EDITORS: Dave Kasik, dave.kasik@gmail.com
Chris Johnson, crj@sci.utah.edu
Mary C. Whitton, mcwhitton@gmail.com

DEPARTMENT: VISUAL COMPUTING: ORIGINS

Adventures of a Government Researcher

Lawrence J. Rosenblum

Visualization, interactive computer graphics, and related topics have been a particularly dynamic area of computer science, producing advances that impact society. Working at times in a research laboratory and at times for two science funding agencies, I held positions at the level of an individual researcher, Director of a VR Laboratory, and funding agency Program Director. This article will discuss some of my experiences, including insight into how science programs in funding agencies are initiated.

SCIENTIFIC VISUALIZATION

I had planned on a career in academia after receiving my Ph.D. in mathematics in 1971 from The Ohio State University specializing in number theory. However, there were few U.S. academic positions for new graduates due to Vietnam war era funding cuts. Those with Ph.D. degrees in pure mathematics were especially hard hit. Historically, almost all would take university positions and there was limited industrial demand for their expertise. Friends were taking positions at branch campuses or community colleges. I ended up receiving a reasonable academic job offer, but by then I was examining nonacademic alternatives. I accepted a government position where I could learn programming and other computer science skills.

Several years later I found myself at the Naval Research Laboratory (NRL), designing real-time data collection systems for the ocean sciences and analyzing collected data. NRL came into existence in 1923 from an idea of Thomas Edison who wrote: "The Government should maintain a great research laboratory-...In this could be developed...all the technique of military and naval progression without any vast expense." NRL produced the first practical radar systems in the U.S. and grew into a large, interdisciplinary research laboratory whose staff included a Nobel Prize winner.

0272-1716 © 2021 IEEE
Digital Object Identifier 10.1109/MCG.2021.3115059
Date of current version 10 December 2021.

My adventure began in the early 1980s when, more out of intellectual curiosity than anticipation of the future, I took a computer graphics graduate-level course taught by Jim Foley at George Washington University and televised into NRL. The ocean and I were not getting along (seasickness). It was time to move my career in a different direction. Developments in graphics hardware and processing power were enabling visualizations not previously possible, although at a significant cost in time and money. The visualizations were leading to scientific advances. Several ocean science research topics seemed likely to benefit.

My largest project involvement, covering multiple oceanographic experiments, was working with oceanographer George Marmorino to examine ocean finestructure. Temperature measurements made with a towed thermistor chain in the seasonal thermocline of the Sargasso Sea were converted to sections of finestructure "activity" by normalizing values of meter-scale horizontal fluctuations by the local vertical gradient. A longstanding conjecture was that ocean frontal regions contained higher amounts of finestructure. Previously, such data had been examined by converting each slice into a numeric value approximating the Richardson number (a measure of turbulence often used to investigate density and turbidity currents in oceans). By instead representing the data as 2-D for each measurement slice, we obtained pictures that showed that larger patches of finestructure began to appear as the front was approached and they followed the temperature gradient through the front. The patches then gradually diminished on exiting the front.[1]

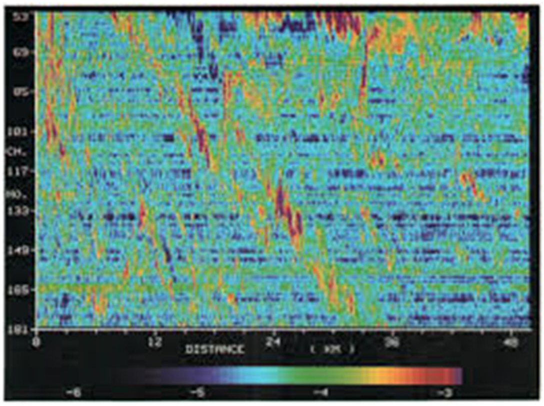

FIGURE 1. Color display of finestucture data.

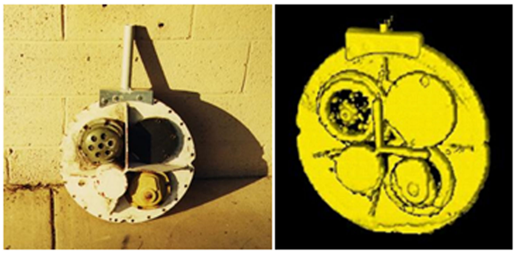

FIGURE 2. Object containing small features (left) and reconstructed object from acoustic lens data (right).

Subsequent research using newly obtained experimental data identified shear instabilities as the cause of the enhanced activity and statistically categorized the finestructure. A few years later one would just output the image onto a color terminal. However, to visualize our data circa 1983 (see Figure 1), I had to purchase and program a Lexidata hardware display system. The Lexidata offered a 256-color display and provided some data manipulation capability at a cost of roughly US$30,000.

