Editor: Theresa-Marie Rhyne

Artistic Collaboration in Designing VR Visualizations

Daniel F. Keefe, David B. Karelitz, Eileen L. Vote, and David H. Laidlaw

Brown University

This article describes some of the lessons we have learned from our collaborations with artists on visualization problems at Brown University's Visualization Research Lab. Over the past several years, we have worked closely with artists to develop, refine, and critique visualizations ranging from archaeological dig data from the Great Temple of Petra site in Jordan to the fluid dynamics and wing bone shape deformations that begin to explain how bats, the only flying mammals, stay aloft.

Perhaps, the most important conclusion we have drawn from this experience is that artists can fill an important role in the visualization design pipeline. In our experience, artists routinely provide a unique source of visual insight and creativity for tackling difficult visual problems. They are also expertly trained in critiquing and refining visual works, an essential task in the iterative visualization process.

The second major conclusion we have drawn from our collaborations with artists is that we need more appropriate design tools to support them and their role. We discuss here the experiences that led us to this conclusion, along with some of the tools we have developed to facilitate working with artists. The lack of appropriate design tools is particularly evident in VR. It's difficult for artists to get involved in visualization design for VR, since, with rare exceptions, you need to know how to program in order to create within the medium. VR is one of the most promising technologies for visualizing today's complex data sets.[1] However, it's also one of the technologies that can most benefit from artistic insight, since guidelines for good visual depiction are far less developed in the unconventional visual space of VR than in more traditional media.

We begin by describing one of our recent major collaborative efforts, a class on designing VR scientific visualizations that was cotaught with professors and students from Brown University's computer science department and the Rhode Island School of Design's (RISD's) illustration department. Many of the experiences and conclusions relayed here are the results of this class. We then discuss three important themes that we derived from our experiences, which are all motivated by a desire to better facilitate artistic collaborations.

Teaching art to computer scientists, computer science to artists, and fluid flow to everyone

Our interdisciplinary visualization class brought artists and computer scientists together to solve visualization problems driven by science. Students worked in teams on visualization and design assignments, following the interdisciplinary Renaissance team model presented by Donna Cox.[2] We began the semester with 2D fluid flow visualization assignments, as shown in Figure 1, and gradually built up to the final projects, which were VR visualizations of pulsatile blood flow through a branching coronary artery. We found more

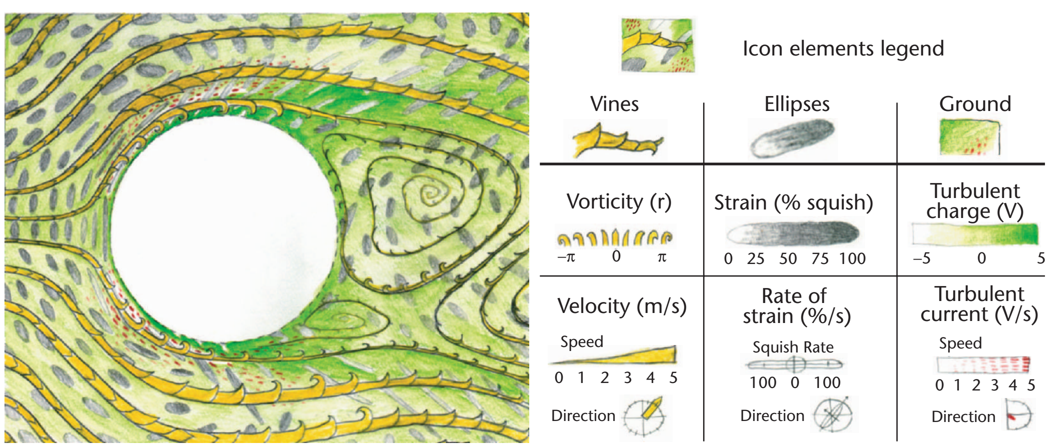

1 Art student's visualization design of 2D steady fluid flow past a cylinder. (Courtesy of Deborah Grossberg.)

Published by the IEEE Computer Society

obstacles to collaboration as we moved toward VR and more complex data. Despite these obstacles, the students learned to collaborate with each other and to value what each discipline (computer science and art) offered to the project, producing some interesting visualizations in the process.

Although artists rarely work with complex scientific data, they do train to convey information effectively through imagery, given the constraints imposed by their media, employers, or audience. In this abstract sense, normal artistic practice is not such a far cry from typical visualization design tasks. The images in Figure 1 show one art student's result from an early visualization design assignment. The students created visualizations and legends that convey eight continuous variables describing a steady, 2D fluid flow in a single picture. This is a difficult visual problem; in fact, the visualization community is still actively researching it. We found that artists were adept at investigating visual problems like this one when we could clearly convey the scientific goals and constraints of the problem.

