Do Visualizations Improve Synchronous Remote Collaboration?

Aruna D. Balakrishnan, Susan R. Fussell, Sara Kiesler
Human Computer Interaction Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15232 USA
aruna@cs.cmu.edu

ABSTRACT
Information visualizations can improve collaborative problem solving, but this improvement may depend on whether visualizations promote communication. In an experiment on the effect of network visualizations, remote pairs worked synchronously to identify a serial killer. They discussed disparate evidence distributed across the pair using IM. Four conditions, respectively, offered (a) spreadsheet only (controls), (b) individual unshared visualizations, (c) view-only shared visualizations, and (d) a full-access shared visualization of all evidence. We examined collaborative performance, use of the visualization tool, and communication as a function of condition. All visualization conditions improved remote collaborators' performance over the control condition. Full access to a shared visualization best facilitated remote collaboration by encouraging tool use and fostering discussion between the partners. Shared visualization without full access impaired performance somewhat and made communication even more vital to identifying the serial killer. This study provides direct evidence of visualization tool features and partner behavior that promote collaboration.

Author Keywords
CSCW, distributed work, empirical studies, information visualization, collaboration, communication.

ACM Classification Keywords
H5.3. Group and Organizational Interfaces: Computer-supported cooperative work, synchronous interaction.

INTRODUCTION
Henry Wallace had killed nine women when he was finally arrested in 1994 for shoplifting at a mall and released [36]. Wallace was finally identified as a serial killer when a detective noticed strangulation in deaths two weeks apart. The department called a meeting to share information, and detectives noted that the last two victims lived in the same apartment complex. As this case illustrates, significant breakthroughs in detective work often come about when someone notices and associates disparate and sometimes unlikely facts or events—colloquially, "connecting the dots." Similarly, in intelligence analysis [17], business innovation [3], and scientific research [33], success may hinge on sharing and seeing linkages in previously unnoticed information.

Our focus is problem solving in which successful task performance, as in the example above, depends on whether individuals share information they have that is crucial to a group's ability to "connect the dots." We argue here that advances in computing that allow collaborators to visualize information open up new opportunities for collaborative problem solving that have failed in the past. For instance, in detective work, a shared map and database of offenses is feasible, potentially improving remote collaboration across investigators. In this paper, we address this possibility and examine the kinds of visualizations that foster collaborative problem solving.

Visualization techniques represent complex numerical and textual information in pictorial or graphical form and allow individuals or groups to perceive and explore patterns in data [1, 32, 41]. By removing the burden of mentally consolidating disparate information, such holistic representations of large amounts of data can help people spot anomalies, see new patterns, and improve their problem solving success [e.g., 21]. Information visualization tools can reduce task completion time and increase productivity on many information retrieval tasks and in data analysis [14, 34, 39]. Information visualizations available to a group can help promote feelings of community and foster discussion in "wiki" websites [40]. Visualizations also have costs—for individuals who must spend time learning how to manipulate them [2] and for groups, who can experience the tradeoffs involved in working alone versus being aware of others' work. Much less is known about the impact of visualizations in collaborative problem solving, although, a few studies have

CHI 2008, April 5–10, 2008, Florence, Italy.

examined collaborative problem solving on such highly complex tasks as investigative analysis [e.g., 4, 18, 38].

Complex problem solving as a collaborative task

Two defining attributes of real-world complex problem solving are that it is ill structured (in the sense that the problem definition is unclear) and, as noted above, that it often involves knowledge or information dispersed across people and groups [18, 33]. For instance, a detective in Illinois investigating a possible serial killer may sift through local cold cases looking for linkages but, unknown to this detective, relevant cases may exist in other states [e.g., 30]. Because of the need for insight and the fact of dispersed information, the success of criminal and intelligence investigations, scientific discovery, medical problem solving, and other important real world problems often depends on opportunistic cross-talk across information sources and serendipity [9, 33]. Collaboration can increase the likelihood that such cross talk and serendipity will occur. Collaboration can increase group performance over that of individual performance in these situations [15] but effective collaboration may depend on the free flow of information among partners [22, 35].

There now exist computer-based visualization tools that support scanning for hidden linkages and sharing dispersed information. Our research question is whether these tools do in fact change problem solving strategies, particularly information sharing, the collaborative relationship among partners, and ultimate collaborative task performance. We studied a type of network visualization application similar to those used in intelligence analysis and criminal investigations (for example, Analyst's Notebook, www.i2inc.com/Products/Analysts_Notebook/default.asp). Our experimental design tested whether the network visualization tool improved collaborative task performance of remote partners who were synchronously solving a complex analytic problem, and the extent to which sharing features in the tool affected the effectiveness of the collaboration.

Information Visualization in Collaboration

Previous studies have shown that visualizations can facilitate information sharing in collocated groups [7, 29]. Mark, Carpenter, and Kobsa [23], in a seminal study of visualization in collaboration, showed that collocated pairs' and remote pairs' use of visualization tools for making bar graphs of statistical data improved their analysis performance over that of participants using the tools alone. Our work builds on these promising results, examining how visualizations aid collaboration.

Visualization tools could aid collaborative problem solving in at least two ways. First, if each member of a group has a visualization of his or her own data, then the individual member's insight into the problem may improve, which in turn would raise the probability of better group performance. If so, visualization tools might not need to provide for jointly viewable or manipulated data, or even promote discussion, as long as they improve the problem solving of individuals in the group. This idea leads to the following general hypothesis:

Hypothesis 1: Access to a visualization tool will increase remote pair performance in complex problem solving.

