Preface

Information visualization has become a major theme during the past ten years for much of the work of the faculty, staff, and students at the University of Maryland's Human–Computer Interaction Lab (HCIL). Our roots are in human–computer interaction (HCI), but increasingly our attention has focused on information visualization. The reason is clear: The overall theme of our work is to improve the experience of people using computers, making that experience more effective and enjoyable. In order to reach that goal, we must create designs that enable users to develop control over the computer and, we hope, to attain a sense of mastery. For this to occur, users must have a fluid and efficient interaction with the computer—and the high bandwidth of visual interfaces is the surest way attain this goal.

A large component of human perception of the world is through sight. As we've said before, *the eyes have it*. There is simply more bandwidth and processing power for input through the human eyes than through any other sensory modality. Sound, touch, and smell are important, but the HCIL's researchers have repeatedly returned to highly visual and dynamic displays as the best way to solve a surprisingly broad set of problems in an equally broad set of domains.

The HCIL is not alone in this belief. The field of information visualization has been maturing, along with its related discipline of scientific visualization. Both have strong ties to the broader field of human–computer interaction, as well as to graphics groups such as the ACM Special Interest Group on Graphics (SIGGRAPH) and many international graphics associations. Information visualization has grown rapidly since 1995, with annual conferences, organized in the United States by the IEEE and in the United Kingdom by the International Conference on Information Visualization.

An important distinction must be made between the more mature field of scientific visualization and the relative newcomer, information visualization. There are certainly overlaps, but scientific visualization researchers deal primarily with three-dimensional physical objects and processes such as blood flowing through heart valves, tornado formation, crystal growth, protein structures, and oil reservoir shapes. They focus on volumes and surfaces, studying formations and flows and asking questions about inside/outside, above/below, or left/right.

By contrast, information visualization researchers are concerned with abstract phenomena for which there may not be a natural physical reality, such as stock market movements, social relationships, gene expression levels, manufacturing production monitoring, survey data from political polls, or supermarket purchases. While both kinds of data sets come from the physical world, instead of dealing with three-dimensional aspects, the users of information visualization tools are interested in finding relationships among variables; discovering similar items; and identifying patterns such as clusters, outliers, and gaps.

Another important discriminator is that scientific visualization users are primarily interested in continuous variables such as density, temperature, or pressure, whereas information visualization users deal with continuous as well as categorical variables, such as gender, race, home ownership, date of birth, state name, and number of bedrooms. Another distinctive feature of information visualization is its attention to discrete structures such as trees and graphs. Of course, there are areas where ideas and applications cross over, but the distinctive aspects of information visualization are important to understand.

The interactive nature of information visualization stems from the use of powerful widgets that enable users to explore patterns, test hypotheses, discover exceptions, and explain what they find to others. Interacting with the data set gives users the chance to rapidly gain an overview, explore subsets, or probe for extreme values. Information visualization tools become telescopes and microscopes that allow users to see phenomena that were previously hidden.

A steadily growing set of books is helping to define the field and support graduate courses in many universities. The classic book by Bertin, *Semiology of Graphics* (1983), has inspired many researchers, while the more recent *Readings in Information Visualization: Using Vision to Think* (Card, Mackinlay, and Shneiderman, 1999) has stimulated numerous developers. The latter book includes 47 early papers from diverse sources with integrative commentaries and an extensive bibliography. Other books on visualization include the fine surveys by Robert Spence (2000), Colin Ware (2000), and Chaomei Chen (1999). Conference proceedings are important resources, and collections of papers on special topics are common in this discipline. Journals devoted to the topic such as *Information Visualization* (*www.palgrave-journals.com/ivs/*) serve to present current research. Guides for practitioners are beginning to emerge (Westphal and Blaxton, 1999).

The broader field of HCI is now also firmly established with major groups such as ACM's Special Interest Group on Computer Human Interaction (SIGCHI), Usability Professionals Association (UPA), British Computer Society Human Computer Interaction Group, and the Association Français pour l'Interaction Homme–Machine (AFIHM). These organizations have significant participation not only from academic researchers, but also from companies and governmental organizations. In fact, the premiere HCI conference (SIGCHI) typically has attendance split equally between researchers and practitioners.

THANKS TO THE HCIL COMMUNITY

On this 20th anniversary of the HCIL, we proudly bring together this collection of work from our colleagues—students, staff, visitors, and collaborators around the world. We hope that by offering this work, along with our reflections on what was important and why, and how the research unfolded as it did, we can shed some light on the often mysterious process of innovation and creation—and encourage others to further advance the field of information visualization.