I also sought to see if animation would be of value. In the early days of scientific filmmaking, there was no desktop system that had the needed capability. Large research labs typically used a DICOMED film recorder at a cost of some US$250,000, often interfaced to a Cray supercomputer. My Branch Head was unwilling to pay for using the NRL Cray, so we wrote software to convert the simulation data stored on magnetic tapes into DICOMED format magnetic tapes using an HP minicomputer. It took roughly 100 hours of tape spinning to produce the DICOMED compatible tapes for a one-minute animation. We found that the animation brought forth science not found by the previous analysis method of thumbing through stacks of contour plots. Examples, such as the entrainment of saltfingers into the path of earlier ones, can be seen in the ocean saltfingering animation.[2]

Because of power requirements for sending acoustic signals through water, acoustical imaging produces under-sampled, low-resolution data compared to optical imaging. However, optical imaging fails in turbid water, a frequent operating environment where imaging of an object from close-up is required. In a joint effort, Ed Belcher at the University of Washington's Applied Physics Laboratory developed a solid acoustic lens, a high-resolution, forward-looking sonar for 3-D underwater imaging. In parallel, Behzad Kamgar-Parsi and I at NRL developed algorithms using volume graphics and digital imaging[3] to reconstruct near-field images from Ed's sonar data (see Figure 2). This research laid the groundwork for future commercialization of near-field sonar imaging systems.

Bathymetric imaging involved sending ships along a track with a sonar that produced 2-D data for each ping. Ships would take parallel data swaths. However, in the 1980s we lacked satellite coverage for obtaining position and ship time was too expensive to obtain overlapping swaths. After several parallel tracks were obtained, a perpendicular path would be traversed. However, due to currents, there could be discrepancies of hundreds of meters between the original swath and the perpendicular one. A joint program between NRL and the University of Maryland developed computer vision algorithms demonstrating the feasibility of automating the registration necessary to correctly align these swaths. This study demonstrated the possibility of producing maps of the ocean floor without expert bathymetrists.[4]

By the late 1980s, scientific visualization was successfully being applied across varied science and engineering topics and a new field was emerging. At NRL, I organized a day-long "Visualization Day," attended by some 200 NRL scientists, where experts from several scientific areas discussed advances and techniques (thanks to Maxine Brown for assistance). I followed up with a Graphics Hardware Demonstration Day and brought Silicon Graphics Inc. (SGI), Ardent, Stellar, and Pixar to NRL. A further demonstration with Ardent used matched-field acoustic data to calculate a complex 3-D color-coded volume (a cube) in two minutes and allowed interactive exploration by slicing through it in real time. Current NRL methodology was to compute the matched field using a DEC computer and required an overnight run. The scientist then spread slice-by-slice printouts on a tabletop to examine individual 2-D shades-of-gray slices of the volume. The goal was to mentally integrate the volume's complex interior, a daunting task. The faster processing and interactive analysis were far superior and several NRL Branches purchased graphics minicomputers as quickly as funding permitted.

In 1988, I accepted an invitation to chair the IEEE-Computer Society's Technical Committee (TC) on Computer Graphics (quickly renamed Visualization and Graphics TC) realizing that visualization needed its own conference and activities. Arie Kaufman, Gregory Nielson, Bruce Brown, and Carol Hunter joined me on an organizing committee to establish the first IEEE Visualization Conference in 1990. With keynote speaker former National Science Foundation (NSF) Director Erich Bloch, keynote panelist SGI-founder Jim Clark, and an attendance that exceeded expectations the conference was successful. One item we got right was anticipating the future and not naming the conference "Scientific Visualization." Over the years, the emphasis has indeed changed and today information visualization and data analytics are the major thrusts. Nielson and I then guest edited an issue of *IEEE Computer Graphics and Applications* (CG&A) devoted to updated versions of the best conference papers and I established CG&A's Visualization Blackboard Department.

Looking back at this decade, what stands out is that a large number of ocean science problems benefitted from the use of computer graphics and computer vision approaches. Some advances came about by applying methods such as animation and color that would be trivial in a few years, but at the time required novel and often expensive hardware and significant effort to use the hardware. In other cases, the algorithmic development was significant and introduced new methodology to the ocean science topic. Similar activities were taking place across many areas of science, and visualization emerged as a research field in its own right. The graphics minicomputer, although not inexpensive, played a significant role in making these advances accessible to researchers in universities and research laboratories.

INTERLUDE

The Office of Naval Research (ONR) was established by an act of Congress in 1946 as the U.S. government's first permanent agency devoted to funding civilian scientific research during peacetime. As a major funder of scientific research, ONR has had a significant role in fostering scientific and technological innovations in a wide range of fields, as well as in maintaining the basic scientific research infrastructure that makes these breakthroughs possible.

As the 1990s rolled around, I was becoming interested in virtual reality (VR). I held discussions with others but failed to arouse interest in initiating VR activity within NRL. However, an interesting career opportunity arose. Since the late 1940s, ONR had maintained a small European Office (ONREUR) based in London (now part of ONR Global) for the purpose of liaison with the European scientific community. These positions were typically one-to-three-year tours covering specialized topics of interest to ONR and filled primarily by university or Navy laboratory personnel. The computer science rotator was coming open and I accepted an invitation for a two-year tour to focus on computer graphics, visualization, and related topics.