Second, prior research suggests that visualization tools may improve collaborative performance because they allow for shared access to data, and encourage information sharing and discussion. In their evaluation of CACHE, a system that supports intelligence analysis via visual data presentation, Billman et al. [4] report that distributed pairs using CACHE collaboratively overcame *a priori* biases and did more effective data analysis. Mark et al.'s [24, 25] video analyses of their experimental data showed that remote pairs using a visualization communicated more intensively than collocated pairs. Their results suggest that communication was necessary to take best advantage of the visualization tools. This argument leads to the second hypothesis.

Hypothesis 2: Access to a visualization tool will increase remote pair performance in complex problem solving when this access increases information sharing and discussion by the pair.

We further asked how fully a visualization tool needs to support shared information and communication. Visualization tools can support different levels of sharing. At the most basic level, each member of a collaboration can visualize his or her own data but cannot see others' visualizations (Unshared Visualizations). Or, collaborators might be able to view their own and others' visualizations but would be able to directly manipulate only their own (Shared View-Only Visualizations); many applications can be shared in this manner. A third possibility is that collaborators have full access to a shared visualization application that allows for viewing everyone's data and jointly manipulating these data (Shared Full-Access Visualization) [e.g., 28]. Full access would support shared information sharing automatically, which might be especially important when collaborators doing complex problem solving do not have the same data [e.g. 4]. Full access also could promote joint attention and may help in the establishment of common ground more than applications that allow only shared views [20, 27].

Hypothesis 3: Access to a shared full-access visualization tool will encourage discussion between partners, and increase remote pair performance in complex problem solving beyond the performance of those using a visualization tool that supports unshared visualizations or shared but view-only visualizations.

METHOD

The design of this study was a single-level factorial, with four visualization conditions. Participants worked in pairs; pairs were randomly assigned to one of the four visualization conditions. The pairs were told they were

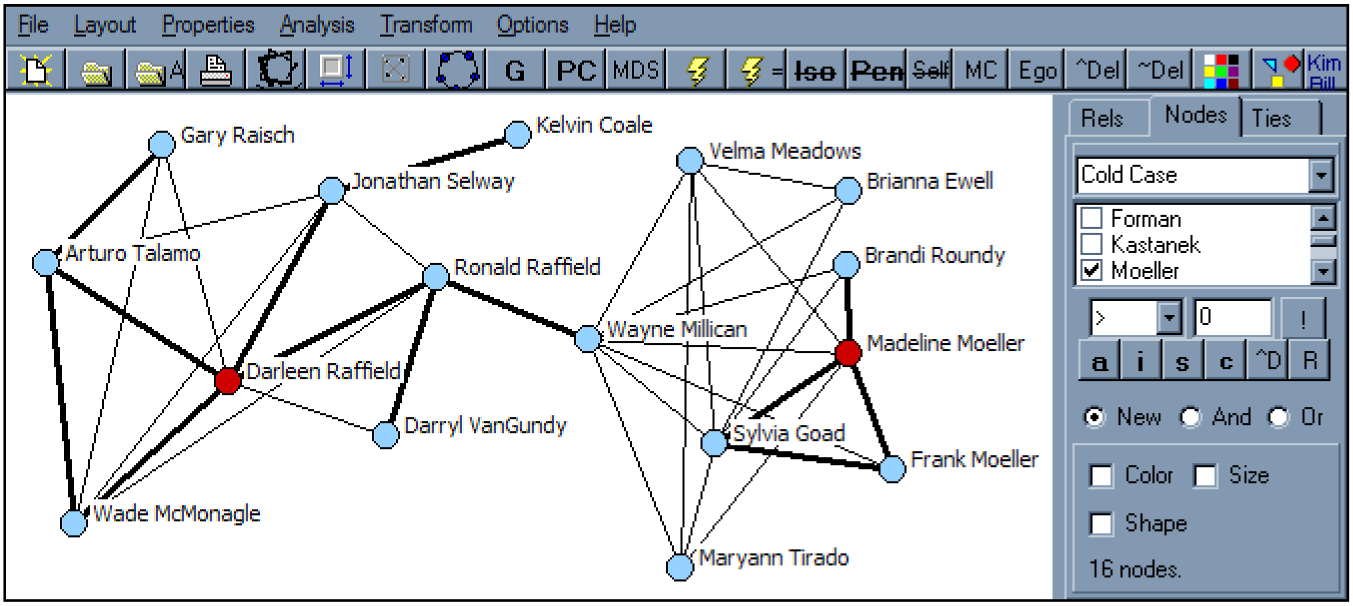

Figure 1. Screenshot of NetDraw, the network diagram tool used by participants.

members of the homicide unit of a local police department and had been assigned to a serial killer task force.

Participants

Ninety-four participants were recruited for a "Detective Mystery Study" (54 female, 40 male; 81% U.S. born; age range 18-64, median age approximately 22). Eighty percent of the participants were undergraduate or graduate students. Participants were paid $15 for their participation. They were told the experiment would last 1.5 hours.

Procedure

For the duration of the experiment, participants were seated apart, such that they could not see their partner or their partner's workstation. Participants role-played a pair of detectives in Zone 5 of a police department, working remotely together to identify a possible serial killer in Zone 5 and to complete reports on their findings.

The participants were trained on the visualization tool and the investigative task, given their detective task assignment, and left to work on the assignment for one hour. After an hour, or when the participants had completed their report, they completed an online survey to elicit their memory of the evidence about the serial killer. The experimenter then debriefed the participants.

Training. The participants practiced first on a comparatively simple problem involving the theft of a laptop computer from a college locker room. They read documents containing evidence relevant to four suspects in the theft and were asked to organize the data using a template that organized the evidence as to the motive, opportunity, and alibi of each suspect. Then they practiced on a more complex problem involving a rash of electronic equipment thefts. The case was constructed to give participants experience scanning and organizing information across crimes. Participants also were shown how to use a timeline and geographic worksheet.