2 **Students prepare for a critique of arterial blood flow visualization designs.**

Artists also excelled in the initial design and conceptualization stages of the scientific research process, often prompting new insights on the part of the scientist team members. Scientists rethought their hypotheses, clarified their experimental goals, and even altered the way they collect data in response to feedback from artists. As Vibeke Sorensen explains in her discussion of the artist's contribution to scientific visualization, this role for artists is a departure from the norm.[3] Artists are typically thought to be useful only in the last stage of the research process, dissemination. However, this is a limited use of the artist. Our experiences support Sorensen's claims that artists can be involved in many more stages of research, with conceptualization being perhaps the most important with regard to scientific visualization.

Collaboration was sometimes difficult to manage. In early assignments, such as in Figure 1, the right tools for the job were colored pencil, oil paint, gouache, watercolors, and software (Adobe Photoshop). In later assignments, the essential tool for the job moved closer to programming.

At this point, the art students often had visual insights to offer, but had difficulty conveying them. It was easy for the nonprogrammers to feel left out of the loop. As Fritz Drury (the RISD illustration professor who cotaught the class) remarked, the programmers are the ones with the ultimate power: They have the final say about what ends up on the screen.

One device that helped us keep artists, computer scientists, and fluid flow researchers on the same page is the critique, a common teaching tool in art classes. All the class work was displayed on a wall, as seen in Figure 2, and as a class, we discussed important design lessons in relation to each work. We critiqued the work both from a visual and a scientific standpoint. Visually, we explored color, scale, form, metaphor, and narrative. Scientifically, we learned about the data we were trying to represent and critiqued the work on the basis of how truthfully and completely the science was represented, given the tasks our scientists wished to perform. We have now adopted *crits* into the visualization development process for many of our projects.

How can artists approach design problems in VR?

As we move from 2D visualizations into more complex 3D situations such as VR, collaboration with artists becomes much more difficult to facilitate. The first theme we derived from our class experiences (along with other collaborative efforts) is that visualization design should occur within the visualization target medium.

This sounds simple, but it has fairly significant ramifications for the visualization media we often use. For example, it's difficult for anyone, and nearly impossible for an artist who is not a programmer, to create visualizations or simply experiment with design ideas in VR.

A starting approximation for designing within VR is to use more traditional, often 2D, media and hope that some of these design ideas will translate to VR. We were forced to take this approach during many of the class assignments. The difficulty is the drastic difference between what we can convey on paper and what we can convey in VR. We use a four-wall CAVE VR display environment for much of our research. So much changes when we enter the CAVE: scale, interaction, stereo vision, vividness of color, and contrast. Due to the drastic differences in the mediums, it's difficult to trust or evaluate traditional designs of VR ideas.

3 **CavePainting visualization design of bat flight data shown using snapshots from a 3D VR program. The bat appears to fly into the page in these snapshots, but viewers walk around the entire model when seen in VR.**

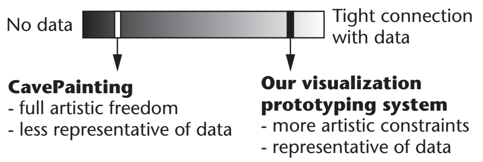

4 **Design tools can have a stronger or weaker built-in connection with data. Tools at both ends of the spectrum are useful.**

Ben Shedd notes a similar dilemma in a comparison between traditional filmmaking and new Imax-style filmmaking.[4] In giant screen films, as in VR, the projected images extend beyond our peripheral vision. This significant change has required filmmakers to begin to invent a new visual language, and prompted Shedd to call for redefining giant screen filmmaking tools. This is one of the issues Shedd explores in his interdisciplinary class at Princeton.

In VR visualization, we are also defining a new visual language, and we deal with a similar lack of appropriate traditional tools to do the job. When designing traditionally with an eye toward VR, we face the problem that a good 2D design does not necessarily translate into a good 3D, much less VR, design. Further, it's difficult to evaluate or propose refinements to a design without actually seeing it implemented in the CAVE. We lose the power of the critique, which we have found so useful. We need to be able to design and critique within VR.

With this motivation in mind, we began exploring ways to work with artists to design visualizations directly within the CAVE. Figure 3 shows some snapshots of one of our VR design results. In this project we are collaborating with Sharon Swartz, of Brown's evolutionary biology department, who studies bat flight from experimental data collected in wind tunnels. Two important clues to understanding bat locomotion are the air flow information surrounding the wing and the pattern of deformations of the wing bones during flight. Artists worked directly in VR to create the visualization design shown in Figure 3.

