Technical Forum

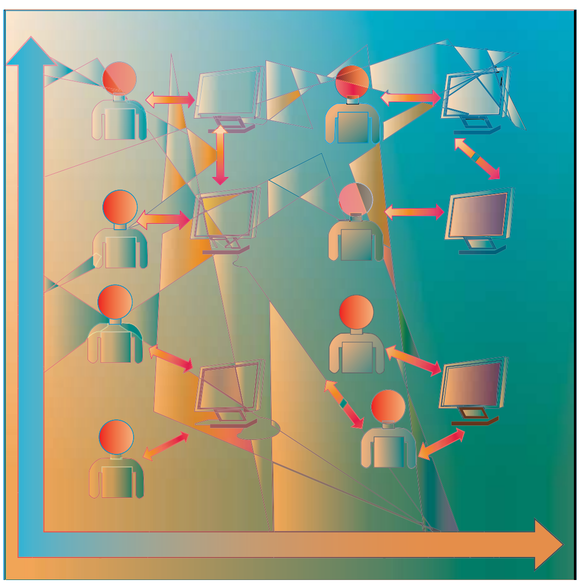

COMMUNICATION-MINDED VISUALIZATION: A CALL TO ACTION

Visualization applications can enable interactions between people in powerful and unexpected ways, as illustrated by the following two personal experiences.

In the spring of 2003, the first author created PostHistory, an application to visualize e-mail archives of individuals.[1] Given the personal nature of the data, it was assumed the archive owner would view the data alone. Indeed, during a usability study, the experimenters carefully explained to participants that no one other than the owner would have access to the visualizations. Yet, as soon as users had access to the application, they began finding ways to share the resulting images. Users mailed screen captures to friends and family; they invited colleagues to sit with them to view the screen images together. This sharing triggered deep reminiscing and long conversations about events in their lives, which the users considered an important benefit of the visualization system.

In the winter of 2005, the second author created NameVoyager, a Web-based visualization of historical data on baby name popularity.[2] NameVoyager helps expectant parents find names for their babies and encourages individual exploration of the baby name data. As it turned out, the data was explored, but not just in isolation. Thousands of visitors to the site engaged in conversations about their findings, using discussion forums and blog comments, collectively identifying trends and anomalies and forming conjectures about the data.

Although not specifically designed for communication, the preceding applications created rich opportunities for users to engage in discussions about the data being displayed. Inspired by these experiences, in this paper we introduce the concept of communication-minded visualization (CMV), a visualization designed to support communication and collaborative analysis. Our emphasis is on the design of the user experience rather than the technical implementation challenges.

We believe that designing for communication is essential because users do not interact with visualizations solely to gain personal insights. An insight that matters usually has to be communicated to others.[3] As Johnson et al. point out,[4] visualization plays an important role in many disciplines, such as biology, physics, and genetics. To harness the power of visualization as a working tool for multidisciplinary teams, designers need to pay close attention to how visualization affects and enables the communication of discoveries and the discussion of ideas within multiple contexts.

As in the preceding examples, communication of visualization findings can take place in a variety of ways, ranging from the pervasive screen capture to elaborate narrated videos. Also ubiquitous is the practice of leaning over someone's shoulder to see what is happening on his or her monitor. It is not uncommon to have up to six viewers looking at the same visualization screen as one person interacts with the data.[5] In business meetings screen captures or videos are projected on a large screen. In presentations to professional conferences video sequences have become more common as a way of making interaction and transition techniques easily understandable to viewers. Finally, printouts are used to share analysis and findings.

The process involved in sharing the visualization data is often cumbersome, including screen captures, an image-editing program, and an e-mail program. There seems to be a gap between the visualization application and the sharing process. Current visualization platforms lack support for communicating a user's findings.

Aside from the fact that communication and sharing capabilities are often external to visualization systems, most ad hoc sharing practices suffer from other drawbacks as well. They often rely on interactions that are not very effective in screen capture or printed form. For example, many popular visualizations, such as Map of the Market from SmartMoney,[6] use tooltips—small windows that contain explanatory text when the mouse moves over a target—which are lost in screen capture form. For three-dimensional visual applications, removing the motion element from the user interface means that the viewer loses one of the strongest depth cues.[7] Videos of an interactive computer session can be hard to follow if the viewer has no advance warning regarding the part of the screen where the keyboard or mouse actions take place. Aside from basic legibility problems of screen captures and printouts, an inability to interact with the application (which applies to canned video as well) may reduce the credibility of an analysis. As a result, ad hoc sharing of noninteractive versions of a visualization is not a satisfactory solution.

Although visualization-driven communication abounds in the real world and although some commercial products have started to explore CMV-style interaction, capturing and communicating visualization interaction and discovery processes have received little attention from the research community. In this paper we propose a conceptual framework within which to pursue CMV issues, and we hope this framework will help ground inquiry in this area as well as encourage the emergence of a community of interest. We lay out the range of issues in the area, associate this topic to related research areas, and provide initial guidelines for CMV design and evaluation.