Participants also were trained how to use NetDraw (see Figure 1), the visualization tool adapted for this study, if they were assigned to one of the three visualization conditions. (Controls were trained on a spreadsheet that contained the same data.) A sample network diagram depicted the connections among the crimes in the second practice case. Participants were familiarized with the concepts of nodes and relationships, and they practiced using search and manipulating the diagram by location, time, and type of theft to give different perspectives on the crimes. Participants were encouraged to ask questions throughout the training. Training, on average, took 30 minutes.

Complex Problem Solving Task

The pairs' task was to identify a possible serial killer in Zone 5. Each participant was instructed to report any other important information that might help their department solve other cases.

Documents and reporting forms. Evidence related to the serial killer was scattered in 15 assorted documents summarizing 6 cold cases and one open homicide, which also functioned as a simple problem solving control task. The documents included witness and suspect interviews, and coroner's reports. There were additional documents on crime statistics by police district zone, a map of the zone and adjacent zones, a bus route map, and a police department organizational chart. Participants also could use an MO (modus operandi) worksheet for recording dates, weapons, and other relevant evidence for each case, a suspect worksheet for recording different suspects, their connection to the victim, and given alibis, and a timeline worksheet for recording when and where each crime took place, intended to support inter-case linkages. Finally, participants were asked to complete two online reports on the results of their investigation, one on their serial killer analysis and another to report any other criminal activity they wanted to convey to the department.

All of the evidentiary documents and reports were available online and could be opened, searched, put in different or new folders, and manipulated freely. To insure that sufficient screen space was available to examine multiple documents at once, the participants each had access to two 17" display monitors placed side by side. Also, the participants were given paper versions of the instructions and worksheets.

Dispersed Evidence. The serial killer was responsible for four of the six homicides in the cold cases folder. Eight pieces of evidence, six within the cold case files and one in the open homicide case file, could be linked to the serial killer: similar blunt force trauma injuries to the victims; victims killed in the evening after they returned from work; victims rode the same bus route; victims lived near the same bus route; offender worked at a local hospital on the bus route; offender had been identified on the bus (alibi for a homicide witness); offender had been seen carrying a tool box on the bus. Identifying the serial killer required conceptually linking these disparate pieces of evidence from different cases rather than simply eliminating a defined group of suspects in one current case folder.

The caseload and evidence for the serial killer were distributed evenly between each member of the pair. To accomplish this, the six unique cold cases and the documents of the current open homicide case were divided between the pair such that each member received 3 distinct cold cases and half of the documentation for the current homicide case.

The open homicide case concerned the murder of a woman named Darlene Raffield. To solve this homicide, participants only had to examine the documents in one folder, review the alibis of witnesses, and evaluate their motives and opportunities to commit the crime. If a pair spent time on this case, they would have less time to focus on the complex serial killer task. In pretesting, we found that individuals who spent more time on the Raffield homicide were less likely to identify the serial killer.

Communication. Participants were given an MSN Instant Messenger (IM) client and encouraged to use the client to talk with their partner. All IM conversations were recorded.

Visualization Independent Variable

Each pair was randomly assigned to one of four conditions, differing with respect to their access to a visualization tool. The tool enabled participants to see social and information network relationships in the data that linked names, places, events and objects, thereby providing a visual analysis perspective to identify the serial killer.

The tool used to create a social network diagram of all persons mentioned in the documents was NetDraw v.2, a software application for drawing 2D social network diagrams available online from Analytic Technologies. Social network diagrams are aptly suited for complex problem solving of the kind we used. The evidence documents contained over 50 unique names and a diagram that represented how each person was connected to various other persons could serve to help participants categorize and group people, and to view how they might be connected across cases. At the start of a session, each participant (except for those in the control condition) received the software set up to show a predetermined social network diagram reflecting the relationships in the documents they had.

Figure 1 is a screen shot of the application. Within the diagram, each circle (a "node") represented a person from the crimes and each line represented a relationship between two people. Victims were represented in red and other persons (such as witnesses and suspects) in blue. (Printed in black and white, victims are black and others are grey.) Thick lines denoted a strong tie, such as married people or coworkers. Thin lines denoted a weak tie, such as two people who happened to be at the same place at the same time (e.g., a waiter serving a restaurant customer or two people who rode the same bus route).

Participants could freely manipulate and move the nodes within the screen, but they could not change underlying relationships. Participants also could search or filter the diagrams based on a set of attributes to reveal people with common characteristics. Searchable attributes included police district zone affiliation, case affiliation, occupation, mode of transportation, time of crime, location of crime, weapon used to injure the victim, and the injured body part of the victim. For example, within the attribute weapon, the three options were handgun, blunt instrument, and poison. If handgun were selected, all victims who were injured by a handgun would be visible on the screen.

The four experimental conditions varied the degree of access that participants had to NetDraw.

No Visualization. In this control condition, the pairs did not have access to NetDraw. To ensure that they received the same information as did participants in the other conditions, they were given Microsoft Excel spreadsheets containing the same relationship information among the persons mentioned in the evidence documents. The names of these people were arranged to form a matrix. Relationships in the matrix were represented by 0, 1, or 2 scores, which reflected no relationship, a weak relationship (such as a witness), or a strong relationship (such as a family member), respectively. Each participant received a spreadsheet that contained the relationship data only for their own cases. The experimenter explained the use of the spreadsheet and the meaning of the numerical data.