The position's responsibilities consisted of arranging (geographically grouped) meetings with leading researchers in visualization, graphics, and related areas, giving talks on U.S. research, and finding mechanisms to report on European activities. During my two years, I distributed by email approximately 45 summaries of European activity under the title "Realization Reports," several of which were reprinted in *Computer Graphics,* ACM SIGGRAPH's quarterly newsletter. Longer articles also appeared in the *ONR European Science Notes Information Bulletin*. José Encarnação, Martin Göbel, and I authored a CG&A article summarizing European VR activities.[5]

ONR and the Air Force Office of Scientific Research annually funded the National Academy of Sciences to arrange a trip for a leading U.S. scientist to various parts of the world. The scientist gave lectures and meetings were held to discuss activities of mutual interest. Computer science was the topic area for 1993. Turing Award laureate John Hopcroft, Cornell University, was to be the featured speaker. The countries to be visited were Poland, Hungary, Russia, and Ukraine. Many challenging issues and interesting experiences arose.

For example, I received a call from ONR asking for help setting up the Moscow visit as they were having little success. I tried to get in touch with Stas Klimenko, whom I knew from the IEEE Visualization Conference. This was 1993 and communicating with colleagues within Russia proved difficult. Fortunately, Stas had an upcoming trip to Germany and we had a detailed phone conversation. Stas brought aboard a colleague who was an associate member of the Russian Academy of Sciences. Arrangements now went smoothly and an interesting agenda was established. Among the sites we visited were Moscow State University, a specialized science high school, and the first declassified Elbrus supercomputer. I cannot fully vouch for my memory nearly 30 years later, but I recall a big motherboard with large, water-filled copper pipes running across the board for cooling.

For our off day before heading to Ukraine, it was suggested that we visit a monastery some 60 miles outside of Moscow. We requested to visit the Kremlin instead. Our guide seemed reluctant but set up a morning visit. On arrival, we found that the Kremlin was closed to visitors due to an appearance by Yeltsin before the Russian Parliament later that day and protests were anticipated. We went around the corner to Lenin's tomb in Red Square and the GUM mall, then to the Pushkin Museum, and finally to the Bolshoi for the evening opera event.

To examine the current status and shortfalls of the visualization field, I organized a workshop of world-wide experts. The workshop was held in 1993 at the Fraunhofer Institute for Graphics Research in Darmstadt, Germany. Updated presentations became chapters in a book and working groups, after months of effort, wrote a concluding chapter on research issues.[6] The book was used to teach many of the first academic visualization classes while awaiting texts designed for the classroom. It was to become the groundwork for an ONR-funded university research program.

This position, with its extensive travel in Europe, sometimes sounds like a mini-vacation. Actually, I never worked harder or longer hours. However, it was an enjoyable, self-contained job. With the sole exception of the Hopcroft trip, there were no assignments, deadlines, interoffice memos, internal meetings, sudden interruptions "because X is needed by Friday," etc. I decided who to visit, what to report, how to distribute the information, and what, if any, extra tasks I would undertake (in my case the workshop and book). Another bonus was living in the West End of London with a seven-minute walk to the office instead of my unpleasant 45-minute automobile commute in heavy traffic back in DC.

VIRTUAL/AUGMENTED REALITY

In April 1994, I returned to NRL and took a position in the Information Technology Division with the goal of establishing a VR Laboratory. There was little understanding of VR and its applicability to Navy problems. I was given a large room with a sizable screen and projector, a Fakespace Boom, and a head-mounted display (HMD) and had part-time access to an SGI computer. I began by bringing in graduate students from the Georgia Institute of Technology (Georgia Tech) and the University of North Carolina Chapel Hill as summer interns. One of them, from Larry Hodges' VR Group at Georgia Tech, produced a complex model of a portion of a decommissioned Navy ship (see Figure 3).

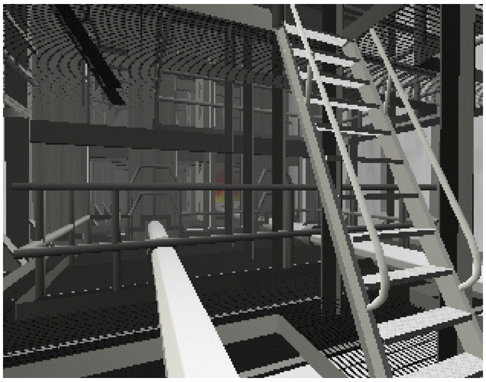

FIGURE 3. Model of an area containing complex geometry such as grated walk-ways, chain-link fences, stairs, and pipes.

For interaction, trackers attached to the HMD maintained the correct display position and orientation. Users could *fly* in the direction they were looking by depressing buttons on a 3-D FliteStik. The virtual environment included a dynamically generated virtual fire made up of approximately 500 polygons. Using a mixture of physics-based modeling and fractal techniques, the fire changed color and transparency levels to simulate the appearance of real flames. Smoke could be added to the environment. This walk-through helped convince management that VR was potentially applicable to Naval problems.

The next step in developing the VR Laboratory was to obtain high-end equipment and research funding. Toward the end of my ONREUR tour, Wolfgang Kreuger at the German National Research Center for Information Technology invited me to visit to see a new tabletop VR display he had developed. The idea of projecting up through a tabletop to get a 3-D image rising above the surface was both simple and innovative. It allowed for *god's eye* VR, where one wanted to see the image in 3-D from a distance rather than walk through it. This was eminently appropriate for command-and-control applications in a command center, a topic of Navy interest. With Kreuger's encouragement, we designed and developed (using NRL's woodworking shop) our own VR Responsive Workbench (VRRW) (see Figure 4). Subsequent discussions with Ian McDowall and Mark Bolas (Fakespace) played a role in their designing and commercializing a version. We were pleased to purchase their first commercial workbench. Having my people spend time in a woodworking shop was not an efficient use of research dollars.

