 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 *Scientific Visualization Conference (SciVis)*, 2015 IEEE Oct 25, pages 87–99. IEEE.
10. 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.
11. B. Meade, C. Fluke, J. Cooke, I. Andreoni, T. Pritchard, C. Curtin, S.R. Bernard, A. Asher, K.J. Mack, M.T. Murphy, D. Vohl. *Collaborative Workspaces to Accelerate Discovery*. Publications of the Astronomical Society of Australia. 2017;34.
12. G.E. Marai, A.G. Forbes, A. Johnson. Interdisciplinary Immersive Analytics at the electronic visualization laboratory: Lessons learned and upcoming challenges. In *Immersive Analytics (IA)*, 2016 Workshop on 2016 Mar 20, pages 54–59. IEEE, 2016.

ABOUT THE AUTHORS

G.E. Marai is an Associate Professor in the Electronic Visualization Laboratory at the University of Illinois at Chicago and an IEEE senior member. Marai holds multiple research awards from NSF and NIH, including an NSF CAREER and NIH R01 awards, and multiple outstanding research awards at conferences. Her research is in visual computing, an area of computer science that handles computing with images and 3D models, as well as the processes that take place at the interface between humans and data that can be represented visually.

J. Leigh is a Professor of Computer Science and the Director of the Laboratory for Advanced Visualization and Applications at University of Hawaii at Manoa, and the former Director of Research at the Electronic Visualization Laboratory. Leigh holds multiple research awards from NSF, and multiple outstanding research awards at conferences. In addition to his research expertise in large-scale data visualization and virtual reality, his work encompasses high performance networking, human augmentics and video game design.

A. Johnson is Director of Research at the Electronic Visualization Laboratory and Associate Professor of Computer Science at the University of Illinois at Chicago. Johnson holds multiple research awards from NSF, and multiple outstanding research awards at conferences. His research focuses on the development and effective use of advanced visualization displays, including virtual reality, high-resolution walls and tables, for scientific discovery and in formal and informal education.