Unshared Visualizations. Each member of the pair had access to NetDraw and a manipulable and searchable social

	1	2	3	4	5	6	7
Collaborative Performance							
1. Pair identified serial killer (0 - 1)							
2. Time spent problem solving (minutes)	- .77 **						
Visualization Tool Use							
3. Visualization selected (min.)	.11	- .10					
4. Visualization tool active (min.) [a]	- .07	.14	.94 **				
Communication							
5. Total IM (# IM lines)	.21	- .26 [t]	.00	- .34			
6. Discuss serial killer (# IM lines)	.27 [t]	- .39 *	.14	- .30	.83 **		
7. Discuss Raffield homicide (# IM lines)	.08	- .01	- .01	- .04	.67 **	.21 [t]	
8. Discuss visualization (# IM lines)	.31 *	- .26	.41 **	- .10	.49 **	.62 **	.27 [t]

[t] p < .10, *p < .05, ** p < .01
[a]Visualization conditions only

Table 1. Correlation of measures of pair performance, use of the visualization tool, and communication (N = 47).

network diagram of the data for the cases that they were given. They could not view their partner's visualization.

Shared View-Only Visualizations. As in the condition above, each member of the pair had access to NetDraw and a manipulable and searchable social network diagram of the data for the cases that they were given. Each participant also had a window within which they could view their partner's social network diagram. They could not search or manipulate this diagram but could view how their partner acted upon it. The diagrams were shared using the Share Applications feature within MSN's Messenger client.

Shared Full-Access Visualization. As in the Unshared Visualization and Shared View Only Visualization conditions, each member of the pair had access to NetDraw and a manipulable and searchable social network diagram of data but unlike the conditions described above, participants shared access to a network diagram that contained data from all the cases. This diagram could be manipulated and searched by both participants in the pair. This diagram was shared via a third computer using TightVNC, an open-source remote desktop software application.

Measures

We had four main sources of data: the final reports participants completed when their hour was up (or earlier if they had completed their analysis), an online posttest survey, IM logs, and WinWhatWhere files that recorded use of the visualization tool.

Individual and Collaborative Performance. Participants' identification of the serial killer were culled from their written reports. The reports were coded for whether they named the serial killer, and whether they named the Raffield offender. Mentioning the name of the correct offender as guilty or a primary suspect who should be arrested was counted as identification. We scored individuals but were mainly interested in the success of collaboration, so both members of the pair had to have named the serial killer for the pair to be coded as having collaborative successful performance.

Visualization Tool Use. Online activities were recorded via WinWhatWhere, a software tool that records the application a participant is using, the time spent with that window as the selected window, all keystrokes, and screenshots of the selected window. Due to resource constraints, one randomly selected participant within each pair was recorded. To estimate tool use, we calculated the total amount of time these participants had the visualization tool as the selected window. Active use was highly correlated with visualization window selection (see Table 1). In analyses, the total minutes the tool was selected and was active were log transformed to adjust for skewness.

Communication. We calculated how much conversation occurred between members of a pair by counting the total number of IM lines they exchanged during a session. An IM line refers to each new line within the recorded IM logs.

Participants' attention to different topics was coded in the IM conversations. IM logs were coded by line for whether the participants were discussing the serial killer task, whether they were discussing the Raffield homicide, and whether they were referring to the social network diagram (See Table 2 for the coding scheme). IM was coded to be related to the serial killer task if the line clearly showed the participants talking about or working on searching for patterns of a serial killer, for example, "Here we have

Topic	Definition	Example
Serial killer task	Pertains to solving the serial killer task or evidence pointing to the serial killer.	"I see a connection between 2 of my cold cases; they both involve a blunt object."
Raffield homicide	Discussion pertaining to solving the Raffield homicide.	"I think the person who poisoned Darlene is Wade."
Visualization	References the visualization tool or the visualization.	"My diagram says that Wayne is involved in the Raffield case." "Move those two out of the way."

Table 2. Conversational coding scheme.

another blunt instrument incident," or "How do we connect these cases?" Discussion of the Raffield homicide was coded if the IM line referenced any person related to the Raffield homicide or if the line clearly showed them thinking about facts relating to the case, for example, "what did Darlene Raffield's boss say?" Because a single IM conversation line could be affiliated with both the Raffield homicide and the serial killer task, these counts were not mutually exclusive. For example, some participants discussed whether the Raffield homicide was connected to the serial killer task. References to the social network diagram were coded if the IM line directly referenced the diagram, for example, if participants used words such as "diagram," "visualization," and "picture," or discussed their active search within the diagram, such as "Watch this" and "See how these pop out?" The percentage of total IM lines during which IM lines referenced the visualization was calculated and log transformed. Over 5,000 lines of IM were coded using the scheme. An independent coder coded 7% of the data (Kappa = .76).

RESULTS

We obtained data from 47 pairs (94 participants), 13 pairs in the No Visualization condition, 10 pairs in the Unshared Visualizations condition, 12 pairs in the Shared View-Only Visualizations condition, and 12 pairs in the Shared Full-Access Visualization condition.

Preliminary Analyses

To insure the task was equally difficult and enjoyable across conditions we administered the NASA TLX workload scale [12], CRT Scales [10], and measured task enjoyment on the posttest. Mean scores did not differ by condition, suggesting that cognitive ability, cognitive load, and enjoyment were equal across conditions. To insure that correctly identifying the serial killer reflected comparable insight across conditions, on the posttest survey we tested participant's recognition memory for the eight pieces of evidence leading to a serial killer (multiple choice questions). Again, there were no differences across conditions.

Table 1 shows the correlations of measures on the pairs. These allow us to examine across all conditions whether visualization-related communication is associated with collaborative success. The table shows that, overall, when pairs identified the serial killer, they also had communicated more about the serial killer and talked more about the visualization. Active use of the visualization tool was not directly associated with communication; this result could be due to partners' opening their visualization window once and then moving to talk and to view documents.

