Teaching Data Visualization in evl's Cyber-Commons Classroom

Andrew Johnson*

Electronic Visualization Laboratory, Department of Computer Science, University of Illinois at Chicago, Chicago, IL 60607

ABSTRACT

CS 424 – Visualization and Visual Analytics I at the University of Illinois at Chicago (UIC) is a 15-week course taken by Computer Science Juniors, Seniors, MS, and Ph.D. students. It is commonly taught in the Cyber-Commons Classroom of the Electronic Visualization Laboratory (evl) which has a large 21' wide touch screen wall for teaching and class projects. Using this wall in ways similar to the multi-blackboard classrooms of the past encourages the students to compare and contrast multiple representations both in class and in their projects. Presentations of all of the projects in class allows the students to see multiple interactive solutions to the same data visualization problem and to compare their relative benefits and shortcomings.

Keywords: Data visualization, Education, Group work, Large displays.

Index Terms: • Human-centered computing~Information visualization • Human-centered computing~Displays and imagers • Applied computing~Education

1 INTRODUCTION

This paper will describe and discuss my experiences over eight years teaching CS 424 – Visualization and Visual Analytics I in the Cyber-Commons [1] classroom at the Electronic Visualization Laboratory (evl) at the University of Illinois at Chicago.

2 THE COURSE

CS 424 – Visualization and Visual Analytics I focuses mostly on information visualization, 2D geospatial visualization, data preparation, and data transforms. The class includes lectures and group work in class, one individual project followed by several group projects, and presentations of all of the projects in the class. We have a follow-on graduate course, CS 524, that focuses more on 3D volumetric scientific visualization and medical visualization.

This 15-week semester long elective course is taken by both graduate students (MS and Ph.D.) and undergraduate students (Juniors or Seniors with experience in Data Structures), mostly from the Department of Computer Science. 25-35 students take the class each Fall, and enrollment ranges from 1/3 undergraduate to 2/3 undergraduate. This course is usually the students' first exposure to data visualization, though they are familiar with modern programming languages and data structures, and some are familiar with user interface design. There are no exams, but the students are expected to work on 3 or 4 data visualization projects – the first on their own, and the others in small groups.

* ajohnson@uic.edu

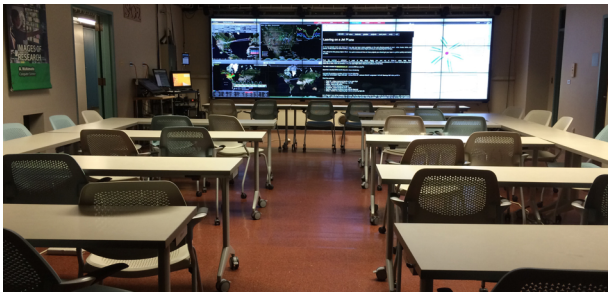

Figure 1: Photo of the Cyber-Commons classroom showing a typical classroom configuration of the display wall with examples on the left, a web page of notes in the center, and a freehand drawing area on the right. The tables and seating are designed to work for both lecture/discussion and group work, and can be moved as needed.

3 THE ROOM

I have taught this course seven times since 2009 when I created it at UIC. Each time I have taught the course I have taught it in the Electronic Visualization Laboratory's Cyber-Commons classroom. See Figure 1. This classroom was designed to be a place where we could apply lessons learned from our hardware research (CAVE, large tiled displays, and CAVE2 [2]) and software research (SAGE and SAGE2 [3]) to our regular teaching activities. We began moving various combinations of projectors and tiled displays into this room in 2004 for classes, and the room reached its current configuration in 2009.

The room is 24' wide x 42' long (7.3m x 12.8m) and seats 40 people at tables, which can be easily moved into different configurations, and 10 more in a row of high chairs at the back. The front of the room is dominated by a 21' wide x 6' tall (6.4m x 1.8m) tiled 17-megapixel LCD screen with a touch screen overlay. We currently can either drive the wall as a single large Windows 10 desktop, or more commonly as a shared SAGE2 display.

SAGE2 is a web browser-based software framework allowing multiple users to simultaneously show, share, and interact with multiple artifacts on large high-resolution displays. In general, SAGE2 acts like the operating system for a digital war room or project room. In the case of the Cyber-Commons classroom SAGE2 gives me a large high-resolution canvas upon which I can stream multiple desktops, and place additional images, movies, pdfs, etc. The instructor can interact with the wall through touch at the wall, or through a laptop computer running a web browser. The students in the class can also connect to the wall through their web browsers to interact.