The rest of this paper is organized as follows. First, we describe a number of commercial and experimental visualization systems that address various communication needs. Second, we highlight established research areas whose concerns are related to those in CMV. We point out the relevant topics in these areas and how these issues emerge in CMV. Finally, we outline proposals for the design and evaluation of CMV applications.

EXISTING VISUALIZATION SYSTEMS

Designers of visualization systems have not completely ignored the role of communication and group sharing. We describe here a number of commercial and experimental visualization systems that have been designed with communication in mind.

One such system is CoMotion**, a commercially available product from Maya Viz, LLC. that allows users to synchronously and remotely jointly perform visual data analysis tasks. In CoMotion each user opens a window that provides a common view of the visualization target. Users take turns interacting with the data in the shared view, chatting by means of instant messaging.

The CoMotion architecture led to Command Post of the Future, an application for the United States military that allows the members of a command unit to share information through a collaborative visualization application. All users are located in a command room in which a large visualization screen is the focal point. In addition, users have individual machines running a copy of the visualization application, and the information they generate—manipulations and annotations of maps—shows up at once on all other users' screens. In 2004, the system was deployed in the field, and military personnel credit the application with providing troops with the highest level of insight and situational awareness they had ever experienced.[8]

Since the mid-1990s, several research projects have explored synchronous remote sharing of scientific visualizations under the rubric of "collaborative visualization."[9] Collaborative visualization systems have become important data exploration tools in a range of scientific fields from medical diagnosis[10] to archaeological excavations.[11] The concerns of this field have primarily been related to the technical problem of faithfully replicating one user's experience for another at a different network location. Brodlie et al.[12] provide an excellent survey of the state of the art.

Visualization sharing can also happen asynchronously. DecisionSite** Posters from Spotfire, Inc. has been designed specifically to support asynchronous sharing of visualizations. The application is a Web-based client that allows users to capture interactive snapshots of analyses and pass them as posters to a co-worker, who in turn can refine the analysis. Users can make notes and set visualization "bookmarks" (pointers to a specific state of the visualization). The notes have associated threads and allow any researcher to see comments made by others. DecisionSite Posters can also be sent using regular e-mail; a recipient of a poster may then view and interact with the poster with a Web browser and even follow the sequence of steps taken by the sender.

DecisionSite Posters was launched in January of 2002 and has seen a slow but steady rate of adoption. According to the company, the product was created in response to customer interest in sharing and collaboration.[5] So far, the communication capabilities in DecisionSite Posters have been used in an unexpected way. Instead of engaging in deeply nested threaded conversations by using the conversation panel, as envisioned by the designers of the system, users have largely used the tool just for presenting their findings to colleagues. The ability to create commentary associated with pointers into the visualization provides an easy way to choreograph a step-by-step presentation. Having such paths coupled with the full-fledged visualization makes it easy for viewers to take advantage of the directed view of the data and at the same time break off, when desired, to freely explore the visualization.

A final example that touches on the question of collocated communication is PhotoMesa,[13] an experimental image browser. Inspired by the desire to have his two-year-old daughter watch him browse without getting lost, Ben Bederson, the creator of the application, set a design goal that all viewers should be able to easily follow what the person interacting with the tool is doing. To ensure that viewers follow the sequence of navigational commands, PhotoMesa employs zooming and highlighting to call attention to actions involving the mouse and other I/O devices.

The visualization systems just described illustrate different communication scenarios. Whereas DecisionSite Posters is designed for asynchronous, remote communication of visualization findings, PhotoMesa is intended for live, collocated sharing of visual data. CoMotion, on the other hand, is an application for synchronous communication between remote users. Though far from exhaustive, these examples begin to show the variety of possible CMV scenarios. In the next section we anchor this diversity of communication settings under a unifying conceptual framework.

THEORETICAL BACKGROUND

CMV poses fundamentally new problems for visualization designers. A CMV application has all the legibility, perception, and layout challenges of any visualization application, in addition to a host of new difficulties related to group communication practices. Of course these collaboration challenges are not limited to visualization applications. Group communication difficulties have been thoroughly investigated in established research areas such as Computer Supported Cooperative Work (CSCW). Although CMV raises a series of questions that are

