Lightning and Thunder: The Early Days of Interactive Information Visualization at the University of Maryland

Catherine Plaisant and Ben Shneiderman, *University of Maryland, College Park, MD, 20742, USA*

The thrill of scientific discovery, the excitement of engineering development, and the fresh thinking of design explorations were invigorating as we participated in the birth of a new discipline: Information Visualization. This discipline, based on graphical user interfaces with pointing devices, became possible as software matured, hardware sped up, and screen resolution improved. Driven by the concepts of direct manipulation and dynamic queries, we made interactive interfaces that empowered users and opened up new possibilities for the next generation of designers. We worked with professionals who had real problems and tested real users to get their feedback. Some projects failed and some papers never got published, but many of the new ideas found their way into widely used commercial products. Our great satisfaction is that our students have spread the community spirit of the Human-Computer Interaction Laboratory as they continue to make further contributions.

The early days of information visualization at the University of Maryland Human-Computer Interaction Laboratory in the 1980s and 1990s was a time of bold ideas that took advantage of increasingly flexible software and the increasingly speedy hardware. The bold ideas triggered new visual designs using color, animation, and interaction, which were controlled by innovative user interface widgets based on mouse hovering, clicking, dragging, and sliding. As touchscreens became more widely available, high precision pointing and other gestures became possible. It seemed that each innovation was a bolt of lightning that flashed through the minds of our team as they refined and applied the ideas in new domains, with new devices, for new audiences. The lightning was visible in conference presentations and journal publications, while the thunder reverberated farther away in successful applications, widely-used toolkits, and improved algorithms.

0272-1716 © 2022 IEEE
Digital Object Identifier 10.1109/MCG.2022.3180012
Date of current version 15 July 2022.

SMALL TIGHT-KNIT COMMUNITY

We worked as a community of students, staff, and faculty, collaborating extensively but each person was clear about what their task was and what topic they were leading. We celebrated successes with champagne and strawberries, putting each bottle labeled with the date on a high shelf, until the campus prohibited alcohol at events. We honestly discussed failed studies, rejected papers, or other setbacks so that we could all learn from mistakes and feel comfortable getting advice about next steps.[1] We gave all our talks in the laboratory before giving them at conferences—and still follow that principle today.

We enjoyed each other's company, leading us to do annual day-long retreats at Maryland State Parks or a heritage farm on the Potomac (see Figure 1). The key question was: What will you accomplish during the next year?

FIGURE 1. 1991 HCIL Laboratory retreat.

Winters were made warmer by lunchtime group visits to the nearby skating rink or when the snow was just right, we departed at 3:30 P.M. to drive 75 mins to Whitetail ski resort, just across the Pennsylvania border, for night skiing and a meal together.

EARLY LABORATORY ENVIRONMENT

Ben started the University of Maryland's Human-Computer Interaction Laboratory (HCIL) in 1983. When Catherine joined in 1987, we still read email on text-only green screens. Soon we relished our new IBM PCs with VGA resolution 640 × 480 pixels. Those "personal computers" were not connected to anything since the World Wide Web did not exist. We used early information retrieval tools, but as data volumes and variety grew, we taught each other to do database searches using clumsy query languages or confusing fill-in forms. These experiences led us to recognize that there were better ways. The HCIL had modest resources, but collaboration gave us strength. We gave our talks using plastic slides placed on overhead projectors. By 1990, we traveled with a bag of VHS tapes to demonstrate interactivity with our recorded demos.

WORKING WITH OUTSIDE PARTNERS

Close collaborations with diverse practitioner partners outside the campus challenged us with meaningful problems (museums, startup companies, NASA, medical professionals, etc.). Working with them proved to be the primary source of inspiration for new directions and ideas, and led to grant proposals or short-term student projects, which tried out solutions to the novel problems. In other cases, our partners provided sample data and helped us refine ideas and validate them, so the resulting crude software applications (which we called "inspirational prototypes") illustrated real situations and often led new partners to contact us to push the idea one step further. In rare cases, we targeted our own needs—such as understanding how our tiny hard disks could be better utilized.

We put substantial energy into validating new ideas with user studies to keep us honest in our claims. Kent Norman in the Psychology Department was a vital early partner. He taught us how to do rigorous studies and choose the right statistical methods.

CLOSE COLLABORATIONS WITH DIVERSE PRACTITIONER PARTNERS OUTSIDE THE CAMPUS CHALLENGED US WITH MEANINGFUL PROBLEMS.

EARLY RESEARCH

The emerging graphical user interfaces, some games, and specialized applications (e.g., air traffic control) suggested a set of direct manipulation principles: visual representation of the objects and actions, interaction by pointing devices to replace keyboard commands, and rapid, incremental, and reversible operations.[2] Driven by these principles, the laboratory's research was focused on novel interaction methods, with projects on hypertext systems to click on highlighted words and images, touchscreen interfaces for home control systems, and information visualization tools with slider controls.