Using this kind of large high-resolution display real estate has several advantages, and is similar in many ways to having a classroom with multiple rolling blackboards or whiteboards, giving more of a sense of context to the material, and leaving enough room to make comparisons and linkages without having to flip back and forth through a PowerPoint deck or scroll back and forth through a web-browser or film roll on an overhead projector. During project

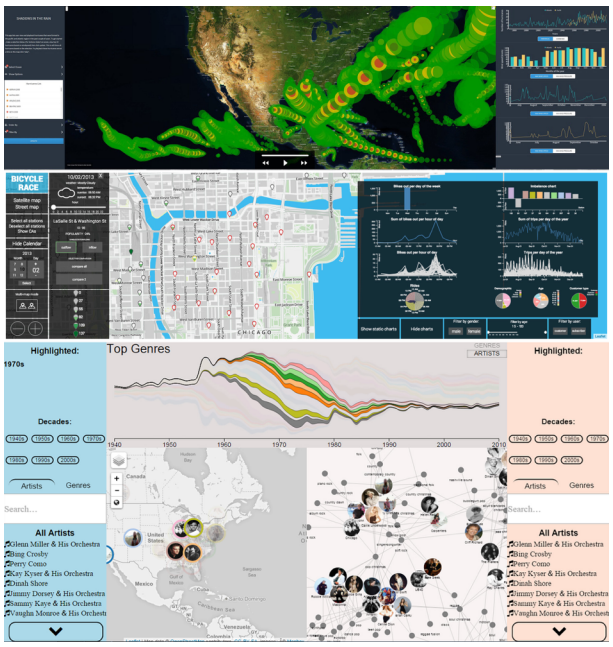

Figure 2: Sample student CS 424 Projects that were written for the classroom wall in D3, simultaneously showing the data in multiple ways – from top to bottom: Hurricanes near the United States, Divvy bike-sharing data in Chicago, and musical artist data.

presentation days the room can function like a grammar school hallway where the teacher posts the solutions of all the students on the wall at the same time so each student can see the multiplicity of solutions and reflect on their solution compared to those of others.

The room is in regular use. Since 2009, in addition to CS 424, I have taught four other courses there, and six other faculty members in Computer Science, Art and Design, Communication, and Physics have taught 19 times there. The room is also used for lab-wide meetings, presentations from external visitors, and very often students can be found there after-hours testing their projects, studying for exams, or practicing presentations for various courses.

4 Class Time and Class Projects

Class time is usually broken up roughly 50/50 between lecture/discussion and small group work on related exercises. The seating in the room is reconfigurable and designed to make it easy for people to work in small groups. During lecture/discussion I typically split the wall into three parts. The center of the wall shows a webpage where I have my general notes and examples for that class. All of these notes are available to the class on the web at the start of the term if they want to follow along on their laptop or tablet. I typically use the left side of the wall as a repository of related visualizations, diagrams, and videos that I can drag out of my notes. The right side of the wall is usually a virtual white board for freehand sketching or writing. See Figure 1. The focus of the class tends to be on the 1/3 of the wall where the action is currently taking place, with related material on the other 2/3. Since I adopted this layout I have been getting more questions and comments relating recent material (that is still visible on the wall) to the current material, and more discussion across more examples.

The students create multiple visualization projects during the course. The first is an individual project, followed by two or more group projects. Typically, the first project deals with standard graphs, then the group projects add on geographic data, network data, dynamic data, etc. All of the students in the class work on the same project at the same time, rather than having individual groups work on individual projects, as I really want the students to be able to compare and contrast their solution with the solutions of others.

Figure 3: Interdisciplinary team of scientists working with multiple visualizations in the Cyber-Commons as a research space – this kind of interaction with multiple linked representations was the model for the projects that the students work on in the class.

In the first four years of CS 424 I used Processing [4] as the language students would use to implement their projects, and then tried KineticJS for one year, but have been using D3 [5] since 2014. I also moved from initially having five projects down to four and now down to three to allow more time for mid-project reviews. In all of these cases the students developed interactive visualizations for the larger display wall, where the projects simultaneously show data in multiple ways, and user interaction modifies all of the linked views. They are given links to the dataset, and a set of requirements for a C, B, and A grade. Typically, the C level uses a subset of the data for overview graphs and a single interactive visualization controlled by a couple of filters; the B level uses the entire dataset, multiple linked visualizations and more options for the filters; and the A level focuses on integrating related external datasets, finding and documenting interesting patterns and trends.