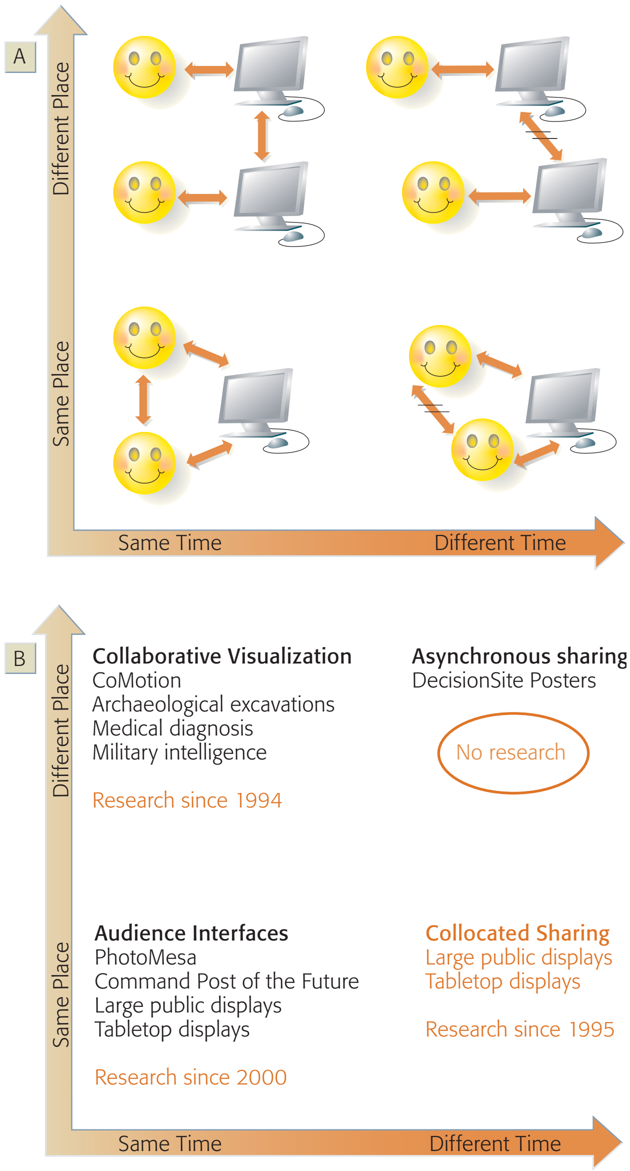

Figure 1
Johansen's time-space matrix applied to CMV: (A) possible CMV scenarios; (B) visualization projects that support collaboration between users

specific to visualizations, it is beneficial to look at the vast CSCW literature for guidance on collaboration/communication issues. In this section we discuss a set of principles that are relevant to the topic of visualization as a communication artifact.

CMV design space

What is the structure of the CMV design space? A good structure should ideally produce an organized view of existing systems while highlighting areas that have been underexplored. We believe Johansen's well-known time-space matrix[12,14] does both (see ***Figure 1A***).

The matrix posits two independent dimensions on which systems may differ: space and time. In some cases people use the system synchronously (at the same time), and in other cases they communicate asynchronously (at different times). Orthogonal to time is space; communication can take place at a single location (same place) or at several different places. The time-space matrix in ***Figure 1B*** lists experimental and commercial products that have been designed with communication capabilities in mind; they populate the four different quadrants produced by the two dimensions. The matrix is a helpful way to organize this diverse set of systems.

Perhaps more important, Figure 1A makes clear that a relatively neglected area exists. Synchronous collaborative scenarios have been much more widely explored than asynchronous visualization-based communication. Except for DecisionSite Posters, virtually no work has been done in this area in either research or industry. We believe that asynchronous communication of visualization-driven insights is a key aspect of CMV and an important research topic. As global usage of e-mail attests, lightweight, asynchronous communication is a ubiquitous and powerful flavor of computer-mediated communication. By not fully exploring asynchronous communication of visualization discoveries and processes, the research community is missing an opportunity to make important contributions to visualization research.

Increasing adoption of CMV applications

There is extensive literature in the field of CSCW regarding the obstacles encountered when new groupware applications are introduced in organizations. Such work raises practical concerns for technology in real-world scenarios, ranging from the social aspects of communication processes to the political agendas of various actors in a group. An early and influential survey of these concerns was published by Grudin.[15] The survey lists eight challenges for developers of CSCW applications, with a careful eye to the different stakeholders of large-scale collaborative systems (purchasers of such systems, management, and users). By analyzing which groups of stakeholders benefit from the deployment of CSCW systems, Grudin concludes that there will always be disparities in the accruing benefits for different groups of users (for instance, between those who use the system to get work done and those who use the system to monitor others' work). Unless the systems address and minimize these disparities, they are doomed to fail. In order to encourage the adoption of CMV systems, it is essential to understand the social dynamics involving collaborative systems.

Since Grudin's survey, collaboration technologies have evolved considerably, creating new social situations and raising additional and unexpected challenges for adoption. The following list of guidelines is partially based on Grudin's survey and also includes more recent concerns involving collaborative systems.

1. *Balancing effort and benefit*—CMV should not significantly increase the level of additional work for individuals who do not perceive a direct benefit from the use of the application.

   Example: There are likely to be different kinds of users that interact with a CMV program. For instance, some users may take full advantage of the visual analysis and communication capabilities of the application, whereas others may use the program simply for information viewing (which is a task that could easily be accomplished by using some other graphical program with which the user is already familiar). It is important that there be no additional work for the viewer-type user in order to use the CMV application. Spotfire's DecisionSite Posters provides a good example of this principle because sharing a particular visualization does not require a viewer-type user to install any special software.

2. *Critical-mass adoption*—The benefits of using CMV has to be sufficiently clear in order to enlist critical-mass adoption. Note that since Grudin's survey the World Wide Web has exploded in popularity, providing an important avenue for amplifying usage once critical mass has been reached.