Our ideas were refined during our collaboration with the Library of Congress to replace the card catalog with 51 touchscreen kiosks in the main reading room and elsewhere. The designers thoughtfully used the oak panels from the card catalogs to build elegant carrels where users could work comfortably. Later, a memorable and large test of these strategies was in spring 1988 when we designed a public access kiosk using our Hyperties tool for the Caesarea (King Herod's Dream) Exhibit, a Smithsonian Museum exhibit.[3] Our Guide to Opportunities in Volunteer Archaeology (GOVA) may very well have been the first touchscreen museum kiosk. It was used by 75,000+ patrons over several months at six museums. Patrons could browse information about volunteer archaeological digs. Text search was possible, but it was hard to get an idea of what was available, what matched flexible criteria, where and when digs were common, or where there were none. William Weiland later built prototypes to query volunteer dig sites using sliders and buttons, but response time remained slow. We were increasingly excited by visions of interactive exploratory data analysis (i.e., being able to pose dozens of queries very quickly through direct manipulation). These explorations led us to understand that interactive information visualization often gave users answers to questions they did not know they had.

> *WE WERE INCREASINGLY EXCITED BY VISIONS OF INTERACTIVE EXPLORATORY DATA ANALYSIS.*

The next sections tell further stories of our early interactive information visualization innovations. Our articles include the technical details—while here we focus on the human stories. Our projects were enriched by working with professional partners and involving talented students who could quickly build something which tickled our minds and helped develop the ideas further. Since we valued rapid prototypes and quick evolution, our motto was to build "something small soon." Many ideas were thrown away in frank but friendly discussions, and successful projects benefited from comments by many participants.

> *OUR MOTTO WAS TO BUILD "SOMETHING SMALL SOON."*

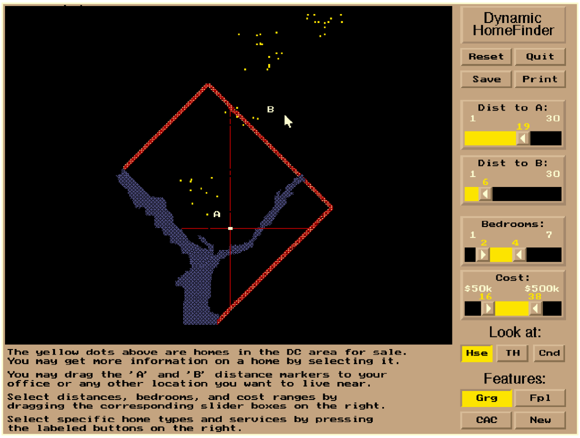

FIGURE 2. HomeFinder dynamic query prototype. A simplified map of Washington DC shows houses for sale. Buttons and sliders are used to rapidly explore what is available. A and B markers indicate work locations.

FROM DIRECT MANIPULATION TO DYNAMIC QUERIES: THE ROAD TO SPOTFIRE

The direct manipulation ideas[2] influenced our movement to do interactive information visualization with maps, diagrams, and scatterplots, manipulated by dynamic queries sliders, guided by the Information Visualization Mantra: *Overview first, zoom and filter, then details-on- demand.*

Dynamic queries[4] and starfield displays introduced direct manipulation to querying and helped bring interactivity to early information visualization (by updating results in under 100 milliseconds.) As users adjusted sliders, buttons, check boxes, and other control widgets the starfield display containing color- and size-coded points or shapes updated rapidly.

Christopher Ahlberg, a visiting student from Sweden during summer 1991, took up a lunch-time challenge and built a slider-based version of a polynomial equations viewer that Ben had built as a graduate student in 1972. Within a week, he had built a dynamic query for the chemical table of elements. Six sliders for attributes such as atomic radius or ionization energy enabled users to filter the elements on the table, showing the clusters, jumps, and gaps. The chemists we invited found the approach fascinating. A study with 18 chemistry students showed faster performance with use of a visual display (versus a simple textual list) and using sliders (versus a form fill-in box).

Later on, Christopher Williamson's HomeFinder showed a map of Washington, DC and 1100 points of light indicating homes for sale (see Figure 2). Users could mark the workplace for both members of a couple and then adjust sliders to select circular areas of varying radii. Other sliders selected number of bedrooms and cost, with buttons for air conditioning, garage, etc. Within a tenth of a second, users could see how many homes matched the queries, and adjust them accordingly. We could not get data digitally, so Christopher Williamson and his friends took the Sunday *Washington Post* classified advertising pages and typed in the data for the 1100 homes. Controlled experiments with benchmark tasks showed dramatic speed-ups in performance and high subjective satisfaction. We were onto something important.

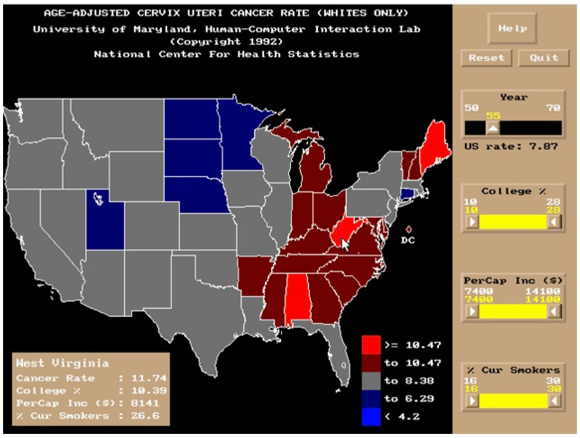

FIGURE 3. Dynamap for Health Statistics. Users first select a cause of death and census data attributes and the interface is created automatically. A time slider reveals changes in death rates over time. Attribute sliders filter the map and allows users to explore possible correlations.