We are deeply indebted to all of our colleagues at the University of Maryland and around the world. The rich intellectual atmosphere and warm, personable climate in which we work has made this book possible. Our close faculty colleagues, François Guimbretière in Computer Science, Kent Norman in Psychology, Allison Druin, and Doug Oard in Information Studies are true partners in the HCIL. Gary Marchionini was an important participant for many years before he moved to University of North Carolina. The HCIL's long-term research scientist, Catherine Plaisant, has been a key developer of many of the ideas in this book. Other contributors have been yearlong visitors and postdoctoral researchers such as Richard Biegel, Khoa Doan, Richard Salter, and Jean-Daniel Fekete.

Ten years ago, THE HCIL published a book called *Sparks of Innovation in Human–Computer Interaction* (1993), containing a selection of the lab's work from its first ten years. This decade, on the other hand, has been so fruitful that we have decided to focus on information visualization, leaving our colleagues to publish other specialized books, such as the collections on children's technology by Allison Druin (1999, 2000).

The work in this book could have been done only with the participation of a terrific staff. Research assistants, like Anne Rose, have stuck with us through thick and thin, contributing broadly to our research. Our newer staff, Aaron Clamage, Allison Farber, Jesse Grosjean, Trina Harris, and Sabrina Liao, have already made their marks, and we are excited about their joining us.

The rhythm of our work is tied to the seasons—fall, spring, and summer semesters. A new semester is a chance to learn from the past, and each affords an opportunity to start fresh. We have had extremely strong computer science doctoral students who have gone on to make notable accomplishments of their own. Master's students and undergraduates have also made important contributions and are co-authors on many papers. Our students, present and past, are what make the HCIL such a dynamic place.

Benjamin B. Bederson
bederson@cs.umd.edu

Ben Shneiderman
ben@cs.umd.edu

College Park, Maryland
January, 2003

Introduction

Our goals in writing and editing this book were to give researchers and students an understanding of how ideas in information visualization evolve and spread. The dissemination of ideas is a fascinating and instructive process, especially if it involves your original ideas. It is a thrill to see someone adopt your ideas or software, often refining them substantially as they apply them to some novel domain. Fortunately, we have had very positive collaborative experiences. We have repeatedly found that by being open with our own ideas and honest about the source of others', a great harmony results—with innovation, excitement, and rapid improvement.

THE IMPORTANCE OF FLOW

Another important theme that has pervaded our work is something that has come to be called *flow* (Mihaly Csikszentmihalyi, 1990). Over the years we have developed an intuition about what makes information visualization (and other) interfaces work well, and we have discovered congruence between these ideas and the concept of flow, an idea from the psychology literature. Though we don't have a strict formula for a successful interface, we know that a few basic approaches do help. In general, we believe it is important that the users stay in control and that the computer offers choices with appropriate feedback for user actions. Conversely, computer-controlled interactions often lead to unpredictable, and therefore unacceptable, interfaces.

We also have learned that people are primarily interested in focusing on their task, and not on operating the interface—and yet so much of a user's experience with a computer is manipulating widgets, resizing windows, and selecting from menus. It is crucial that computers give users prompt and informative feedback at every step along the way. Finally, users must stay engaged in the task for their experience to be effective in the long run. This means that the interface must not be too complex or confusing as to alienate users, nor so simplistic or condescending as to make them bored.

A computer interface that strikes the right balance can enable users to concentrate on the task at hand. The computer becomes a "tool" in the best sense of the word—an extension of the user's body. Time passes quickly, and the users develop a sense of control and confidence while making progress toward their goals.

When people experience this kind of focus, they sometimes refer to "being in the flow." Some psychologists refer to this as *optimal experience*, a shorthand that describes the best experience that one can hope for.* And though it may first seem far afield from computer work in information visualization, as researchers let us consider it our ideal: to create computer interfaces that enable users to forget they are using a computer and think only of the important work they are accomplishing.

This book is about that process in innovation during the last ten of the lab's twenty years as we concentrated on the field of information visualization, a subfield of HCI.

EVALUATING OUR WORK

How do we assess our progress? Are we any nearer to our goal of creating interfaces that support *flow*? Tough internal assessments—critiquing each other, challenging assumptions, and demanding evidence help prevent us from falling into traps of wishful thinking.

External reviews from colleagues add to our continuing assessments. We send drafts of papers to colleagues, invite visitors to see our work in progress, and engage potential users to try our software. We appreciate good feedback, taking to heart constructive comments that push us to refine our work.