At the same time, the VR Lab submitted a proposal for funding for a Virtual Reality Immersive Room which we called a GROTTO (see Figure 5) to an NRL-wide

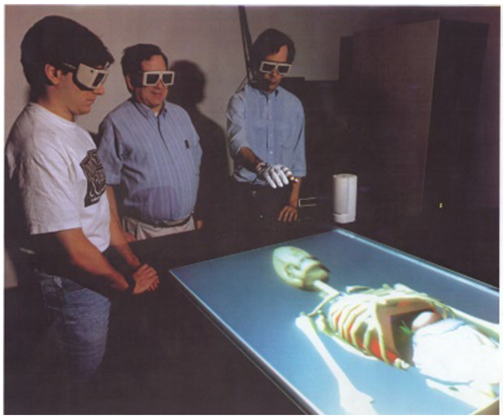

FIGURE 4. Initial demonstration of the NRL-fabricated VRRW, Fall 1994.

FIGURE 5. GROTTO is used for (top left and clockwise) studying materials properties, CAD/CAM, a ship walkthrough, and computational steering.

equipment funding program that was designed to support the fabrication of expensive new facilities. Our $700,000 proposal included a high-end SGI computer, four high-end projectors to place images on three walls and the floor, and miscellaneous other items. In addition to the technical merit, our Division Superintendent enthusiastically supported the proposal because the Division had never submitted to this funding pool but was taxed to support it. The proposal was funded.

At this price only a handful of similar facilities were being built. The SGI's annual maintenance agreement cost nearly 10% of the SGI's purchase price, a problem for us as project funds paid for it. Fortunately, technology improvements helped out. A couple of research groups developed hardware for PCs that allowed using the PC to drive and sync the display. Next, a commercial company manufactured a card for this purpose. Shortly thereafter, PC internal hardware could handle the task. We now purchased PCs to run the GROTTO, a substantial saving.

Even at the cost reduction from using PCs, the price for an immersive room facility remained high enough that few research labs and universities purchased them. This resulted in the lack of a growing, active user community. In turn, this may have contributed to the failure to produce enough examples of the immersive room being successfully demonstrated to generate a demand for research to overcome technical limitations.

Our VR projects over the next few years included:

- Using the VRRW to examine a preliminary design of a proposed new type of ship. Among the items noted were missing stairs and a previously unnoted subtle double curvature of the hull.
- Using the GROTTO and VRRW for the visualization of scientific data. We found VR excellent for demonstrations and collaborative study, but it did not produce scientific advances at the level that color and animation had in the 1980s.
- Conducting a formal experiment to quantify the benefits of using VR for training shipboard firefighters. A user study showed that VR had value for training.
- Experimenting with distributed VR for engineering design.
- Developing an interoperable VR system for maintaining a common operating picture given display devices.
- Developing a VRRW system for the command center of a Marine Corps warfighting experiment in which we participated (see Figure 6).
- Visualizing multidimensional environmental uncertainty using data from a large ONR-funded ocean structure and acoustics experiment. Results are discussed in Schmidt *et al.*[7].

The above projects represent a fairly broad subset of the set of application areas. Research teams elsewhere were experimenting with the value and efficacy of VR with similar results. Thus, this work represents a reasonable snapshot of where the field stood. The move from the research laboratory to the commercial world has been slow but ongoing.

In 1991, Steve Feiner (Columbia University) received an ONR Young Investigator Award for research on "Automated Generation of Three-Dimensional Virtual Worlds for Task Explanation." I had included his work in talks I gave in Europe about U.S. computer science research while at ONREUR and followed the

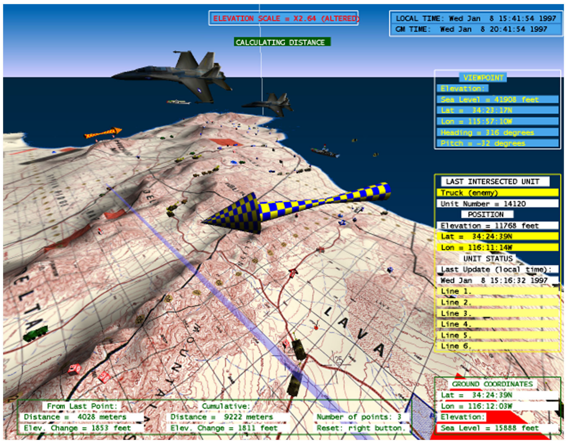

FIGURE 6. Third-dimensional display on the VRRW used in the command center during a warfighting experiment.

FIGURE 7. NRL BARS system.

subsequent developments. I later discussed it in ONR funding-related meetings and expected others to pick up on it, but this did not happen. Therefore, in 1997, I worked with Feiner and VR Lab scientist Simon Julier on a proposal to ONR to explore how mobile augmented reality (AR) could be developed to improve battlefield performance. Recent events in Somalia, as later depicted in the movie *Blackhawk Down* where the ability to give troops in the field a heads-up representation of street names and routing instructions would have saved lives, strengthened our presentation. The project was funded. In general, demonstrating how a novel technology may be applicable to resolving shortcomings that are clearly identifiable in a high-profile, real-life event enhances the case for funding.