For example, when looking at NOAA hurricane data in the North Atlantic and Eastern North Pacific there is room to show maps of the hurricane paths in both oceans at the same time, along with traditional graphs showing the number of storms per year, ratios of intensities, as well as controls for which hurricanes are shown, and what range of dates to show in a static view or played back as an animation. In the Divvy bike-sharing data the visualizations show a map of the city with the various starting and ending locations of trips, graphs breaking down the data by day of the week, hour of the day, rider demographics, and controls to filter the data by location, time, or demographic data. The user has multiple entry points to the data, and might be interested in how the seasons or a particular day's weather affects ridership, where do people bike from before a Cubs game, are the people renting by museum campus casual users or subscribers, or what are the stops with the biggest imbalances of incoming and outgoing bikes at different times of the day.

Writing applications for this very wide high-resolution wall is pretty easy with modern scalable graphics libraries, as the students can do much of their work on their laptops running a standard web browser, and then come into the classroom to test on a web browser running on the larger wall to make sure the elements function at an appropriate human scale in terms of readability and accessibility.

While the students are constantly working with personal devices such as their laptops, phones, and tablets, I want them to also think about co-located multi-user group experiences. The large wall encourages the students to write applications that support multiple simultaneous users, allowing the users to see the data in multiple representations, and seeing different paths or stories in the data simultaneously, whether created by one or more users – e.g. the

bottom image in Figure 2 shows a musical artist and genre visualization where two users can set independent filters and see information about all those artists and which they have in common.

These usage patterns that we want to support come from our experiences working with teams of interdisciplinary scientists who need multiple visualizations to be able to work effectively, as their disciplines look at the same data in different ways. These scientists also want enough room to show enough representations to force themselves to confront parts of the data that do not match their current explanations, as they do not converge to an incorrect solution too quickly. See Figure 3.

I also want the students to try creating interactive visualizations for platforms that are not commonplace. The technology they encounter in school is not going to be the technology they use after they graduate, so I want them to think about how to work on a novel platform, like a large wall. In past User Interface Design Courses I used handhelds in the 1990s and tabletop displays in the 2000s for the same reason. Since the code they are writing is scalable and runs in a web browser, it is pretty straightforward to take these same projects they write for the large wall in class and break them up into smaller views that can run on typical laptops or tablets.

After a month of the students working on their first individual project, I usually devote two weeks in class for the students to demonstrate their solutions, talk about the rationale behind their decisions, and receive feedback from the class. Aside from the benefits of seeing and hearing about a variety of solutions to the same problem, this also helps the students gain more experience in presenting their work effectively to others and answering questions about that work, and they get feedback from the class on those presentations. These presentations also help the students to form compatible teams for the later projects. Each of the later projects has one week where all of the groups present their projects to the rest of the class and answer questions. Overall, while there is a range in quality of the solutions, many of them are very well done, both in terms of their functionality and their user interface. After all the groups present their work, we spend some time showing a representative image of all of the projects at the same time and talking about the common and less common choices that the various groups made. Finally, there is one more short presentation near the end of class where each student presents a recent interactive visualization that they found particularly interesting, to help show the variety of interactive visualizations out there.

In addition to the presentations, the students must create a public web page documenting each project they work on, including discussing how they manipulated the data, instructions on how to use the application, and interesting features they found in the data. This page must include a link to a short narrated YouTube video demonstrating the use of their project. These web pages and videos help with grading, as we know what the application is supposed to do, act as an archive of past projects for new students to look at, and give the students something to add to their portfolio when they go looking for a job.

Typically, the students break up the project work with one person focusing on the data processing, and the others on creating the different interactive visualizations. They do much of the work individually, along with some separate integration testing, before starting to evaluate the overall application in the classroom. Sitting in on some of these sessions at the wall I could see that part of the work was basic functionality testing and scaling elements for readability and usability, but as the students brought up different examples from the data, the discussion often turned to what they were seeing in the visualizations. Having multiple eyes on the multiple visualizations at the wall seemed to stimulate more conversation about what they were seeing, and whether certain features in the visualizations were really in the data, whether the data made sense, how to best compare things at different scales, dealing with missing data, comparing different ways to represent the same data, etc. These discussions tended to be most interesting in the groups with students from more diverse backgrounds.

Figure 4: Teaching a smaller seminar class in CAVE2 where everyone has equal simultaneous access to the larger surrounding screen space through the web browser on their laptop.

The large wall is also very handy for looking at code, as a lot of code can be shown readably on the screen at the same time, allowing the students to see the code in context, especially when related code can be spread across several files, and the code can be shown simultaneously with its results. When debugging, as we learned from our classic CAVE days, the big walls make it easier to have more eyes looking at the problem, and often lead to faster solutions.

Other departments on campus are building similar setups. We helped our Pathology Department built a room very similar to evl's Cyber-Commons on our medical campus, as they also have need to show and interact with multiple high-resolution images in their classes. Our Communication Department has also deployed a smaller seminar room display made from three large monitors.