   Example: NameVoyager provides a clear benefit to expectant parents. It has been suggested that other types of users saw benefits as well: a chance to socialize, explore, or just cause trouble.[2] Once an initial core of bloggers—the critical mass—began discussing the Web site, usage soared. The NameVoyager experience suggests that it is worth taking into account the various types of users and that the Web can be a powerful tool for increasing usage of the CMV application.

   In *The Social Life of Information*, Brown and Duguid argue that much of the value of real-world processes and tools is derived from the social context that surrounds them.[16] As the NameVoyager and PostHistory[1] experiences suggest, data visualization has the potential to generate contextual social activity. As described in Reference 2, the discussions surrounding NameVoyager contain a huge amount of detailed analysis that greatly enhances the value of the data contained in the visualization itself.

3. *Support for social processes*—As much as possible CMV should complement established social processes so as not to disrupt existing practices and discourage users who are crucial to its success. In other words, because features that support group processes may be used relatively infrequently, CMV should be seamlessly integrated with heavily used features.

   Example: The emergence of CMV does not mean that users will stop conducting visual analysis on their own and start exploring visualizations in groups all the time. CMV should be seen as an additional option for doing data analysis. As such, it is crucial that these applications be well integrated into the environments in which users now carry out individually most of their data exploration.

   Specifically, if e-mail is the main way in which a group of people communicates, it makes sense for the visualization application to be integrated with e-mail at some level—the entire application need not work *within* e-mail, but e-mail should be part of how users communicate about the visualization without having to switch contexts. Few visualization applications work well in the context of e-mail, despite the vast amount of collaboration that occurs in that medium. Users of visualization tools are often reduced to sending screen captures or even elaborate textual directions for what to look at. Exceptions are current Web-based mapping programs, such as Google Maps[17] or MapQuest.[18] A combination of location and zoom level can be bookmarked through an ordinary URL, which is easily sent by e-mail or instant messaging.

The fact that DecisionSite Posters requires that users move from their standard desktop analysis tools to a special, stripped-down Web client may be one of the reasons why it is not being used more frequently for collective data analysis. Given this lack of integration, it is perhaps not surprising that the software is currently used mainly for presentation purposes.

4. *Privacy*—Data sharing leads to diminished privacy. As collaborative systems become more powerful and data streaming more seamless, excessive exposure becomes a problem. By their very nature, visualization systems reveal patterns and connections in large data sets that might not have been easily perceivable before. Consider the privacy concerns that arise when a revealing visualization is shared among dozens, maybe hundreds of other viewers. CMV systems should give users nuanced control of how much of their data sets they might like to share with others. Such control could affect both the data segments that are shared and the ways in which data are presented (e.g., the ability to make names anonymous, to aggregate data points to higher levels of granularity, etc.).

   Example: Erickson et al.[19] argue for social translucence, suggesting that total transparency is undesirable in some situations. In the case of visualization, there are several scenarios where it could be better to show less information rather than more. For instance, a user may wish to share a visualization of an e-mail in-box, but with the proviso that all names are made anonymous. In the realm of public data, a user study conducted by Viégas and Smith[20] shows that even visualizations of publicly available data such as Usenet conversations have the potential to raise privacy issues. It may be desirable to implement "privacy filters" that protect sensitive data. It may also be important to create "audience filters" that gear different views of a visualization to different kinds of audiences.

Information foraging theory for groups

Finally, it is worth asking how standard theories of the value of information visualization may apply to group usage. In this section we consider a well-known framework for analyzing visualizations and other information interfaces, *information foraging theory*.[21] Based on an analogy with strategies used by animals foraging for food, this theory describes how people navigate when searching for information. Central concepts in the theory are the idea of an information patch (a small, easily explored subset of a large collection of data) and that of information scent—hints that a subset of data might be a valuable or relevant information patch. In the foraging model, a user alternates between searching for information patches (found by their "scent"), and exploiting patches (typically reading a document or examining particular data points in detail).

> ■ In spite of advances in data visualization techniques, not enough has been done to support visualization-based communication ■

How might this model apply to an information search performed by a group rather than an individual? Several investigators have studied how group information foraging may benefit an individual. Collaborative filtering can be described as a group foraging activity, and Reference 22 advocates creating "history-enriched" objects so that individuals may learn from the previous searches of others. Neither of these models, however, addresses the question of how a group may successfully forage for data.

Here we provide a very brief sketch of a theory of information foraging for groups and how it may support CMV. Imagine that many people are analyzing a data set that contains an important undiscovered pattern. What is the optimal exploration strategy for the group? Assuming that communication between people is efficient, the best strategy is to avoid redundant searching; that is, to follow an approach in which individuals explore distinct parts of the data set in parallel. To maximize the speed of discovery, the searchers should look first at patches with a high information scent. When a discovery is made by one person, it can then be quickly communicated to the rest of the group.