The VR Lab's funds were supplemented by an NRL Base Program Award. Our AR funding was now nearly a million dollars per year, allowing for research into several challenging problems. The program, called the Battlefield Augmented Reality System (BARS), investigated how multiple mobile AR users on foot could cooperate effectively with one another and with personnel in combat operations centers. Several challenges became apparent: building and maintaining environmental models of a complex and dynamic scene, managing the information relevant to military operations, and interacting with this information. To achieve such a system, the software architecture to encapsulate these features had to be developed. Although this also required high-fidelity tracking of multiple mobile users, our primary focus was on the information management and interaction components.

The display space for a mobile AR system is limited. In order to utilize the technology in a 3-D urban environment, methods were needed to determine what to display. Based in part on the user's spatial relationship to items of interest, algorithms determined the information that is most relevant to the user. User interface component design determined how the selected information should be conveyed, based on the kind of display available, and how accurately the user and objects of interest could be tracked relative to each other. Our work on view management introduced an efficient way of allocating and querying space on the viewplane, dynamically accounting for obscuration relationships among objects relative to the user. AR authoring tools were developed to allow content developers to create richer AR experiences. A key concept was to combine a 2-D media timeline editor, similar to that used in existing multimedia authoring systems, with a 3-D spatial editor to allow authors to graphically position media objects in a representation of the 3-D environment.

Two systems, one based on consumer grade hardware and the other using embedded computers, were developed. There was a direct tradeoff of capability and weight versus usability. Both systems used Sony Glasstron optical see-through head-worn displays, and a loosely integrated tracking solution consisting of a real-time kinematic GPS receiver and an orientation sensor. The first demonstration of BARS (see Figure 7) was in November 1999. NRL and Columbia demonstrated some of this joint work at ISWC 2000, showing the new backpack systems. At SIGGRAPH's Emerging Technologies in 2001, we demonstrated integration with wide-area tracking in a joint effort with InterSense.

Domain analysis revealed that one challenge of urban operations is maintaining understanding of the location of forces that are hidden by urban infrastructure. This is called the "X-ray vision" problem: Given the ability to see "through" objects with an AR system, how does one determine how to effectively represent the locations of the occluded objects? VR Lab researchers developed visualization techniques that could communicate the location of graphical entities with respect to the real environment. We implemented a range of graphical parameters for hidden objects and were the first to study objects at far-field distances of 60–500 meters, identifying visualization parameters such as drawing style, opacity settings, and intensity settings that could compensate for the lack of being able to rely on a consistent ground plane. A user study was conducted in which several VR Lab members and Virginia Polytechnic Institute researchers Debbie Hix and Joe Gabbard examined which of the numerous graphical parameters were most effective.[8]

In addition, we developed a robust event-based data distribution mechanism for mobile augmented reality and virtual environments. It was based on replicated databases, pluggable networking protocols, and communication channels. This was demonstrated using our mobile augmented reality systems, immersive and desktop-based VR systems, a 2-D map-based multimodal system, handheld PCs, and other sources of information such as external data servers.[9]

Even with an effort of this scope, many critical AR issues remained, including highly accurate registration needed for certain applications and improved computer and display hardware. The Sony Glasstron was the best display available but it had significant limitations, including insufficient brightness so that the display would wash out in bright sunlight. For important demonstrations, I not-so-fondly recollect pleading for the weather to be overcast (for a clear display) but not raining (as the system components were exposed). A detailed description of our development of mobile AR systems is given in Rosenblum *et al.*[10] When I left NRL in 2004, the AR research continued under the leadership of Mark Livingston for another decade, albeit at a smaller funding level. The BARS research began to transition to other agencies that would take the technology to prototype development and ultimately into training and operations systems. NRL supported these programs with focused research on human factors for the systems being developed. A BARS research thread continues currently with a study by Livingston on how to deal with cluttered views in virtual environments.

Both VR and AR have developed significantly in the decades since this work was performed. VR did not prove especially valuable for scientific visualization. However, driven in part by the gaming industry, low-priced systems are now used across many applications. Immersive rooms were not particularly successful. One reason was that, given the expense, management expected them to be funding sources. Instead, they were funding sinks, at least for those organizations that purchased them as a visualization aid to support other research groups but still expected them to be self-funding.

As we conjectured, AR has proven to be important with many applications now forthcoming. It is rewarding to see that, 25 years later, an extended version of the BARS concept has been implemented. Microsoft has recently been awarded a contract worth 22 billion US dollars over 10 years to deliver more than 120,000 such devices to the U.S. Army (see: “US Army awards Microsoft $22 billion IVAS contract to deliver Augmented Reality technology to US Soldiers,” *AR and VR Industry News and Intelligence,* April 1, 2021).