5 Student Reaction

Student reaction has been generally positive, though course ratings from the graduate students have been higher than those from the undergraduates. The graduate students are more likely to embrace the possibilities of more exotic technology like a large wall, and tend to look for jobs that employ more futuristic technologies, while the undergraduates prefer more common platforms with more guaranteed employment. As expected, the graduate students have more experience working in teams, start earlier on their projects, test more often on the actual hardware, and spend more time revising and documenting their work.

The class has been a great place to find and recruit graduate and undergraduate research assistants to work in evl, as we can see who is good at creating visualizations for novel platforms, working in teams, getting their work done on time, and effectively presenting their work to others. One unexpected outcome from the course was the number of students taking their D3 data visualization experience from the course to various visualization hackathons across the country, and doing rather well at them.

6 Seminar Courses

In CS 524, the graduate level visualization course, as well as other human-centered computing graduate courses at UIC, which are smaller seminar courses with 10-20 students, the Cyber-Commons wall has the additional benefit of allowing and encouraging the students in the class to contribute to the discussion by dragging and dropping relevant content on their laptop's local web browser to add it to the classroom wall. This makes discussion much more lively as content from multiple sources can be shared immediately with the entire class, and questions can be answered much more quickly.

I have also taught these seminar courses in the more intimate CAVE2 environment (a 22' (6.7m) diameter cylinder) with a much larger amount of screen space, which comes much closer to the idea of a digital war room or project room. See Figure 4. While focus and discussion tend to center around a certain set of documents (images, videos, papers, etc.) on the wall, the students can be adding relevant images to other parts of the wall, and then either move that new content into the area of discussion or ask the group to turn and look at this new material without disturbing the current material. This gives a larger canvas on which to add and position related material for discussion, and relevant material stays visible longer.

7 Variations on the Themes

There are many different possible variations on these themes. Rooms with whiteboard walls and a few good HD or 4K short throw projectors can give enough space to show multiple simultaneous interactive visualizations for instruction and discussion or project review. For more permanence and easier viewing outside of class, large printouts can be made and hung on the walls. For project review each team can set up their solution on a laptop or tablet in the classroom and one team member can host that solution while others rotate around to visit other groups at regular intervals. If the class has a large number of laptops or tablets, then small groups can bring up related visualizations on multiple devices to discuss their advantages and disadvantages.

8 Conclusions

A major impetus for creating the Cyber-Commons classroom, and to regularly hold classes in it, was to try and move beyond the limitations of single-projector multi-media classrooms, and to see how teaching and class projects in visualization, user interaction, and other human-centered computing areas would change given interactive screen real-estate similar to the multi-blackboard classrooms of the past. This extra real estate encourages comparison, and helps highlight the advantages and disadvantages of different representations and methods of interaction. Current libraries like D3 make it easy to scale projects from laptops to these larger displays, or to smaller handheld displays. Presenting various different solutions in class lets the students see different possibilities that they may not have considered, and to discuss those solutions with their authors. While classrooms like the Cyber-Commons may remain exotic in the near term, the focus on group projects, multiple representations of the same data, having the students present and defend their work in class, and comparing multiple solutions to the same problem, transcend the hardware and could be applied to a variety of environments.

Acknowledgements

This publication is based on work supported in part by the National Science Foundation (NSF) awards CNS-0959053, OCI-0943559, ACI-1441963; The Electronic Visualization Laboratory at the University of Illinois at Chicago also receives funding from Argonne National Laboratory, the State of Illinois, and Pacific Interface on behalf of NTT Network Innovation Laboratories in Japan. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies and companies.

SAGE, SAGE2, CAVE, and CAVE2 are trademark(s) of the University of Illinois Board of Trustees

References

[1] J. Leigh, M. D. Brown. Cyber-Commons: Merging Real and Virtual Worlds. *Communications of the ACM*, 51.1:82-85, January 2008.

[2] A. Febretti, A. Nishimoto, T. Thigpena, J. Tolandis, L. Long, J. Pirtle, T. Peterka, A. Verlo, M. Brown, D. Plepys, D. Sandin, L. Renambot, A. Johnson, J. Leigh, CAVE2: a hybrid reality environment for immersive simulation and information analysis, *Proceedings of SPIE 2013 - The Engineering Reality of Virtual Reality 2013*, San Francisco, CA, February 3-7, 2013.

[3] L. Renambot, T. Marrinan, J. Aurisano, A. Nishimoto, V. Mateevitsi, K. Bharadwaj, L. Long, A. Johnson, M. Brown, J. Leigh. SAGE2: a Collaboration Portal for Scalable Resolution Displays, *Future Generation Computer Systems* volume 54, January 2016, pp. 296-305.

[4] processing.org

[5] d3js.org