Since the bat data assumes symmetry between the two wings, the artists chose to represent different aspects of the data on each side of the bat. On the left side of these images, flow close to the wing is described by color and texture along the wing surface. Vortex cores and vortical structures in the flow behind the bat are also represented. Changes in bone shape at two distinct times during a wing beat cycle are shown on the right side of the images along with a 3D trace of an important bone joint through the wing beat cycle.

The basis for our VR design tools is the CavePainting program,[5] a tool intended for artists to use inside the CAVE environment to create free-form 3D objects. Artists have described it as a form of zero-gravity sculpture. Artists interact with the system by moving a tracked paintbrush prop through the air to create 3D paint strokes. The paintings are actually 3D models, since each brush stroke exists in 3-space. The system's intuitive interface makes it easy for artists to pick up and quickly begin modeling in the CAVE.

Working directly in the CAVE with a tool like CavePainting has several benefits. The most important is that the design can be easily critiqued and refined with proper attention to the nuances of the target medium. In practice, we have gained valuable insight from these critiques.

We have made several alterations to our initial bat visualization designs based on feedback from Swartz and her collaborators after meeting for critiques in the CAVE. During these critiques we have even been able to quickly sketch modifications to designs and discuss them immediately.

Using CavePainting to design visualizations also lets us investigate, refine, and converge on a successful visual design at an early stage in the process. With the usual approach of implementing before visual refinement, it might take weeks or months of implementation before we discover our design is flawed from a visual standpoint, and once we notice a problem and brainstorm another design, it could take another few weeks before we are ready to visually critique that new design.

Thus, particularly in VR, where implementation can be difficult and time consuming, putting visual design decisions in series with implementation can extend the time between iterations on a visual design. Designing directly in VR, on the other hand, lets us converge upon a visually successful design early in the implementation process. We can quickly work through many more