RESEARCH FUNDING (ONR)

As part of my work as VR Lab Director at NRL, I took a two-day per week detail to ONR to provide expertise in graphics/visualization. This placed me in a position where I could participate in internal competitions for research programs to fund academia and industry. For example, a winning proposal to the Defense University Research Instrumentation Program resulted in two million dollars being awarded to university researchers working in graphics, visualization, VR, and related areas for the acquisition of equipment to augment current or develop new research capabilities. Another internal competition offered the potential to fund university research. To compete, I used the book[6] from the 1993 ONREUR workshop to demonstrate the need for research advances, emphasizing the chapter on volume graphics. Funding was obtained to award six three-year grants to university researchers for modeling algorithms, including volume graphics, large-scale models, and level sets.

With BARS demonstrating the potential of mobile AR, Andre van Tilborg, then the Director of the Mathematical, Computer, and Information Sciences and Technology Division at ONR, asked me to assemble an AR research program to complement the BARS research. Funding for university grants came from the ONR base research program. Other funding sources included relevant portions of a Multidisciplinary University Research Initiative and certain U.S. Congress directed funds. Combined with BARS, I believe this program was the largest funded mobile AR research effort in its time period. Funding varied from year to year and totaled several million dollars a year for five years (in addition to the BARS funding). Topics included tracking and registration, usability of mobile AR systems, 3-D urban terrain reconstruction, and displays. Not all efforts were successful (for example, a retinal scanning display never became the see-through display that AR required), but in total, the research provided a groundwork for advancing the field.

DATA AND VISUAL ANALYTICS (DAVA)

In 2004, recent changes in management both at ONR and NRL had made it unlikely that BARS funding would continue at the current level. Therefore, I decided to accept a position at the National Science Foundation (NSF) to serve as lead for the graphics and visualization program in the Computer and Information Science and Engineering (CISE) Directorate. NSF Program Directors (PDs) bear the primary responsibility for carrying out the agency's overall mission to support innovative and merit-reviewed activities in basic research and education within the PD's expertise that contribute to the nation's technical strength, security, and welfare.

One graphics area I thought merited accelerated development was computational photography. To this end, I added a specific statement to the annual solicitation for proposals on this topic and it resulted in several excellent ideas that were funded. This was a minor accomplishment with minimal effort, but it indicates how a PD can influence a field.

Even before arriving at NSF, I had one idea that I hoped to implement. Visual analytics is defined to be the science of analytical reasoning facilitated by interactive visualizations. The goal is to use visual methods to enhance understanding of massive, complex datasets. Such datasets arise in science, engineering, and commerce, as well as in intelligence analysis and emergency response. Analyzing these datasets is a major intellectual challenge and existing methods focused primarily on static visualizations. However, the analysis of large, complex datasets often requires advances from disparate areas of mathematics, statistics, and computational science prior to but integrated with the visualization stage. A dataset may be so large and complex that it becomes unmanageable using interactive visual techniques. When this occurs, the idea is to transform it via mathematical algorithms into a dataset of lower dimensionality that better enables analysis tools. Now apply visual analytics techniques for data understanding, possibly including a feedback loop utilizing the dimensionality reduction algorithms. This approach is potentially beneficial even if one could utilize the original dataset for visual analysis.

At the time, there was little push in the visual analytics community to expand the field to import these computational techniques and little awareness among the mathematics and computational science communities of the role they could play. I discussed my idea with Pat Hanrahan of Stanford University and then with Joe Kielman of the Department of Homeland Security (DHS) and Jim Thomas of the National Visualization and Analytics Center. Pat, Joe, and Jim were developing an application-specific program in visual analytics for DHS. All of us agreed a foundational research program was needed and we decided to try to establish a joint NSF/DHS effort. DAVA was clearly going to impact a wide range of societal issues and thus was a suitable topic for NSF.

How does one go about establishing a funded multiagency research program? The first step was to link DHS and NSF by cross-serving on DHS and NSF review panels and other related activities. Next was a workshop to further define the scientific goals and the research needed to reach these goals. This was particularly important and difficult. Generally, such a workshop would bring in area experts with much commonality, but here we wanted to involve people from outside the graphics/visualization arena. Researchers who had never heard the term "visual analytics" would have to be convinced that they had a vital role to play.

Pat agreed to serve as Lead Organizer of the Mathematics of Visual Analysis Workshop held at the Mathematical Sciences Research Institute at the University of California, Berkeley in October 2006. Co-organizers were William Cleveland, Sanda Harabagiu, Peter Jones, Leland Wilkinson, and Nancy Chinchor. The list of invited attendees consisted of a mixture of visualization experts and those from relevant areas of mathematics, statistics, and computer science. The goal was to define mathematical and computational problems involved in DAVA as well as to engage the mathematical community and introduce them to this application area.

The Workshop was successful. We agreed that the program would have NSF-style research focused on developing the underlying science for DAVA without concern for any particular application. Reviews of grant proposals would be organized by NSF and use the NSF structure for award decisions based on input from external review panels of subject area experts. NSF's Mathematics Division joined the effort as did a second CISE Division. Both contributed additional funding that gave us a budget of approximately two and a half million dollars a year for five years. PDs from both groups (Tie Luo, Ephram Glinert, Maria Zemankova) joined the management team. A memorandum of agreement was created and signed by NSF and DHS executives indicating what was to be accomplished and organizational obligations.