In a tightly controlled organization, this process may be formalized, with people explicitly dividing the work of inspecting a data set—consider a group of lawyers sifting through a box of subpoenaed documents. In many other contexts, however, individuals may wish to communicate their discoveries, but may not be able or willing to coordinate the details of their searching behavior. In such a situation, the problem is to find a system in which individuals are drawn to promising patches (to avoid wasteful searches) while making sure that not everyone searches the same patch (to avoid redundancy).

Visualizations may be an excellent solution to this problem. First, visualizations are good at providing information scent: an outlier in a scatter plot or a rectangle with an unusually bright color in a treemap (a visualization of hierarchical data) give hints that drilling down may be worthwhile. At the same time, because visualizations are inherently nonlinear and provide only hints rather than direct relevance ratings, they lend themselves to parallel exploration. In a ranked list, all users would redundantly look at the top items; but (as the second author has observed many times!) not everyone will drill down into the same parts of a treemap. Thus according to this argument, visualizations are an effective way for a group to make discoveries—but only if those discoveries can be communicated easily and quickly.

The group perspective on information foraging theory provides a helpful complement to the other theoretical underpinnings described in this section. For instance, tightly coupled synchronous collaboration is sometimes assumed to be richer or better than decoupled, asynchronous work. But if the optimal strategy for a group information search is for individuals to look at different pieces of a data set, foraging theory would predict that asynchronous collaboration may be preferable even when synchronous work is feasible—providing additional motivation to study the different-time/different-place corner of the matrix in Figure 1A.

RESEARCH DIRECTIONS: SHALLOW VERSUS DEEP SHARING

Visualization-based communication raises design considerations that suggest new areas for investigation. As previously mentioned, users typically share visualizations by distributing images—through screen captures, printouts, and so forth. We refer to this kind of sharing as *shallow* because it is limited to the duplication of pixels, which represent a partial view of data. *Deep* sharing, on the other hand, refers to sharing of data in its entirety. Whereas asynchronous scenarios have been relegated to shallow sharing, deep sharing has taken place mostly in synchronous applications. Consider, for instance, collaborative visualization projects where every user has access to a fully interactive copy of the visualization. We believe that intelligently sharing visualizations asynchronously—that is, with annotations and tailored playback capabilities—poses an important challenge for developers and designers.

From our observation of how sharing of visualizations takes place,[3] two crucial capabilities emerge for users to successfully communicate: establishment of a common ground for visualizations and the ability to direct the viewers' attention (deixis). The psychological process of grounding—the process by which a group of people determines that they are all referring to the same thing, object, or proposition—has a long history of research.[23] A key concern is how people use language to negotiate the precise boundaries of what is being referred to in speech. In fact, CSCW studies show that when users are remotely located, their ability to establish common ground is one of four key social-technical conditions required for effective distance work.[24] The other three conditions are: coupling of work, collaboration readiness, and collaboration technology readiness. In CMV, we are interested in creating capabilities for users to establish common ground with respect to a visual domain. A big part of visual grounding may be the deictic facilities provided in the visualization system. As Hill and Hollan[3] indicate, the use of pointing behaviors in visualizations can be much more complex than simply directing viewers' attention. Instead, pointing may convey meanings as diverse as the height of a measurement, intervals, groupings, difference, and position. In addition, both Reference 3 and Reference 24 observe that there is a strong interaction between deictic behavior and memory. By pointing to the former location of now-absent graphics, users make use of a group's memory of an image in order to support ongoing discussion.

A challenge for CMV application designers is how to provide deep sharing while also providing effective deictic and grounding capabilities. One way to partition the CMV design problem space is to use the time and place matrix presented before. By applying the notions of shallow and deep sharing to the

■ Communication-minded visualization (CMV) is a new perspective that recognizes the important role played by social interactions during graphical analysis ■

matrix, it becomes clear that the more decoupled the sharing action is, the more challenging grounding and deixis become. We examine now what it means to design for deep sharing, while taking into account grounding and deictic capabilities in each one of the quadrants of the time-place space.

Same time and same place: Visual cues

As the most tightly coupled sharing activity in the entire time and place matrix, same time and same place situations need the least in terms of grounding and deictic features. Nevertheless, the viewer has to recognize actions (inputs) and their consequences (outputs). Tracking of inputs and outputs becomes more vital the larger the number of collocated participants. Above a certain group size, most participants become viewers, while only a few can interact with the visualization application.

It is extremely common for a small group of people to analyze data together, gathered around a single computer screen. At times this can be awkward; watching someone else operate a computer can be irritating, much like watching someone else use a remote control for a TV. One reason for this may be that changes in the display occur abruptly and unexpectedly, and the viewer loses track of what the active user is doing. The classification of spectator interfaces by Reeves et al.[25] is helpful here: they suggest that for the most understandable spectator interface, both the actions taken by the active user and the effects they generate should be easy to follow.