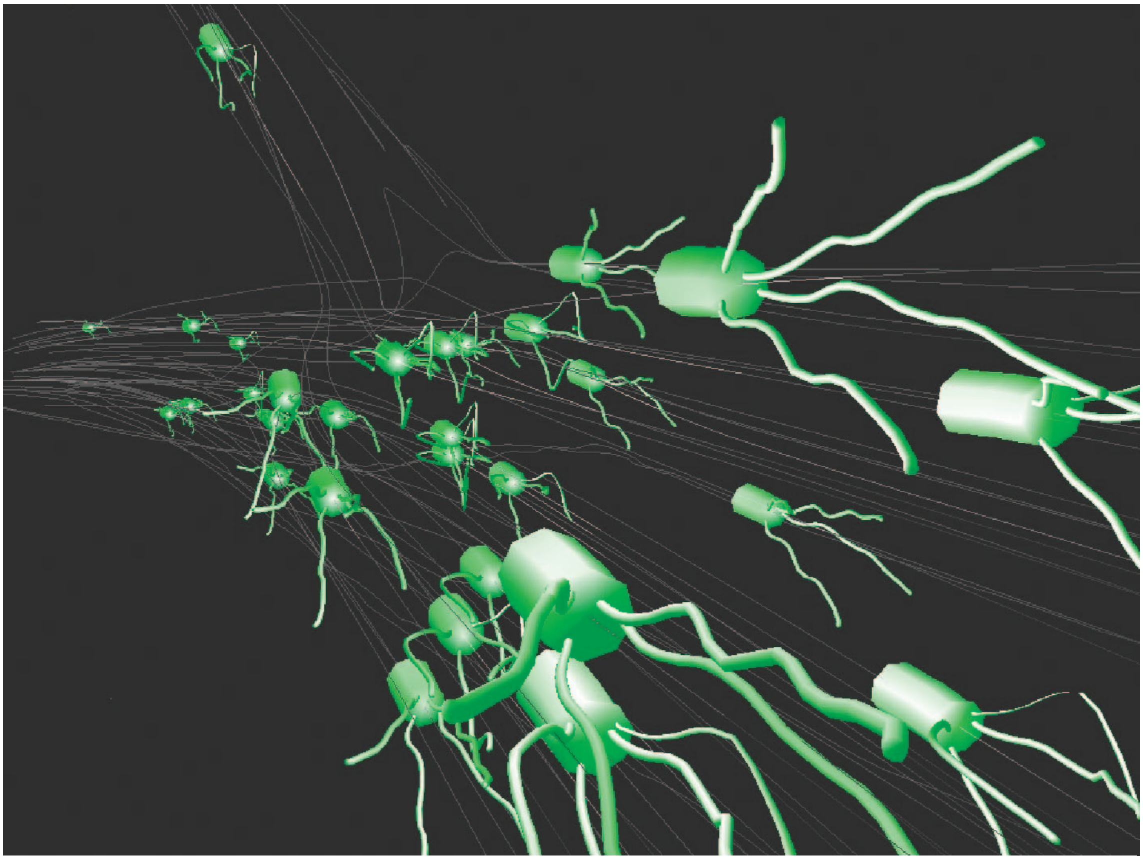

5 **Result from our visualization prototyping system. An artist sketched the 3D icons, then connected them to an arterial blood flow data set so that they morph in direct response to the data.**

iterations of the design because we do not have to wait for them to be implemented before critiquing them in the CAVE.

Where's the data?

The second theme that has emerged from our collaborative work is that we should incorporate varying levels of data involvement in the design process. In the bat visualization design shown in Figure 3, there is no programmed link between the visuals and the bat flight data. Designs such as this lie at one end of the spectrum shown in Figure 4. Despite the lack of a low-level link with the data, this type of design is extremely useful. The designers have imagined some representative data and sketched it out. The visualization is not far fetched; the designers have seen previous attempts at bat visualizations and talked with the scientist about her goals. Essentially, they know enough about the data structure to paint a typical situation so that they can meet with the scientist and critique the visualization idea in the CAVE.

The danger in going too far in the design process without a program-level connection of the visuals to the data is that we might converge upon a design that works well for our perception of the data but not as well with the actual data. In an effort to explore this issue, we built some design tools to explore the other end of the spectrum of data involvement. As we see in Figure 4, tools like CavePainting lie at the far left end of the spectrum, with no program-level connection to the scientific data. The visualization prototyping system we describe next is much closer to the right side of the spectrum, where data plays a key role in generating the visualization design.

Our visualization prototyping system lets an artist draw icon-based 3D visualizations that are completely driven by the underlying scientific data.[6] Figure 5 is a snapshot of one such visualization design. The squid-like icons represent data values within a fluid-flow data set of pulsatile blood flow through a branching coronary artery.

In this design, the squid's tentacles morph in response to data values. At high speeds, they straighten out and the squid appears quite streamlined. At lower speeds, they flail out to the sides, as the squid assumes a sluggish posture.

This tool has been useful in evaluating several different designs for arterial blood flow visualization. Since we are working with time varying, pulsatile fluid flow, the ability to see the design animated—with icons flowing down the artery and changing shape in response to the data—is critical in evaluating the design's success. This would be a difficult display to realize without a program-level link to the flow data. Despite the success of this approach in achieving these animated visualization designs, we have had difficultly moving beyond these relatively simple cases to the more complex ones required for many of our driving scientific problems.

This experience illustrates the tradeoff that exists in many design systems based on the role they provide for data. Given plenty of preprogrammed connections to

6 3D visualization design for the bat flight problem inspired by the Miró painting The *Gold of the Azure* (see http://www.bcn.fjmiro.es/angles/_coleperm/_salaanys6070/lordelatzur.html).

data, design tools can produce visualization designs that are so representative of the data that they can be trusted and critiqued as completely accurate visualizations. However, preprogrammed connections to data can be constraining to the artist. For example, in our current implementation of the prototyping tool, the icons must be drawn in a special way to establish a solid correspondence for our morphing algorithm. This means that the artist must have this in mind while working on the design. Creating complex designs—for example, icons that respond to six different variables—can become almost impossible to manage cognitively. Again, these difficult design tasks are the ones our driving scientific problems require and the ones in which we can most benefit from artistic insight. We need to continue to develop intuitive design tools that provide this type of solid connection to the data, but also allow artists to work naturally.

Tools for VR design work

The final theme that has emerged in our collaborations is the need to support continued, evolving work with VR tools. This has been evident in two areas. First, getting started in VR is difficult. Often our artists make several preparatory sketches or studies before entering the CAVE to work. We need to make it easier for them to build on those sketches when they get to the CAVE, rather than shutting out the real world and concentrating only on VR. Second, we need to facilitate returning to a design to rework and refine it. The real-world problems with which we anticipate artists will work are sufficiently complex that they will require many design iterations to complete. Tools to facilitate artistic collaboration in visualization must be accessible to artists in these ways if we want to support artistic involvement in difficult visual problems.