The next step was to structure the research program. We titled the program Foundations of Data and Visual Analytics (FODAVA) in order to clarify that the program was focused on new mathematical and computational methods to be integrated with visual analytics and that, while some tie to visualization methods was expected, the proposals were not to emphasize visualization methods. There would be one large five-year grant to a "Lead" institution. This institution would perform research across several areas but would also have major responsibilities for integrating other FODAVA-partner research. In close collaboration with the FODAVA-partners, they would lead in developing activities such as workshops, tutorials, short courses, and summer programs for students. There were 17 three-year FODAVA-partner research grants awarded to universities, with the additional provisions that some integration with visual methods was required. Principal Invesigators (PIs) would actively work in close coordination with the Lead institution to provide exposure for the topic. Georgia Tech submitted the winning Lead proposal with Haesun Park as Director. The three-year grants were awarded primarily in years one and two of FODAVA, with extensions possible for years four and five. Full details of the program, including the PIs research topics and slides from the annual program reviews, can be found at https://fodava.gatech.edu/.

In addition to standard problems in developing a research program, FODAVA required bringing together disparate, unrelated groups of researchers, a difficult but critical undertaking. I will briefly illustrate. I wanted to get my NSF colleague and frequent lunch partner Sankar Basu involved. For a month, I heard "I'm not a visualization person, what can I contribute?" I finally convinced him that the underlying research was primarily in mathematical and computational science and he joined the team. His insight and assistance proved invaluable. I asked why Mauro Maggioni (Duke) had applied for a grant. Our 2006 Berkeley Workshop co-chair Peter Jones, then his Department Chair at Yale, suggested he examine the solicitation. Mauro realized that there were challenging problems within his discipline. Figure 8 shows the user interface for one of his research outputs utilizing wavelets. Leland Wilkinson stated he had lost interest in visualization, but that FODAVA was moving in an exciting and important direction and he was reinvigorated.

Haesun Park and her coauthors received the Test of Time (10 years) award at VIS 2020 for their VIS 2010 paper "iVisClassifier: An interactive visual analytics system for classification based on supervised dimension reduction." Given high-dimensional data and cluster labels, linear discriminant analysis gives a reduced dimensional representation and yields a good

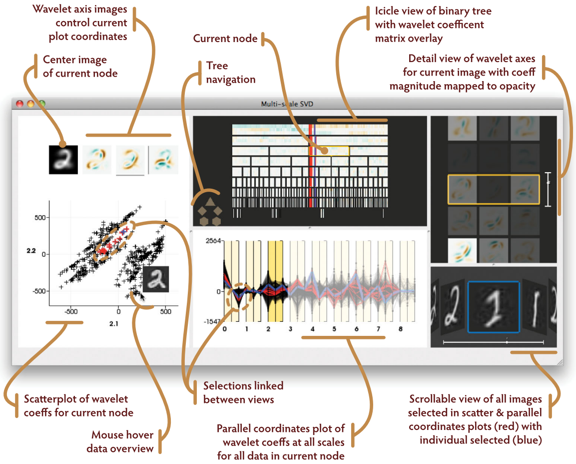

FIGURE 8. User interface for geometric multiresolution. (Source: Mauro Maggioni, Eric Monson, and Rachael Brady, Duke University; used with permission.)

overview of the cluster structure. iVisClassifier interacts with the reduced dimensions using parallel coordinates and a scatter plot. Haesun was surprised she won because 10 years earlier reviewers were telling her that her research submissions were interesting but lacked innovation in visualization. Haesun's experience is a good illustration of FODAVA's programmatic difficulty. Yes, we had to convince other domain scientists and engineers that this was an important problem that merited their time and effort. However, it was equally difficult to convince the visual analytics community of the necessity of bringing computational science researchers into the fold. Some "got it," but many did not. Haesun's award indicates improvement over the last decade.

Incidentally, Haesun told me she put together the winning Lead proposal because of conversations she and I had when we briefly overlapped at NSF. She had no previous knowledge of visual analytics and was unlikely to have examined the FODAVA solicitation or realize how well her work would mesh with DAVA.

The FODAVA program succeeded in establishing DAVA as a research discipline with expanding interest. As one example, the NSF-funded Statistical and Applied Mathematical Sciences Institute (SAMSI) held a SAMSI-FODAVA Workshop on Interactive Visualization and Analysis of Massive Data in December 2012 led by the FODAVA-Lead team, bringing in many researchers who were new to the topic. The workshop investigated "the mathematical, statistical, and algorithmic issues in efficient representation and transformation of data, scalable and dynamic algorithms for real-time interaction, visual representation in limited screen space, performing evaluations, and applications."

The movement of DAVA into the academic curricula has been dramatic. Early in FODAVA, the FODAVA-Lead group used materials from FODAVA activities to assemble a new course (CSE 6242 Data and Visual Analytics - DAVA) taught in Georgia Tech's School of Computational Science and Engineering. It was first taught in the spring of 2013 with 35 students attending. By the 2016/2017 academic year, it was taken by 429 students. DAVA became an advanced core class in a new MS Analytics degree program which is now jointly offered in both on- and off-campus versions by a three school/college consortium at Georgia Tech. In academic year 2020/2021, the course was taken by 2157 students (about 600 on-campus). Now recognized as a critical pillar supporting discovery and innovation, DAVA courses are being developed throughout academia. The data scientists needed to exploit the large, complex data sets generated by science, engineering, and societal problems are being educated!