Individual and Collaborative Performance

We first examined performance on the simple problem—the Raffield homicide. We did not ask pairs to solve this case, but about one-third of the pairs did so anyway. We believe they did so in part because it was an easy way to get something done when the pair had trouble identifying the serial killer. Consistent with this argument, the correlation between identifying the serial killer and solving the Raffield homicide was r = -.20. There were no differences across conditions solving the Raffield homicide, suggesting that having a visualization tool does not influence performance on a simple problem.

We next examined individual performance on the serial killer case. According to our arguments, if visualization improves individual performance then that improvement might translate into a greater likelihood of collaborative success. Because the dependent variable, solving the serial killer, is a discrete variable, the appropriate analysis is a logistic regression [16]. This regression assesses whether visualization conditions predict the dichotomous outcome (identified the serial killer or not). The logistic regression analyses at the individual level showed a highly significant influence of condition on whether individuals identified the serial killer (logistic regression Likelihood Ratio $\chi 2 = 12.1$, $p < .01$, $df = 3, 93$) with the No Visualization condition different from the other conditions ($\chi 2 = 5.75$, $p = .01$).

We predicted in Hypothesis 1 that using a visualization tool would increase collaborative performance over performance in the control condition. We defined collaborative success when both members of the pair reached consensus and correctly identified the serial killer. We conducted analyses at the pair level, whose results are shown in Figure 2.

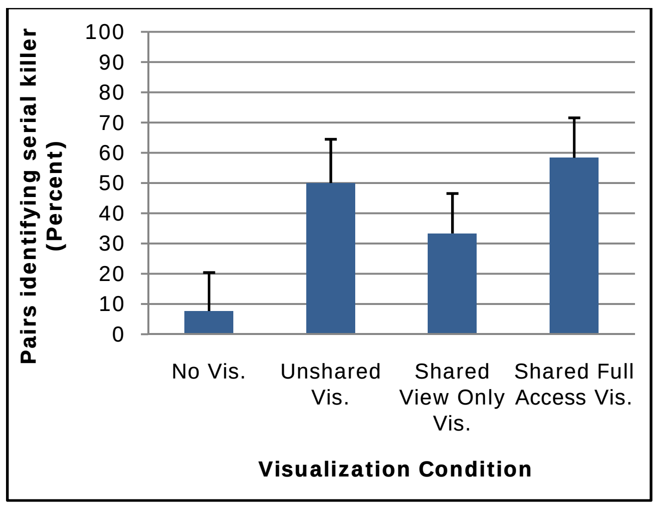

Figure 2. Percent of pairs solving the serial killer task by condition.

Pairs in all three visualization conditions outperformed pairs in the No Visualization condition, as predicted. Only 7.7% (SE = 12.7) of pairs in the No Visualization condition identified the serial killer, whereas 50% (SE = 14.5) of pairs in the Unshared Visualization condition, 33.3% (SE = 13.2) of pairs in the Shared View Only Visualization condition, and 58% (SE = 13.2) of pairs in the Shared Full Access Visualization condition identified the serial killer (logistic regression Likelihood Ratio $\chi 2 = 9$, $p < .05$, $df = 3, 46$). Student's t tests revealed significant differences at the $p <$

.05 level between the two best conditions—Full Access Visualization and Unshared Visualizations versus the No Visualization controls.

To test whether visualization helped the collaboration over and above the help it gave individual members of each pair, we conducted a *nominal pairs* analysis [43]. For this analysis, we compared actual problem solving success of each pair in each condition with the simulated success of all possible other pairs in the same condition. The idea here was to compare these nominal (in name only) pairs with the actual pairs, to evaluate the extent to which collaboration really mattered when visualization was given to pairs. The results of this analysis can be seen in Figure 3. They show that, controlling for condition, performance was worse by nominal pairs than by actual pairs (logistic regression Likelihood Ratio $\chi 2 = 3.04$, $p = .08$). In nominal pairs, the top mean performance in the Shared Full Access Visualization condition was only 48% (SE = 14.4). These analyses indicate that although visualizations aided individuals, collaborative performance was superior and benefited from using the visualization tool.

The results support Hypothesis 1 and show that visualization increases collaborative performance but the comparatively weak performance of the pairs in the Shared View-Only Visualization condition suggests that features of the tool do matter. The next section delves into the different tool use and communication in the three visualization conditions, and tests of Hypotheses 2 and 3.

Visualization Tool Use and Communication

We predicted that access to a visualization tool would increase remote pair performance in complex problem solving when this access increased information sharing and discussion by the pair (Hypothesis 2). The first step was to establish that participants with access to the visualization tool actually used it. They did. In the No Visualization condition, on average, participants spent 2.7 minutes with the spreadsheet selected. By contrast, in the Visualization conditions, participants on average spent 5.7 minutes with the network diagram opened (F [3, 43] = 4.1, p = .01, d = .63). Pairs in the two Shared Visualization conditions used the visualization tool more than did the pairs in the Unshared Visualization condition (F [2, 44] = 3.36, p < 0.05, d = .57). As shown in Figure 4, Shared View-Only Visualization pairs used the visualization tool the most (M = 6.84 minutes, SE = 1.02), followed by Shared Full Access Visualization pairs (M = 5.14 minutes, SE = 0.89), and then Unshared Visualization pairs (M = 2.83 minutes, SE = 0.54). A contrast revealed that this difference was significant when comparing both shared conditions against the unshared condition (F [1, 30] = 6.37, p < 0.05, d = .69). These results indicate that access to shared visualizations does encourage tool use.

Hypothesis 3 stated that the Shared Full Access Visualization would promote discussion and joint problem solving. We thus tested whether the participants in the Visualization conditions, particularly in the Shared Full Access condition communicated differently than those in the other conditions. We found no overall effect on the total amount of IM conversation in the pairs, but a significant effect on talking about the network diagram versus the spreadsheet (F [3, 43] = 2.8, p < .05, d = .52; see Figure 5).