Looking at CavePainting again as an example of an artistic design tool, we can see that it can be difficult to start on a VR design. When the program begins, we walk into the CAVE—a dark, blank room of projection screen walls. We carry a tracked paintbrush prop and a pair of glasses. Once we put on our stereo glasses, it's too dark to see any paper or other real objects we might have brought in with us. By default, we start with a completely blank canvas and no external inspiration, something designers almost never want to do.

One approach that has been helpful to us is to import our design inspiration, often 2D work such as sketches or paintings, directly into the VR design program. In Figure 6, we see a 3D CavePainting design inspired by a Miró painting.[7] One of our designers saw the painting, and it gave her the idea for visualizing the bat data set. We cut out subregions of the Miró and imported them into CavePainting as brushstroke textures. Then our designer could work directly with elements of the inspirational imagery to create the 3D design she imagined. This gave her a jump start on her 3D design and helped her quickly create a coherent design.

The ability to return to a design and refine it again and again is just as important as starting with something in VR. The design task is necessarily an iterative one, with critiques by other designers, implementers, and scientists all playing an important role in refining each iteration.

Normally artists refine work in two ways. First, they add additional layers of clarification. In painting, for example, additional layers of paint conceal what lies below. A rough outline of a face can be laid down as a place holder for a much more complex rendering to come later, applied with additional paint layers. Second, they create many studies of an idea, sometimes ending up with a studio full of renderings and re-renderings. At the end of this period the idea is clear enough in the artist's mind that she feels ready to produce a final work.

These approaches are not at all mutually exclusive, however, we have difficulty supporting either with our current design tools. In the first case, we can add some additional layers of clarification with the CavePainting system, but this can have the effect of distorting the original form. We are a little closer to supporting the second style of refinement, which amounts to letting an artist quickly reel off many sketches before creating a final work. However, it's unclear how to refer back to several studies while working on a new piece, since each design is usually intended to be viewed in the full space of the CAVE. These issues are among the most important to address before working closely with artists on design problems, since they can be frustrating and limit the amount of real design work that can be accomplished.

Conclusions

We have had many exciting and fruitful collaborations with artists, and we are convinced that they have a place in the visualization design pipeline. One of the driving motivations in our recent work has been to consider what an artist would do for 8 hours a day if hired by a visualization lab. Given current visualization practice, this is a difficult question to answer. However, the key seems to lie in enabling an artist to get involved in design at a level that goes deeper than simply turning knobs of existing visualization techniques. We anticipate that artists will be hired to fill positions in exploratory visualization. That is, rather than merely making a picture pretty or visually clear for publication, we see artists as having key roles in working closely with scientists to design novel visual techniques for exploring data and testing hypotheses.

We hope we have illustrated the potential for this type of collaborative visualization work, along with presenting some of the lessons we have learned along the way in our collaborations with artists. We also hope to have further motivated the need for additional research in design tools that can be easily targeted toward visualization problems. ■

Acknowledgments

We thank the Brown and RISD students from our visualization design class, as well as RISD illustration professor Fritz Drury. We also thank our collaborators Sharon Swartz and Peter Richardson. This work uses the Virtual Reality Peripheral Network library, which is supported by the National Institute of Health's National Research Resource in Molecular Graphics and Microscopy at the University of North Carolina at Chapel Hill. This work was partially supported by the National Science Foundation (CCR-0086065).

References

1. A. van Dam et al., "Immersive Virtual Reality for Scientific Visualization: A Progress Report," *IEEE Computer Graphics and Applications*, vol. 20, no. 6, 2000, pp. 26-52.
2. D. Cox, "Using the Supercomputer to Visualize Higher Dimensions: An Artist's Contribution to Scientific Visualization," *Leonardo*, vol. 21, no. 3, 1988, pp. 233-242.
3. V. Sorensen, *The Contribution of the Artist to Scientific Visualization*, California Institute of the Arts, 1989; http://visualmusic.org/text/scivi1.html.
4. B. Shedd, "Exploding the Frame: Seeking a New Cinematic Language," *Proc. Society of Motion Picture and Television Engineers 135th Conf.*, 1994; http://members.aol.com/sheddprod2/ExplodingtheFrame.html.
5. D. Keefe et al., "CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience," *Proc. ACM Symp. Interactive 3D Graphics*, ACM Press, 2001, pp. 85-93.
6. D. Karelitz, D. Keefe, and D. Laidlaw, "Using CavePainting to Create Scientific Visualizations," *Proc. IEEE Visualization*, IEEE CS Press, 2003, p. 109.
7. J. Dupin, *Miró*, Abrams, 1993.

Readers may contact Daniel F. Keefe at dfk@cs.brown.edu.

Readers may contact editor Theresa-Marie Rhyne at tmrhyne@ncsu.edu.