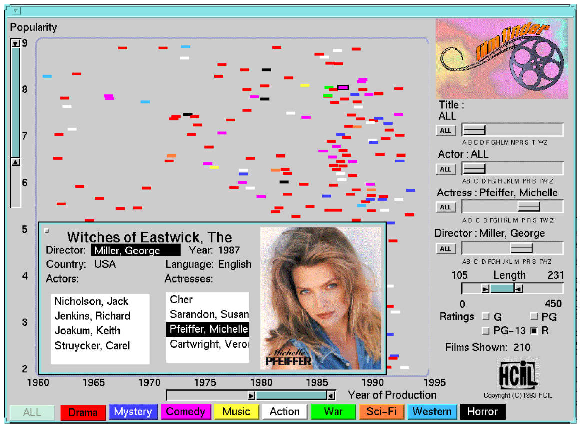

FIGURE 4. FilmFinder prototype showing a collection of films. Two numerical attributes are selected for the *X*- and *Y*-axes, providing an overview of the dataset. This design led to the *Overview first, zoom, filter, then details-on-demand* mantra.

> CONTROLLED EXPERIMENTS WITH BENCHMARK TASKS SHOWED DRAMATIC SPEED-UPS IN PERFORMANCE AND HIGH SUBJECTIVE SATISFACTION. WE WERE ONTO SOMETHING IMPORTANT.

Today, we still smile when using web-based tools such as Kayak to search for hotels, because the interface closely resembles the HomeFinder.

When we presented those demos at the annual HCIL Symposium, they generated interest from several government agencies. The National Center for Health Statistics asked us to work on the Cancer Atlas data, which had only been published in book form.

In the summer of 1992, Vinit Jain worked with Catherine Plaisant to build dynamic queries on choropleth maps[5] (see Figure 3). Users could animate the map over 20 years of cancer rate data and filter the map by demographic parameters such as income or college education.

We were pleased that Alan MacEachren's team refined and expanded the cancer map into a modern web version, while we steered Ph.D. student Haixia Zhao's research toward making interactive choropleth maps accessible to users with visual disabilities with sonification.

The concept of a generic two-dimensional scattergram with zooming, color coding, and filtering was first applied in the FilmFinder.[6] At the time, our laboratory was supported by IBM to develop a film catalog for interactive TV applications. Catherine was prototyping novel interfaces to rapidly browse movie posters by categories but searching for movies remained a challenge. In a brainstorming session with 8–9 attendees—including Christopher Ahlberg who joined us for a second summer, Ben asked each person to describe a possible interface for finding a film from a library of 10,000 videos. As the variants of traditional approaches with command lines and menus were rejected, Ben proposed a dynamic query approach using a two-dimensional layout with years on the *x*-axis and popularity on the *y*-axis. The idea was quickly accepted and refined.

By the next morning, Christopher had a prototype showing 1500 films with color-coded spots (red for drama, white for action, etc.). As the weeks passed and other students built components, Christopher integrated them into the FilmFinder (see Figure 4). A range slider allowed filtering by the film's length and buttons allowed selection by ratings. A click on one of the spots produced a pop-up box with details of each film and a picture of one of the actors/actresses. The pictures for about 50 actors were grabbed from the net, and Michelle Pfeiffer (Christopher's choice) became the default for the others. Demos were carefully scripted around the pictures that we had.

Christopher Ahlberg returned to Sweden to work on his excellent Ph.D. dissertation on Dynamic Queries at Chalmers University. He developed an enhanced UNIX implementation called Information Visualization and

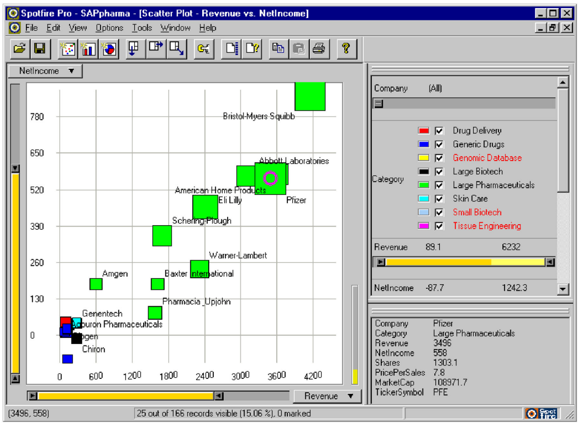

FIGURE 5. Spotfire, first commercial product.

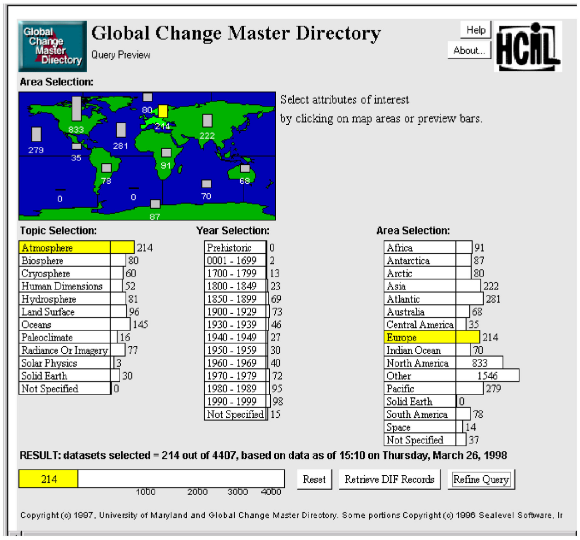

FIGURE 6. Query Preview prototype for NASA. Distributions of the number of available datasets are shown on a map and for three main attributes. Selecting a value updates the distributions and the total count.