As described in Reference 2, a natural hypothesis is that visualizations should have as expressive an interface as possible. For example, *animated transitions* might be an optional enhancement for a single user, but might be critical for a viewer to maintain a sense of orientation. This was the hypothesis used in the design of PhotoMesa.[13] Similarly, changes to the widgets that control a visualization should be clearly visible, and keyboard shortcuts may need to be accompanied by visible redundant cues.

Same time and different place: Ability to point

Similar to collocated users, users working synchronously and remotely also need clarity regarding inputs and outputs. An addition critical issue, however, is the ability to point to elements in the visualization. When two people are working at the same computer, it is common to see them pointing to different parts of the screen.[3] Pointing is a problem for other kinds of synchronous remote work, but it can be especially acute for the unstructured displays of a visualization tool. With a spreadsheet it is easy to say, “Look at cell E3,” but with a typical graph layout tool it may be difficult to verbally identify a particular node.

One solution is to allow one person to move a cursor on the other’s screen, or to allow drawing of annotations. A more sophisticated idea in line with deep sharing is to enable complex types of pointing that are “data aware.” For example, in working with a treemap it is natural to point to a particular tree node and its set of descendants. One could imagine an interface in which a user right-clicks on a node to get a menu of possible highlighting options: put focus on the node, its children, all descendants, and so forth.

Different time and different place

Asynchronous and remote situations represent the pinnacle of communication decoupling. As such, establishing common ground plays an even more crucial role here than it does in any other quadrant of the time and place matrix. In fact, all considerations for building common ground discussed up to this point apply here as well. Visual cues, animated transitions, and pointing mechanisms should all be taken into consideration when designing CMV tools for asynchronous remote collaborations.

Playback

If a playback feature is available, it may be desirable to allow users to edit a session, picking out only a few key frames and discarding any false starts or superfluous navigation. This implies a capability to edit sequences of a visualization session. On a technical level, playing back a visualization implies the ability to completely reconstruct the state of the application from some sort of token sent from one user to another. Such a token could be a pointer to a central server where information is stored or in simple cases, could itself contain all the necessary information.

An important issue here, related to the challenges defined by Grudin,[15] is that the communication capabilities are most likely to be used when they work well with existing systems. For example, if users are passing a token representing state, it is helpful if this is a simple text string (much like a URL) that can be transferred by e-mail or instant messaging. In Reference 2 a simple example of such a scheme is discussed.

Annotation

Another important method of asynchronous communication is annotation. In many situations it may be helpful for users to be able to point to objects and add some sort of information. The added information may be anything from a single word to a long discussion. In nonvisualization settings, simple annotations have proved powerful: consider the Web site del.icio.us,[26] essentially an annotation service for the Web. Annotations raise difficult questions. A particularly thorny one is how to handle data that changes. What should happen to an annotation, for example, when an annotated data point is deleted? As with many systems that involve group creation of meaning (software development or Wikipedia), it may be helpful to have a versioning system. Another related design question for visualizations is whether the annotations themselves can become part of the visualization.

Information foraging

The previously mentioned information-foraging model suggests some additional designs for asynchronous communication. One of the implications of the theory is that users should spread their attention over the data space, and one can imagine several ways in which a system could support this behavior. For example, Reference 2 describes a scenario termed "antisocial navigation" in which the visualization shows which views have been seen by many users, thus encouraging new users to try different ways of analyzing the data.

A second requirement for efficient group data foraging is that users should be able to communicate discoveries quickly and easily. To support the communication of discoveries, a system might allow users to register interest in different parts of a data set. An individual might, for instance, make use of a "watch list" mechanism by registering to be notified of any change in a specified list of items in a large hierarchical data structure.

Different time and same place

The prototypical example of different-time and same-place communication is the discourse that occurs on a public display screen. All of the issues previously described are likely to be relevant, but there is an additional interesting design question. If the space is the same, then it makes sense to design the physical surroundings to augment the visualization. At a very simple level, it may make sense to place a printer, pens, and tape at the public display, so that users can print out key frames, draw annotations, and stick them to the walls. An open question is whether there are deep-sharing techniques that exploit this physical space.

EVALUATING CMV APPLICATIONS

It is, of course, important to validate the benefits of visualization-based communication. The case for CMV is promising but by no means proven. It frequently happens, as in the zooming transitions of Reference 13, that a feature is posited to help group interactions, yet is never tested in a group setting. It would also be valuable to study existing deployments of visualization applications such as DecisionSite Posters and CoMotion in order to discover exactly how often their special collaborative features are used. There has been arguably too little academic evaluation of real-world deployments of commercial systems, perhaps due to the "not invented here" syndrome. Another lesson from the CSCW field may be that important insights can come from studying commercial products such as Microsoft Excel**.[27]

The premise that the communicative aspect of data visualization is an important part of its value leads to questions about evaluation as well. One of the perennial problems in visualization research is the difficulty of evaluating designs. The most common approach to such an evaluation is a laboratory study, in which a small number of volunteers perform carefully specified tasks. Whereas such studies produce replicable, statistically significant results, there is a widespread feeling that they do not reliably test the true value of visualizations (Reference 4 names evaluation as one of 10 grand challenges facing the field). In a survey of visualization evaluations, Plaisant[28] notes the importance of reporting on long term use in natural settings (something that is hardly ever done now). The CMV perspective suggests some possible approaches to these challenges.