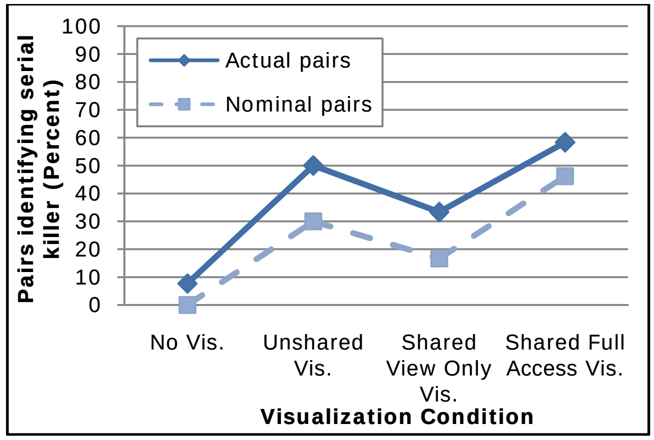

Figure 3. Percent of actual and nominal pairs solving the serial killer task, by condition.

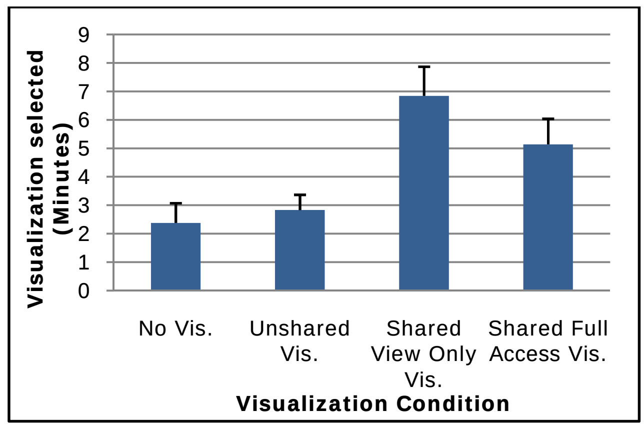

Figure 4. Mean number of minutes during which participants had the visualization selected, by condition.

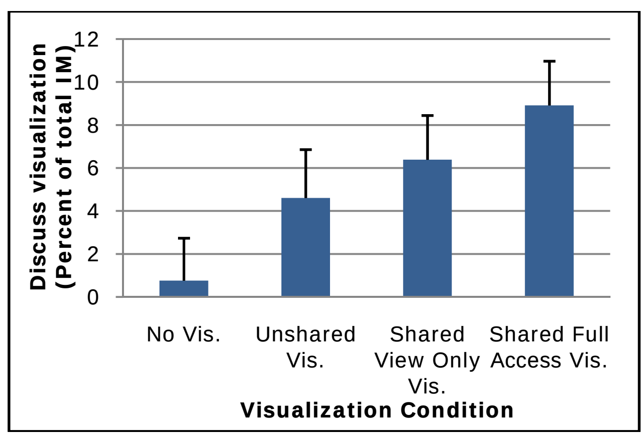

Figure 5. Mean percent of total IM lines during which pairs discussed the visualization, by condition.

According to a Student's *t* test, pairs in the Shared Full Access Visualization condition talked significantly more about the network diagram (9% of IM lines) than did pairs in the other Visualization conditions. (5% of IM lines) or pairs (talking about the spreadsheet) in the No Visualization condition (<1% of IM lines).

How was talking about the visualization relevant to identifying the serial killer? To examine this question, we looked at whether those who identified the serial killer talked differently with their partners than those who did not in the three Visualization conditions. (Because we did not manipulate communication directly, these are correlational analyses.) The analyses showed interesting relationships: Across the three Visualization conditions, controlling for condition, the higher the percentage of pair discussion about the network diagram, the higher was the percentage of discussion about the serial killer (F [1, 30] = 7.9, $p < .01$, $d = 1.1$). The more pairs discussed the serial killer, the more likely they were to identify the serial killer (logistic regression Likelihood Ratio $\chi 2 = 6.2$, $p < .05$). These analyses point to the visualization as a contributor of solutions to the complex serial killer case.

DISCUSSION

We studied the impact of a shared visualization tool on difficult collaborative problem solving (overall, even after one hour, only 36% of pairs solved the case collaboratively). The visualization tool made a significant difference, improving not only individual performance, but also collaboration.

Although differences among the tools tended to be overshadowed by the positive impact of having any visualization at all, the Shared View-Only Visualization was comparatively unhelpful. That is, having full manipulable access to the shared visualization (Shared Full Access Visualization condition) encouraged pairs to use the tool and fostered more discussion and better performance—58% of the pairs solved the serial killer case. By contrast, when pairs had a tool that gave shared views but no ability to manipulate others' data (Shared View-Only Visualization condition), there was a dip in performance—only 33.3% of pairs solved the serial killer case.

We wondered if the mere awareness of the partner's visualization (but no ability to manipulate it) was a distraction to these pairs, explaining why mere awareness was worse than not seeing the partner's view at all (Unshared Visualization condition). A decade ago, Gutwin and Greenberg [11] proposed that collaborative systems involve "mixed focus" whereby individuals must attend to their own work and to that of others. They analyzed a groupware concept map editor, showing that the tradeoffs of a shared view may include less individual flexibility and more attention to coordinating with others, detracting from accuracy.