CONCLUSION

In 1995, after seven years of work, Andrew Wiles, a number theorist, produced a proof of Fermat's Last Theorem (probably the most famous problem in the history of mathematics). He was asked if he knew all along he was going to succeed. His answer was "Maybe I wouldn't get the whole thing . . . I just felt I was making progress." This approach to problem solving is perhaps intellectually closer to how one would create a symphony than to problem solving via the scientific method. In a far humbler way, I took a similar approach.

Looking back, my achievements came by asking questions such as: What new methods could provide better solutions to problems than current practice? Is there an emergent technology that is likely to be important and relevant? What is limiting the use of this approach? Exploring potentially breakthrough ideas is more intellectually challenging and adventurous than making (often minor) improvements to current practice and will frequently produce significant advances.

My direct involvement with research was rewarding. However, I was in the right place at the right time to initiate funded programs that helped to grow important research fields. My work as a Program Director at funding agencies was an equally rewarding aspect of my career and I encourage others to seek out such opportunities.

(Portions of this article rely on decades old memories and readers should be aware that omissions or inaccuracies are possible.)

ACKNOWLEDGMENTS

The list of highly capable coworkers, senior researchers, and administrators with whom I interacted are far too long to acknowledge individually. I thank them all. Simon Julier, Mark Livingston, and Ed Swan displayed dedication and technical excellence as my Project Leaders at the NRL VR Lab. A special thanks to Andre van Tilborg who selected me for the ONREUR position and provided inspiration and funding support for the development of BARS. I also thank Editor David Kasik for his editing skills and constructive comments.

REFERENCES

1. G. Marmorino, L. Rosenblum, J.P. Dugan, and C.Y. Shen, "Temperature fine-structure with a patches near an upper ocean density front," *J. Geophys. Res.*, vol. 90, no. C6, pp. 11799–11810 and 11843–11845, Nov. 1985.
2. L. Rosenblum and C. Shen, Video: "Double diffusive convection saltfingering," ACM Siggraph Video Rev., no. 49, (video #11), Jul. 1989.
3. L. Rosenblum and B. Kamgar-Parsi, "3-D reconstruction of small underwater objects using high resolution sonar data," in *Proc. Symp. Auton. Underwater Veh.*, Jun. 1992, pp. 228–235.
4. B. KamgarParsi, L. Rosenblum, F. Pipitone, L. Davis, and J. Jones, "Toward an automated system for a correctly registered bathymetric chart," *IEEE J. Ocean. Eng.*, vol. 14, no. 4, pp. 319–325, Oct. 1989.
5. J. Encarnação, M. Göbel, and L. Rosenblum, "European activities in virtual reality," *IEEE Comput. Graph. Appl.*, vol. 14, no. 1, pp. 66–74, Jan. 1994.
6. L. Rosenblum *et al.*, *Scientific Visualization: Advances & Challenges*, London, U. K.: Academic Press Ltd., 1994.
7. G. Schmidt, S.-L. Chen, A. Bryden, M. Livingston, L. Rosenblum, and B. Osborn, "Multidimensional visual representation for underwater environmental uncertainty," *IEEE Comput. Graph. Appl.*, vol. 24, no. 3, pp. 56–65, Sep. 2004.
8. M. Livingston *et al.*, "Resolving multiple occluded layers in augmented reality," in *Proc. Int. Symp. Mixed Augmented Reality 2003*, pp. 56–65, Oct. 2003.
9. D. Brown, S. Julier, Y. Baillot, M. Livingston, and L. Rosenblum, "An event-based data distribution mechanism for collaborative mobile augmented reality and virtual environments," *Presence*, vol. 13, no. 2, pp. 211–221, Mar. 2004.

10. L. Rosenblum, S. Feiner, S. Julier, J. E. Swan II, and M. Livingston, "The development of mobile augmented reality," in *Expanding the Frontiers of Visual Analytics and Visualization*, J. Dill, R. Earnshaw, D. Kasik, J. Vince, and P. Wong eds. New York, NY, USA: Springer, 2012, pp. 431–448.

LAWRENCE (LARRY) J. ROSENBLUM is a retired U.S. federal government employee who is currently a Visiting Senior Research Scientist at the University of Maryland Institute for Advanced Computer Studies, College Park, MD, USA. He received the Ph.D. degree in mathematics (number theory) from The Ohio State University, Columbus, OH, USA, in 1971. He was the recipient of the IEEE VGTC Visualization Career Award, the IEEE VGTC Virtual Reality Career Award, a DHS/National Visualization and Analytics Center Award for leadership in the field of visual analytics, and an IEEE-CS Outstanding Service Award for co-founding the Visualization Conference. He was among the initial inductees to the newly formed IEEE Visualization Academy in 2019. He served on the Editorial Boards of *IEEE Computer Graphics and Applications*, IEEE Transactions on Visualization and Computer Graphics, and several other journals. He is a Senior Member of the IEEE and a Golden Core 500 Awardee of IEEE-CS. Contact him at ljrosen@verizon.net.

Contact department editor Dave Kasik at dave.kasik@gmail.com, department editor Chris Johnson at crj@sci.utah.edu, or department editor Mary C. Whitton at mcwhitton@gmail.com.