Exploration Environment (IVEE) that was flexible in reassigning the axes to other attributes and had a scroll bar to permit large numbers of sliders.

IVEE allowed the importation of arbitrary flat files[6] and, therefore, was applied to many datasets, including data from the State of Maryland's Department of Juvenile Justice (DJJ). We remember vividly the visit of DJJ's team to see the new prototype. They and others had studied their data carefully, using varied statistics, but suddenly the manager saw many incomprehensible flaws in the data. He could not believe we had received the real data. In five minutes of exploration with the help of HCIL researcher Anne Rose, the team had identified a multitude of alarming issues—for example, the large number of juvenile offenders with a reported age under 5—because case workers had entered zero for the age when they did not know the real age of the juvenile. They were literally "seeing" their data for the first time, so they worked intensely with Anne for the rest of the day.

THEY WERE LITERALLY "SEEING" THEIR DATA FOR THE FIRST TIME.

In Sweden, Ahlberg gathered his friends and found venture capital to start a company in mid-1996, named IVEE Development and later renamed Spotfire. This commercial version of the starfield display increased user control, offered greater flexibility, and supported larger databases (see Figure 5).

We enjoyed watching Ahlberg's transformation into a successful business person. Spotfire became a leader in visual data mining and information visualization, focusing at first on the challenging pharmaceutical drug discovery task. Sixteen of the twenty big pharmaceutical companies became major adopters. After growing to 200 employees, the company was acquired by TIBCO in 2007, which continues to develop and market Spotfire (www.tibco.com/products/tibco-spotfire).

QUERY PREVIEWS TO PREVENT ZERO-HIT RESULT SETS

In parallel, we worked with the nearby NASA Goddard Space Flight Center, whose researchers regularly attended the HCIL Symposium. They challenged us to help their scientists find the remote sensing datasets they needed from storage in multiple large repositories spread over the country. These networked repositories were searchable with a form fill-in interface. However, usage statistics revealed that the majority of queries returned nothing. Those "zero-hit queries" were the result of overly restrictive searches or searches for data that was not available. Equipped with our new mantra of *Overview first, zoom and filter, then details-on-demand*, we knew what to do, but the data was too large to use dynamic queries even with Spotfire at the time. Instead, we used a two-step approach, using interactive *query previews* to reduce the dataset size, followed by a traditional dynamic query in the style of Spotfire. Instead of showing individual results within 100 milliseconds, the query preview shows the result set size and distributions over major attributes such as topic, year, and area (see

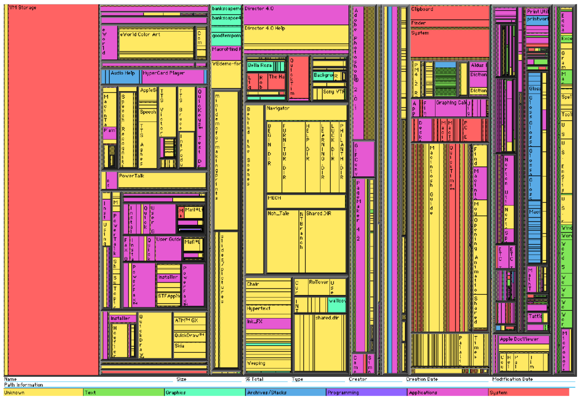

FIGURE 7. Treemap overview of the HCIL lab's hard-disk, color-coded by file type.

Figure 6).[7,8] This eliminated wasted time with empty and overly large result sets. Even the stand-alone query preview was useful to learn what actually was in the NASA archives.

While our prototypes worked only with local files, this approach worked well on the World Wide Web, leading to implementations at NASA and elsewhere. Later, query previews inspired Marti Hearst to develop *faceted search,* which merged query previews with database search.[9] The bars were dropped to save space, but the numerical counts remained. Faceted search is now the user interface found in most e-commerce websites today, allowing customers to quickly find the light bulb, tent, or hiking boots that match their needs.

LIGHTNING IN THE FOREST: TREEMAPS FOR VISUALIZING HIERARCHIES

Treemap was born from a simple need. In 1990, the HCIL's 80 Megabyte hard disk was often full, so Ben became obsessed with the idea of producing a compact screen-filling visualization without scrolling a lengthy directory tree structure. Since the disk was shared by 14 users it was difficult to determine who was using the space. Finding large or duplicated files and old versions that could be deleted, were difficult tasks.