First, if the success or failure of a visualization system depends on how it lets groups collaborate and communicate, then it may make sense to perform tests on groups rather than individuals. This suggests a number of variations on current practices. Most obviously, one might simply test the performance of groups of two or more people using

■ The CMV design space is the well-known time-space matrix ■

same-time visualization. The approach may be further refined by measuring not only the direct performance, but also how well a viewer follows the visualization session. One might have one active user perform a series of actions and then test whether a viewer is able to repeat those actions. More subjective measures may be revealing as well. If a visualization has a "magical" or "secret" interface according to the classification in Reference 25, a viewer may feel that the active user is uncooperative or is hiding information. These ideas are speculative, but that is exactly the point: the CMV perspective points to an entire set of untried ideas that need to be evaluated.

It is also interesting to note that the field of CSCW has faced a similar problem: Grudin[15] points to several obstacles to assessing the value of systems for collaborative work. The main culprit is that the success of such a system depends upon complex social and environmental factors. As a result, it is extremely difficult to perform a laboratory study that will predict performance "in the wild." It is also, according to Grudin, hard to generalize from one deployment to another because so many factors may be different.

To summarize, if a major benefit of visualization is its ability to catalyze discussion, then the parallel with CSCW research may not be a coincidence. It may be useful to look to CSCW for ideas on how to overcome evaluation difficulties. Given the experience of CSCW researchers, we believe studies of visualization should include more studies of real-life deployments, with careful ethnographies to understand the surrounding environmental influences on success or failure.

CONCLUSION

Although improvements in computing technology (with faster computers allowing for smoother graphical rendering and new hard drives able to store vast data sets) have led to significant progress in information visualization, little has been done to support visualization-based communication practices. We propose CMV, a new perspective to remedy this oversight. This perspective recognizes the critical role played by conversations, be they synchronous or asynchronous, and social activity surrounding graphical data analysis.

Our survey of existing techniques shows that there are many ways in which designers can encourage visualization-based communication and that there is a need for a unifying framework to help our understanding of these techniques. In several cases, such as PostHistory and NameVoyager, affordances that led to social activity were inadvertent. In other cases, advanced ideas for CMV exist in commercial products and are neither widely known nor well studied. One of the goals of this paper is to point out the common themes that run through these disparate systems.

One source of ideas for studying CMV comes from CSCW. We have highlighted three theoretical constructs that we believe are especially relevant. Jonathan Grudin's work on problems in collaborative systems has natural applications to communicative visualizations. The standard partitioning of the design space into a synchronous/asynchronous and same-place/different-place matrix proves to be helpful here as well. Designing for visual grounding and deixis is one of the key aspects for the success of CMV. Finally, adapting information-foraging theory to groups provides a helpful framework for analysis and suggests new design ideas.

With these theoretical constructs in hand, we have suggested several research directions. First, in designing visualizations, we advocate a deep-sharing approach, in which multiple users all have access to the full system. At the simplest level, this may mean designing expressive interfaces so that two people at the same computer can easily follow each other's actions. At a more complex level, it can mean designing systems that are "bookmarkable," where the full state of the system may be specified by a URL-like token, allowing annotations and asynchronous conversations.

Finally, we have offered suggestions for new ways to assess visualizations. Evaluation has traditionally been difficult in this area—in fact, finding new ways to evaluate visualizations is listed by the National Science Foundation as a "grand challenge." We believe that one reason traditional evaluation methods have been insufficient is that they do not take into account group usage, and therefore, miss a significant portion of the benefits (and problems) present in a visualization system.

ACKNOWLEDGMENTS

Many people made critical suggestions to the development of the ideas in this paper. Jesse Kriss and Steve Rohall made many contributions, and conversations with John Patterson were influential as well. We also thank the many other members of the Collaborative User Experience group who have helped refine these ideas and brought related work to our attention.

**Trademark, service mark, or registered trademark of Maya Viz, LLC., Spotfire, Inc., or Microsoft Corporation in the United States, other countries, or both.