Overall, we did not see that tradeoff. In the Visualization conditions, those who solved the problem actually spent less time on task, whereas in the No Visualization condition, solving the case was correlated $r = +.15$ with time spent. Discussion was positively correlated with solving the case in the two best visualization conditions ($r = .50$) and negatively correlated with time spent problem solving ($r = -.45$). Talk was uncorrelated with solving the case in the less effective Shared View Only condition. We can only speculate, but possibly, when each member of the pair had his or her own visualization and could only stare at the other person's diagram and manipulations, the tradeoff proposed by Gutwin and Greenberg prevailed. Alternatively, the two nonintegrated diagrams of data might have violated the "proximity compatibility" principle of display design [42], and confused pair members.

Our study is a step in the direction of understanding how visualizations can aid collaboration. Our nominal pair analysis (see Figure 3) showed that real collaboration was valuable on this task but we do not know exactly how pairs came to aid one another, for example, whether they formed a common mental model of the problem [8, 26] or whether they simply tried harder because the visualization was fun and motivating [40]. Future research will be needed to study these potential consequences of the use of collaborative visualization tools. How joint representations are created, perceived and given meaning is still being explored and understood as different from linguistic and gestural cognitive processing [5, 6, 44].

Limitations

This study cannot be generalized now to other genres of visualization tools, to other task types (such as decision making), or to other remote collaborative settings. For example, sharing information through IM may have introduced barriers to the effective flow of information or made visualizations particularly effective in a way they would not otherwise be effective. Previous studies have shown that IM provides an effective channel of communication between partners [e.g., 31] but an audio chat feature could help us understand the role of different channels in the use of visualization tools.

Participants were all given predrawn social network diagrams. One might argue that taking a more active role in creating the diagram would aid pairs in understanding their data [37]. However, a recent trend in using social network diagrams for analysis is to use diagrams automatically generated from an existing dataset. Oftentimes the datasets have millions of different records from which diagrams are produced, so a real challenge is how to engage users in helping to create them.

This study examined synchronous interactions. In distributed teams, colleagues often do not work simultaneously. Asynchronous collaborative visualizations can encourage to knowledge discovery [13]. Asynchronous communication and access to the information visualization tools would be most similar to our Unshared Visualization condition. Pairs did quite well in this condition (50%

solution rate). Thus our findings suggest that asynchronous teams would benefit from the use of such tools to solve complex problems.

Design Implication

This study has just one main implication for designing tools for complex problem solving—that is, create a visualization tool. Our study also suggests that, for most effective collaboration, visualization tools should have the capability of being jointly manipulated and should facilitate the integration of data from different collaborators.

CONCLUSION

Information visualization in the form of a network diagram aided both individual and collaborative complex problem solving. Real collaboration improved the performance of pairs over statistical pairings, particularly if pairs (a) had an integrated visualization that both could manipulate, and (b) when pairs discussed the visualizations they received. Doing so led to more relevant discussion of evidence and higher solution rates.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation Grant IIS-0325047. We thank Peter Scupelli and Gail Kusbit for their assistance in developing experiment materials.

REFERENCES

1. Andrews, K., & Heidegger, H. (1998). Information Slices: Visualization and Exploring Large Hierarchies using Cascading, Semi-Circular Discs. *Proc. Information Visualization 1998*. NJ: IEEE press.
2. Ariely, D. (2000). Controlling the information flow: Effects on consumers' decision making and preferences. *Journal of Consumer Research*, 27, 233-248.
3. Baron, R. A. (2006). Opportunity recognition as pattern recognition: How entrepreneurs "connect the dots" to identify new business opportunities. *Academy of Management Perspectives*, *20*(1), 104-119.
4. Billman, D., Convertino, G., Pirolli, P., Massar, J. P. and Shrager, J. (2005). Collaborative intelligence analysis with CACHE: Bias reduction and information coverage. Unpublished manuscript. Palo Alto Research Center, CA. Retrieved from http://cscl.ist.psu.edu/public/users/gconvert/mypapers/hcic2006_BillmanEtAl.pdf.
5. Cheng, P., Lowe, R., & Scaife, M. (2001). Cognitive Science Approaches To Understanding Diagrammatic Representations. *Artificial Intelligence Review*, *15*(1), 79-94.
6. Clancey, William J. (1994). Situated cognition: How representations are created and given meaning. In Lewis, R. and Mendelsohn, P. (Eds.), *Lessons from Learning* (pp. 231-242). Amsterdam: North Holland.
7. Edelson, D., Pea, R., Gomez, L. (1996). Constructive in the Collaboratory. In B.G. Wilson (Ed.), *Constructivist Learning Environments: Case Studies in Instructional Design*. Englewood Cliffs, NJ: Educational Technology Publications.
8. Fiore, S., Salas, E., Cuevas, H., & Bowers, C. (2003). Distributed coordination space: toward a theory of distributed team process and performance. *Theoretical Issues in Ergonomics Science*, *4*(3), 340-364.
9. Fraidin, S. N. (2004). When is one head better than two? Interdependent information in group decision making. *Organizational Behavior and Human Decision Processes*, *93*(2), 102-113.
10. Fredericks, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, *19*, 25-42.
11. Gutwin, C., & Greenberg, S. (1998). Design for individuals, design for groups: Tradeoffs btween power and workspace awareness (pp. 207-216). *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW '98), Washington, D.C. NY: ACM Press.
12. Hart, S. G., & Staveland, L. E. (1988). Development of a multi-dimensional workload rating scale: Results of empirical and theoretical research. In P. A. Hancock & N. Meshkati (Ed.), *Human mental workload* (pp. 139-183). Amsterdam: Elsevier.
13. Heer, J., Viegas, F. B., & Wattenberg, M. (2007). Voyagers and voyeurs: Supporting asynchronous collaborative information visualization. *Proc. CHI 2007* (pp. 1029-1038), NY: ACM press.
14. Hendrix, T. D., James H. Cross, I. I., Maghsoodloo, S., & McKinney, M. L. (2000). Do visualizations improve program comprehensibility? Experiments with control structure diagrams for Java. *SIGCSE Bulletin*, *32*(1), 382-386.
15. Hill, G. (1982). Group versus individual performance: Are N+ 1 heads better than one. *Psychological Bulletin*, *91*(3), 517-539.
16. Hosmer, D. W., & Lemeshow, S. (1989). *Applied Logistic Regression*. NY: John Wiley & Sons.
17. Huerer, Jr., R. J. (1999). *The psychology of intelligence*. Washington DC: Center for the Study of Intelligence, Government Printing Office.
18. Johnston, R. (2005). *Analytic culture in the U.S. intelligence community*: *An ethnographic study*. Washington DC: Center for the Study of Intelligence, Government Printing Office.
19. Klahr, D., & Simon, H. A. (1999). Studies of scientific creativity: Complementary approaches and convergent findings. *Psychological Bulletin*, *125*, 524-543.
20. Kraut, R. E., Fussell, S. R., Brennan, S. E., & Siegel, J. (2002). Understanding effects of proximity on collaboration: Implications for technologies to support remote collaborative work. In P. Hinds & S. Kiesler (Eds.), *Distributed work*, 137-162. Cambridge, MA: MIT Press.