Tree-structured node-link diagrams grew too large to be useful, so we explored ways to show a tree in a space-constrained layout. While puzzling about this in the faculty lounge, Ben had the lightning bolt experience of splitting the screen into rectangles in alternating horizontal and vertical directions as the algorithm traverses down the levels. This recursive algorithm seemed attractive. It took him a few days to convince himself that it would always work and to write a six-line algorithm. Choosing the right name took probably as long, but the term "treemap" described the notion of turning a tree into a planar space-filling map.[10]

Finding an effective visualization strategy took only a few months but producing a working piece of software took over a year. Brian Johnson and David Turo's implementation added features and refinements such as zooming, hue/saturation control, and many options and labeling controls (see Figure 7). They conducted user studies to understand the benefits and limitations of treemaps.

The excitement was great. Like many innovators we thought millions of users would be using it within a few years. Still, training took about 15 mins in order to demonstrate the benefits of treemaps, and we knew that people only got excited when they could see their own data and gain insights from it.

THE EXCITEMENT WAS GREAT. LIKE MANY INNOVATORS WE THOUGHT MILLIONS OF USERS WOULD BE USING IT WITHIN A FEW YEARS.

A German visitor, Alexander Jungmeister worked with Dave Turo to build a stock portfolio visualization that showed client portfolios, industry groups, stocks, and trades. A Japanese visitor, Asahi Toshiyuki, built his own innovative treemap interface to implement the Analytical Hierarchy Process in decision making. Another success story for treemaps was their inclusion in a satellite management system for Hughes Network Systems

Enthusiasm for treemaps grew very slowly, reminding us that many users needed to develop visualization literacy to appreciate how a visual representation could help them with their work. One devoted analyst from a major train company first failed to get our treemap application accepted by the 225 managers in his company. When he refined the application so that the treemap appeared when the managers turned their computers on and was automatically customized to their personal needs, managers became very enthusiastic. Eventually, the University of Maryland's Office of Technology Commercialization signed 17 licenses to corporate users, which supported two undergrad students to work with customers, fix bugs, and make refinements.

During summer 2000, the HCIL resumed work on treemaps, when Raghuveera Chalasani developed a Java version, further expanded by a team of undergrads supported by a grant from Chevron-Texaco. Catherine and Ben Bederson worked with Cheryl Lukehart, who was charged with exploring the potential of information visualization for the company. The grant also supported the work on SpaceTree, giving us the challenge of effectively exploring the enormous organization chart resulting from the merger of Chevron and Texaco.

Since we did not patent treemaps, they thundered out to many organizations and researchers, which developed their own versions, often providing significant improvements. The squarified treemap algorithm (with aspect ratios close to one) proved to be aesthetically appealing and eased the inclusion of text labels, accelerating adoption.

The goal of ordered treemaps was finally achieved by a novel algorithm that nicely balanced square-like nodes while preserving order. The 2001 article, "Ordered Treemap Layouts" was written with Martin Wattenberg and Ben Bederson[11] and won the IEEE Visualization Test of Time Award in 2021.

PEOPLE GOT IT IMMEDIATELY AND FROM THERE THEY COULD IMAGINE A MYRIAD OF OTHER USE FOR TREEMAPS.

Independently, Martin Wattenberg developed a particularly well-designed application that made treemaps a widely used business tool and triggered a rapid spread of treemaps. The MarketMap allowed anyone familiar with the stock market to see how the market was doing at a glance. It only had 2 levels to show the 11 industry sectors and 500 stocks, color-coded red for down and green for up, and a control to change the time range. In this application, no training was required. People got it immediately and from there they could imagine a myriad of other uses for treemaps. A startup company, Hive Group, licensed treemap from the University and developed a commercial version, which remains active as Visual Action (https://www.visualaction.com/).

We were also delighted by the Food & Drug Administration's use of treemaps to show adverse drug event reports. During our first visit to the FDA (to discuss time sequence analysis tools), Catherine was hugged tightly by Ana Szarfman who said "thank you for treemap!" Initially quite puzzled, we later understood what Ana meant when we walked by many FDA offices and saw people looking intensely at treemaps, which they called Sectormaps. The Sectormaps interface helps pharmacovigilance experts monitor thousands of adverse drug reactions and spot higher than expected numbers of reports, organized in a hierarchy of human body systems.

Treemaps eventually spread through most information visualization applications and toolkits. In 2014, we felt the satisfaction of our 20+ year efforts when Microsoft included treemaps in the widely-used Excel spreadsheet. Another success, demonstrating the scalability of treemaps, was made by our long-term collaborator, Jean-Daniel Fekete, who found ways to rapidly display treemaps with a million items.

The artistic side of treemaps led Ben to create a portfolio of treemaps that used real data, with layouts and color palettes influenced by Mondrian, Gene Davis, Hans Hoffman, and others. The twelve 24 by 36 inch prints were exhibited in our hallways. A second printing appeared at the National Academy of Science in Washington, DC, and a third printing is in the collection of the Museum of Modern Art in New York. The full size images are free to download at the Treemap Art Project (http://treemapart.wordpress.com).