CITED REFERENCES

1. F. B. Viégas, D. Boyd, D. H. Nguyen, J. Potter, and J. S. Donath, "Digital Artifacts for Remembering and Storytelling: PostHistory and Social Network Fragments," *Proceedings of the 37th IEEE Hawaii International Conference on System Sciences*, Big Island, HI, January 5–8, 2004, p. 40109a.
2. M. Wattenberg, "Baby Names, Visualization, and Social Data Analysis," *Proceedings of IEEE Symposium on Information Visualization*, Minneapolis, MN, October 23–25, 2005, pp. 1–7.
3. W. C. Hill and J. D. Hollan, "Deixis and the Future of Visualization Excellence, *Proceedings of IEEE Visualization Conference*, IEEE, New York (1991), pp. 314–319.
4. *NIH-NSF Visualization Research Challenges Report*, E. Johnson, R. Moorhead, T. Munzner, H. Pfister, P. Rheingans, and T. Yoo, Editors, U.S. National Institutes of Health, Bethesda, MD (2006).
5. L. Bauerle (lead developer at Spotfire), personal communication, January 10, 2006.
6. M. Wattenberg, "Visualizing the Stock Market," *Proceedings of ACM Conference on Human Factors in Computing Systems*, Pittsburgh, Pennsylvania, May 15–20, 1999, pp. 188–189.
7. C. Ware, *Information Visualization: Perception for Design*, Morgan Kaufmann Publishers, San Francisco, CA (2000).
8. S. Roth, "Visualization as a Medium for Capturing and Sharing Thoughts," Capstone Address in *Proceedings of IEEE Symposium on Information Visualization*, Austin, TX, October 10–12, 2004, p. xiii.
9. V. Anupam, C. L. Bajaj, D. Schikore, and M. Schikore, "Distributed and Collaborative Visualization," *Computer* **27**, No. 7, 37–43 (1994).
10. Y.-P. Chui and P.-A. Heng, "Enhancing View Consistency in Collaborative Medical Visualization Systems Using Predictive-Based Attitude Estimation," *First IEEE International Workshop on Medical Imaging and Augmented Reality (MIAR '01)*, Hong Kong, China, June 10–12, 2001, p. 0292.
11. H. Benko, E. W. Ishak, and S. Feiner, "Collaborative Mixed Reality Visualization of an Archaeological Excavation," *The IEEE International Symposium on Mixed and Augmented Reality (ISMAR 2004)*, Arlington, VA, November 2–5, 2004, pp. 132–140.
12. K. W. Brodlie, D. A. Duce, J. R. Gallop, J. P. R. B. Walton, and J. D. Wood, "Distributed and Collaborative Visualization," *Computer Graphics Forum* **23**, No. 2, 223–251 (2004).
13. B. B. Bederson, "PhotoMesa: A Zoomable Image Browser Using Quantum Treemaps and Bubblemaps," *The 14th Annual ACM Symposium on User Interface Software and Technology (UIST 2001)*, Orlando, Florida, November 11–14, 2001, pp. 71–80.
14. R. Johansen, *Groupware: Computer Support for Business Teams*, The Free Press, New York (1988).
15. J. Grudin, "Groupware and Social Dynamics: Eight Challenges for Developers," *Communications of the ACM* **37**, No. 1, 92–105 (1994).
16. J. S. Brown and P. Duguid, *The Social Life of Information*, Harvard Business School Press, Boston (2000).
17. Google Maps, Google, Inc., http://maps.google.com.
18. MAPQUEST, MapQuest, Inc., http://www.mapquest.com.
19. T. Erickson, D. N. Smith, W. A. Kellogg, M. Laff, J. T. Richards, and E. Bradner, "Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of 'Babble'," *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 1999)*, Pittsburgh, PA, May 15–20, 1999, pp. 72–79.
20. F. B. Viégas and M. A. Smith, "Newsgroup Crowds and Authorlines: Visualizing the Activity of Individuals in Conversational Cyberspace," *Proceedings of 37th IEEE Hawaii International Conference on System Sciences*, Big Island, HI, January 5–8, 2004, p. 40109b.
21. P. Pirolli and S. K. Card, "Information Foraging," *Psychological Review* **106**, No. 4, 643–675 (1999).
22. A. Wexelblat and P. Maes, "Footprints: History-Rich Tools for Information Foraging," *Proceedings of ACM Conference on Human Factors in Computing Systems*, Pittsburgh, PA, May 15–20, 1999, pp. 270–277.
23. H. H. Clark and S. E. Brennan, "Grounding in Communication," In *Perspectives on Socially Shared Cognition*, L. B. Resnick, R. M. Levine, and S. D. Teasley, Editors, American Psychological Association, Washington, D.C. (1991), pp. 127–149.
24. G. M. Olson and J. S. Olson, "Distance Matters," *Human-Computer Interaction* **15**, No. 2/3, 139–178 (2000).
25. S. Reeves, S. Benford, C. O'Malley, and M. Fraser, "Designing the Spectator Experience," *Proceedings of*

ACM Conference on Human Factors in Computing Systems, CHI 2005, Portland, Oregon, USA, April 2–7, 2005, pp. 741–750.

26. del.icio.us, Yahoo! Inc., http://del.icio.us.
27. B. A. Nardi, *A Small Matter of Programming: Perspectives on End User Computing*, MIT Press, Cambridge, MA (1993).
28. C. Plaisant, “The Challenge of Information Visualization Evaluation,” *Proceedings of the ACM Working Conference on Advanced Visual Interfaces (AVI 2004)*, May 25–28, 2004, pp. 109–116.

Accepted for publication April 24, 2006.
Published online September 29, 2006.

Fernanda B. Viégas
IBM Research
Cambridge, Massachusetts

Martin Wattenberg
IBM Research
Cambridge, Massachusetts