The next level of assessment comes from anonymous reviewers of conference papers, journal articles, and grant proposals. We discuss rejections and try to learn from them. Even when we disagree with reviewers, we try to examine how we might have told the story more effectively. Some of our strongest papers have been shaped by tough reviewing processes.

Published papers are the clearest signs of our progress—they have been validated by peer review, and they are publicly available. Several members of our group appear high on the list of authors ordered by frequency of publication in HCI papers, conference presentations, and books (*www.hcibib.org/authors.html*).

Another imperfect but useful metric of success is the number of references to a paper. The NEC Citeseer (*citeseer.nj.nec.com/directory.html*) index has a special section on human–computer interaction, and we are proud that the most cited paper for many years has been one of our works on information visualization. Similarly, when PARC researchers studied reference patterns in information visualization, they found that their group's papers were cited most frequently, but our lab came in second.

Citations in academic papers are one manifestation of our influence, but downloading our software is also quite validating. More than 30,000 individuals have downloaded PhotoMesa, one of our image browsers (Chapter 2).

Our software also influences commercial and government applications. Many of our ideas have become part of larger success stories, such as SmartMoney's MarketMap (*www.smartmoney.com/marketmap*), the Hive Group's treemaps (*www.hivegroup.com*) (Chapter 6), and Spotfire's (*www.spot-*

**One popular book called Flow by Mihaly Csikszentmihalyi (1990) summarizes the body of work in understanding human optimal experience.*

fire.com) visualization tools (Chapter 1). Contributions to important national and international projects include the Visible Human Explorer for the U.S. National Library of Medicine, NASA's Earth Science Information Partnerships, and the Library of Congress's American Memory Web site.

Finally, an important internal measure of the HCIL's success is the frequency with which our students graduate and join companies, universities, or government agencies where they make valuable contributions. It's especially satisfying to see young, often shy or quiet students become self-confident professionals who are valued by employers and respected by colleagues. As a community, we are gratified when former students return to tell us how much their time at the the HCIL influenced them both professionally and personally.

WORKING WITHIN A BROADER COMMUNITY OF SCIENTISTS

It is difficult to rank or even list all the people in our professional networks, so we must begin with an apology to anyone we have left out in this discussion. Those who want completeness can examine the hundreds of references in our papers. However, we cannot honestly review our work without reflecting on the influences of our colleagues. In several sections, we include more details related to that topic, but this opening mentions a few of the major groups whose influence cuts across many of the sections.

Our strongest and most enduring bonds have been with the community of researchers at the Xerox Palo Alto Research Center, now simply PARC (*www.parc.com*). Our contacts have been mostly with the user interface group and its related teams, especially the manager and long-term researcher, Stu Card. Stu is a leader in theory-driven thinking and research and is a remarkable innovator, as testified by his numerous patents and papers. The PARC group has also included key people such as Jock Mackinlay, George Robertson (now at Microsoft Research), Ramana Rao (now at Inxight), Peter Pirolli, Mark Stefik, and many others.

As Microsoft Research grew, we enlarged our contacts with George Robertson, Mary Czerwinski, and others. Other industry groups include those at AT&T-Bell Labs and spinoff groups, including Stephen Eick, Andreas Buja, and Stephen North. We have also enjoyed long-running interactions with Clare-Marie Karat and John Karat at IBM. Special mention goes to Nahum Gershon of Mitre, who has been an effective champion and organizer for information visualization conferences and journals.

University colleagues include Steven Roth at Carnegie-Mellon University, Steve Feiner at Columbia University, George Furnas at University of Michigan, John Stasko and Jim Foley at Georgia Tech, Andries Van Dam at Brown University, Jim Hollan now at UCSD, Saul Greenberg at University of Calgary, Robert Spence at Imperial College, Keith Andrew at the University of Graz, and Alfred Inselberg at Tel Aviv University. Another special mention goes to Edward Tufte at Yale University, who is well known for his independently published books (1983, 1990, 1997) and for his annual public lecture tour—we regularly pay for students to attend when he swings through the Washington, D.C. area.

There are many others, but these people form the core of our community. We jointly write books and articles, organize conferences, and participate in workshops—all to promote information visualization to broader circles. Seeing each other for a beer or dinner once a year is important, and the continuity of contact is maintained by email. We tell our latest stories, probe for their new ideas, and seek each other's respect. These colleagues are who we turn to validate our innovations, to ask for reviews of our draft papers, and to be our partners in proposals.