LIFELINES FOR PERSONAL HISTORIES

So many of our projects are connected! Remember how the FilmFinder demo attracted interest from the Maryland DJJ? During our early meetings and visits, our partners described how hard it was for case workers to rapidly evaluate the record of a juvenile offender. When a juvenile is brought by the police late at night, deciding whether to send the kid home or to a detention center often has to be made within minutes, and is based on reviewing the youth's record. However, reviewing dozens of screens full of tables and textual reports in the middle of the night can be challenging. On that day, after a few questions and viewing sample data, we remember looking at each other as if we both knew the answer: timelines as overviews! Paper mockups were quickly drawn, first prototypes were developed by a team led by Anne Rose, and, in a flash, Lifelines was born.[12]

Lifelines summarized all the data contained in a youth record on a single screen. A zoomable timeline contained multiple sections called facets, one for cases, one for placements, caseworker assignments, documents, and medical events. For each line segment, we used line thickness and color to map different case attributes, such as thickness for the severity of the accusation and color to show the evolving status of the case (see Figure 8).

As always, the prototypes needed many cycles of user testing and iterative improvements to become

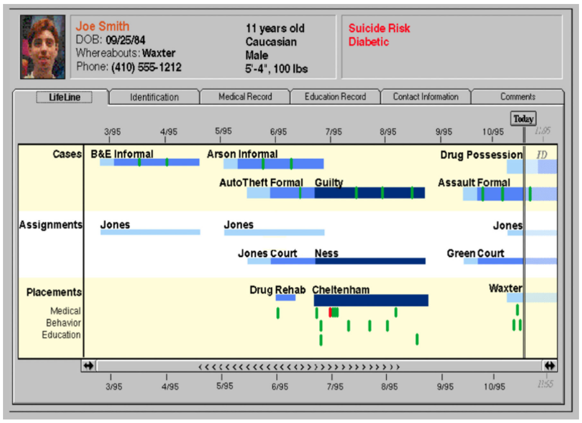

FIGURE 8. Juvenile record summary in Lifelines.

fully effective and caseworkers were pleased with the results (even if—ironically—concerns were voiced about how embarrassing it was to see many delays or cases remaining open when they should have been closed). Grad student Diane Alonso from the Psychology Department helped us conduct a thorough summative evaluation of Lifelines compared to the traditional data tables. The results showed clear benefits in task completion time, error rate, and satisfaction. Reviewers thanked us for confirming those benefits which may now seem obvious but at the time were still debated.

Sadly, the DJJ did not implement Lifelines. It had planned to do so but the year 2000 bug (Y2K) ate all their technology budget. Then management changed, and their attention moved to other priorities.

In the meantime, we clearly saw the potential of Lifelines to be useful to display medical records. Anne Rose worked with student Rina Levy and her husband Warren (a cardiologist) to mockup a quick example that showed diagnoses, medications, tests, interventions, etc. John Karat from IBM saw Catherine's presentation of Lifelines at the CHI'96 conference and immediately found funding so we could work together. At the time, IBM was developing one of the earliest Electronic Health Record systems for Kaiser Permanente in Colorado. We worked intensely with Rich Mushlin of IBM who knew the intricacies of medical data and Aaron Shneider, a practicing internist at Kaiser. We selected realistic sample data and created in-depth prototypes which were presented at the American Medical Informatics Association conference in 1998 (see Figure 9).[13] That article continues to be cited regularly, while the design has inspired dozens of implementations, especially in the medical domain.

Several years later, Dr. Mark Smith—the director of the emergency medicine department of the Washington

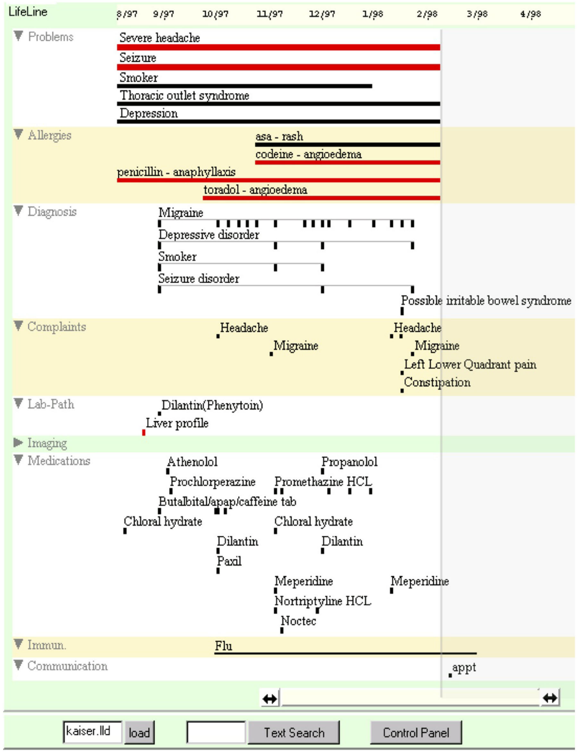

FIGURE 9. Medical health record in Lifelines.

Hospital Center in DC heard Ben talk about the laboratory's research and challenged him to summarize the temporal patterns found in collections of medical records as opposed to a single record. Mark carefully described one particularly dangerous liver failure (contrast induced renal deficiency) that troubled clinicians in the hospital. He wanted to characterize how often the complex temporal pattern of this phenomenon occurred in his hospital. After getting access to sample data accompanied by guidance from medical researchers, David Wang developed the prototype Lifelines2, which could show thousands of patient histories at once. Again and again, Mark Smith challenged us with new requirements. At the same time, he made himself available to discuss the problem in detail, provided de-identified data, and connected us with colleagues who offered related research questions and provided feedback on prototypes.

Mark's next challenge was to ask: Where do patients go after leaving the emergency room, and where does their journey within the hospital take them (ICU? General floor? Death?)—which led to LifeFlow developed by Krist Wongsuphasawat.[14] Later questions from the U.S. Army Pharmacoviligance Center and the National Children's Hospital led Megan Monroe to add support for interval data (for drugs) and a visual search interface to produce EventFlow (https://hcil.umd.edu/eventflow). We came to understand the importance of having meaningful questions, sample data, and feedback from practitioners who cared about answers to their questions.[15] Our remarkable students were able to rapidly build potent prototypes and then refine them based on user needs.

Reflecting back, it became clear that our physician partners (Aaron Shneider at Kaiser Permanente, Mark Smith at Washington Hospital Center, and later Seth Powsner at Yale School of Medicine) have been of tremendous importance to the success of our medical projects.

REFLECTING BACK, IT BECAME CLEAR THAT OUR PHYSICIAN PARTNERS HAVE BEEN OF TREMENDOUS IMPORTANCE TO THE SUCCESS OF OUR MEDICAL PROJECTS.

One of our frequent strategies was to start new students by asking them to help a more senior student and learn from them, and then develop their own unique direction often with the help of another new student who joined our team. Fan Du started by helping Sana Malik, who had helped Megan Monroe, who had helped Krist Wongsuphasawat, who had helped David Wang. Each of them tackled a unique problem, while also contributing to a progressively more polished EventFlow tool, which was eventually made available for licensing by the University's Office of Technology Commercialization.

MORE LIGHTNING AND THUNDER

The HCIL had other researchers working on other projects, some of which were tied to visualization. Chris North developed the first interactive visualization to explore the anatomical images of the National Library of Medicine's Visible Human Project.[16] Ben Bederson arrived in 1998 and became deeply involved with several projects. He dazzled us with his smooth zooming applications based on his remarkable graphical toolkits[17]: Pad++, Jazz, and then Piccolo. Bederson's zoomable user interfaces[18] led to dozens of applications for photograph collections, tree structures, networks, and the widely used International Children's Digital Library. Another application of the Picolo toolkit came in Harry Hochheiser's TimeSearcher to explore time series data in ways that were never possible before.[19] Users could look for spikes, sinks, rising or falling trends, and spot anomalous patterns in stock market, citation, medical, or genomic data.

SPARKS OF INNOVATION

In 1993, ten years after its founding and nearly 30 years ago, the laboratory published a book: *Sparks of Innovation in Human-Computer Interaction*. The title was inspired by a colleague from Carnegie-Mellon University who commented that "we understood the art of fire making, and that the largest benefit to our visitors was learning how to spark fires in their own organizations." As an introduction, Ben wrote about the "Maryland Way,"[1] which he captured in seven succinct points:

1) choose a good driving problem;
2) become immersed in related work;
3) clarify short-term and long-term goals;
4) balance individual and group interests;
5) work hard;
6) communicate with internal and external stakeholders; and
7) get past failures. Celebrate success!

Later, Niklas Elmqvist added:

8) Transfer results to a general audience.

These eight points are a useful starting point, but each project required some tuning to fit our goals, the needs of sponsors, the uniqueness of each student, and our complementary styles. Catherine focused more on work with partners, usability, and validation. Ben focused more on keeping students on track to generate novel research and finish their Ph.D. theses. Ben often pushed for boldness, Catherine for realistic and practical solutions. We discussed progress and plans during weekly meetings with each student or small group.

The collaboration between the two of us was a key part of our careers and produced more than 100 jointly authored articles. Our strong supportive partnership was an asset during difficult times, such as with paper or grant rejections. When overwhelmed, a short meeting was often enough to sort priorities, decide what to drop, review short-term and long-term plans, and proceed with renewed clarity and energy. However, ours was hardly a quiet gentle collaboration. We often had strong, but respectful disagreements, sometimes surprising our students. We hope that they learned to speak up and appreciate diversity as a strength.

CONCLUSION

The eight points in the Maryland Way are good take away messages, but two more thoughts remain important. First, the underlying concepts of direct manipulation were the foundation for most projects, as they steered us to satisfy user needs for comprehensibility, predictability, and controllability. Second, when we are asked what contribution we are most proud of, the answer is clear: Our students! They carried the lightning and thunder out to other academic research groups, industrial laboratories, and government agencies. We make an effort to stay in touch with them, celebrate their accomplishments, and welcome them back for alumni events. Academic life is often about ideas and publications, but for us, the enduring satisfaction is in the relationship with our students.

THE UNDERLYING CONCEPTS OF DIRECT MANIPULATION WERE THE FOUNDATION FOR MOST PROJECTS, AS THEY STEERED US TO SATISFY USER NEEDS FOR COMPREHENSIBILITY, PREDICTABILITY, AND CONTROLLABILITY.

WEB RESOURCES

Videos:
hcil.umd.edu/earlydaysvis

List of Ph.D. students:
hcil.umd.edu/phd-alumni

Related projects' webpages:
Dynamic Queries: www.cs.umd.edu/hcil/spotfire/
Treemaps: www.cs.umd.edu/hcil/treemap-history/
Others projects:
hcil.umd.edu/research-archive/#visualization

ACKNOWLEDGMENTS

We feel fortunate to have had such a productive and happy research partnership over 35+ years. There is so much more to say, but we want to express our deeply felt thanks to all our sponsors, research collaborators, and the faculty, staff, and students of the HCIL for their diverse and creative contributions. We also thank our VIS colleagues who pushed us along the way with constructive criticism, challenging questions, and enthusiastic encouragement.

REFERENCES

1. B. Shneiderman, "Supporting the process of innovation: The Maryland way," in *Sparks of Innovation in Human-Computer Interaction*. Norwood, NJ, USA: Ablex, pp. 1–10, 1993. [Online]. Available: http://www.cs.umd.edu/hcil/pubs/books/maryland-way.shtml
2. B. Shneiderman, "Direct manipulation: A step beyond programming languages," *Computer*, vol. 16, no. 8, pp. 57–69, Aug. 1983.
3. D. Brethauer, C. Plaisant, R. Potter, and B. Shneiderman, "Three evaluations of museum installations of a hypertext system," *J. Amer. Soc. Inf. Sci.*, vol. 40, pp. 172–182, 1989.
4. B. Shneiderman, "Dynamic queries for visual information seeking," *IEEE Softw.*, vol. 11, no. 6, pp. 70–77, Nov. 1994.
5. C. Plaisant, "Facilitating data exploration: Dynamic queries on a health statistics map," in *Proc. Annu. Meeting Amer. Stat. Assoc. Govern. Statist. Section*, 1993, pp. 18–23.
6. C. Ahlberg and B. Shneiderman, "Visual information seeking tight coupling of dynamic query filters with starfield displays," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, 1994, pp. 313–317.
7. C. Plaisant, B. Shneiderman, B. K. Doan, and T. Bruns, "Interface and data architecture for query preview in networked information systems," *ACM Trans. Inf. Syst.*, vol. 17, no. 3, pp. 320–341, 1999.
8. S. Greene, G. Marchionini, C. Plaisant, and B. Shneiderman, "Previews and overviews in digital libraries: Designing surrogates to support visual information-seeking," *J. Amer. Soc. Inf. Sci.*, vol. 51, no. 3, pp. 380–393, Mar. 2000.
9. K. P. Yee, K. Swearingen, K. Li, and M. Hearst, "Faceted metadata for image search and browsing," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, 2003, pp. 401–408.
10. B. Shneiderman, "Tree visualization with tree-maps: A 2-dimensional space filling approach," *ACM Trans. Graph.*, vol. 11, no. 1, pp. 92–99, 1992.
11. B. Bederson, B. Shneiderman, and M. Wattenberg, "Ordered and quantum treemaps: Making effective use of 2D space to display hierarchies," *ACM Trans. Graphics*, vol. 21, no. 4, pp. 833–854, 2002.
12. C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman, "LifeLines: Visualizing personal histories," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, 1996, pp. 221–227.
13. C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, and B. Shneiderman, "LifeLines: Using visualization to enhance navigation and analysis of patient records," in *Proc. Amer. Med. Inf. Assoc. Annu. Fall Symp.*, 1998, pp. 76–80.

14. K. Wongsuphasawat, J. A. G. Gómez, C. Plaisant, T. D. Wang, M. Taieb-Maimon, and B. Shneiderman, "LifeFlow: Visualizing an overview of event sequences," in *Proc. SIGCHI Annu. Conf. Hum. Factors Comput. Syst.*, 2011, pp. 1747–1756.
15. B. Shneiderman, *The New ABCs of Research: Achieving Breakthrough Collaborations*. Oxford, UK: Oxford Univ. Press, 2016.
16. C. North, B. Shneiderman, and C. Plaisant, "User controlled overviews of an image library: A case study of the visible human," in *Proc. first ACM Int. Conf. Digit. Libraries*, 1996, pp. 74–82.
17. B. B. Bederson, J. Grosjean, and J. Meyer, "Toolkit design for interactive structured graphics," *IEEE Trans. Softw. Eng.*, vol. 30, no. 8, pp. 535–546, Aug. 2004.
18. B. B. Bederson, "The promise of zoomable user interfaces," *Behav. Inf. Technol.*, vol. 30, no. 6, pp. 853–866, 2011.
19. H. Hochheiser and B. Shneiderman, "Dynamic query tools for time series data sets, timebox widgets for interactive exploration," *Inf. Vis.*, vol. 3, no. 1, pp. 1–18, 2004.

CATHERINE PLAISANT is research scientist emerita at the Institute for Advanced Computer Studies, University of Maryland, College Park, MD, 20742, USA. Her work spans the interface development lifecycle, with contributions to requirements gathering, interface design, and evaluation. She received the IEEE VGTC Visualization Career Award in 2020. She is the corresponding author of this article. Contact her at plaisant@umd.edu.

BEN SHNEIDERMAN is an emeritus distinguished university professor in the Department of Computer Science, founding director (1983–2000) of the Human-Computer Interaction Laboratory, and a member of the UM Institute for Advanced Computer Studies (UMIACS) at the University of Maryland, College Park, MD, 20742, USA. He is a Fellow of the AAAS, ACM, IEEE, NAI, and the Visualization Academy and a member of the U.S. National Academy of Engineering. Contact him at ben@cs.umd.edu.

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