21. Larkin, J., & Simon, H. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*, *11*(1), 65-100.

22. Lavery, T., Franz, T., Winquist, J., & Larson, J. (1999). The Role of Information Exchange in Predicting Group Accuracy on a Multiple Judgment Task. *Basic and Applied Social Psychology*, *21*(4), 281-289.

23. Mark, G., Carpenter, K., & Kobsa, A. (2003a). Are There Benefits in Seeing Double? A Study of Collaborative Information Visualization. *Proc. CHI 2004* (pp. 840-841), NY: ACM Press.

24. Mark, G., Carpenter, K., & Kobsa, A. (2003b). A model of synchronous collaborative information visualization. *Proc. Information Visualization 2003* (pp. 373-381), NJ: IEEE Press.

25. Mark, G., Kobsa, A., & Gonzalez, V. (2002). Do four eyes see better than two? Collaborative versus individual discovery in data visualization systems. *Proc. Information Visualization 2002* (pp. 249-255), NJ: IEEE Press.

26. Mohammed, S., & Dumville, B. (2001). Team mental models in a team knowledge framework: expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, *22*(2), 89-106.

27. Monk, A. (2003). Common ground in electronically mediated communication: Clark's theory of language use. In J.M. Carroll (Ed.), *HCI Models, Theories and Frameworks: Towards a Multidisiplinary Science* (pp. 265-289). San Francisco, CA: Morgan Kaufmann.

28. Pang, A., & Wittenbrink, C. (1997). Collaborative 3D visualization with CSpray. *Computer Graphics and Applications*, *17*(2), 32-41.

29. Ryall, K., Forlines, C., Shen, C., & Morris, M. (2004). Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. *Proc. CSCW 2004* (pp. 284-293). NY: ACM press.

30. Safarik, M. E., Jarvis, J., & Nussbaum, K. (2000). Elderly female serial sexual homicide. *Homicide Studies*, *4*, 294-307.

31. Scupelli, P., Kiesler, S., Fussell, S.R., & Chen, C. (2005). Project view IM: a tool for juggling multiple projects and teams. *Proc. CHI 2005* (pp. 1773-1776). NY: ACM Press.

32. Shneiderman, Ben. (1996). The eyes have it: A task by data type taxonomy for information visualizations. *Proc. Visual Languages 1996* (pp. 336–343). NJ: IEEE Press.

33. Simonton, D. K. (2003). Scientific creativity as constrained stochastic behavior: The integration of product, person, and process perspectives. *Psychological Bulletin*, *129*, 475-494.

34. Stasko, J., Catrambone, R., Guzdial, M., & McDonald, K. (2000). An evaluation of space-filling information visualizations for depicting hierarchical structures. *International Journal of Human Computer Studies*, *53*(5), 663-694.

35. Stasser, G., & Titus, W. (1987). Effects of information load and percentage of shared information on the dissemination of unshared information during group discussion. *Journal of personality and social psychology*, *53*(1), 81-93.

36. State of North Carolina v. Henry Louis Wallace, 2000. URL: www.aoc.state.nc.us/www/public/sc/opinions/2000/241-97-1.htm. Downloaded June 16, 2007.

37. Suthers, D., & Hundhausen, C. (2001). Learning by Constructing Collaborative Representations: An Empirical Comparison of Three Alternatives. *Proc. Euro CSCL 2001* (pp. 577-592). Maastricht, the Netherlands: Maastricht McLuhan Institute.

38. Tolcott, M. A., Marvin, F. F., and Bresoick, T. A. (1989). *The confirmation bias in military situation assessment*. Reston, VA: Decision Science Consortium.

39. Veerasamy, A., & Belkin, N. (1996). Evaluation of a tool for visualization of information retrieval results. *Proc. SIGIR 1996* (pp. 85-92). NY: ACM Press.

40. Viégas, F., Wattenberg, M., & Dave, K. (2004). Studying cooperation and conflict between authors with history flow visualizations. *Proc. CHI 2004* (pp. 575-582). NY: ACM Press.

41. Wattenberg, M. (1999). Visualizing the stock market. *Proc. CHI 1999* (pp. 188-189). NY: ACM Press.

42. Wickens, C. D., & Carswell, C. M. (1995). The proximity compatibility principle: Its psychological foundation and its relevance to display design. *Human Factors*, 37, 473-494.

43. Wright, D.B. (2007). Calculating nominal group statistics in collaboration studies. *Behavior Research Methods*, 39, 460-470.

44. Zhang, J., & Norman, D. (1994). Representations in Distributed Cognitive Tasks. *Cognitive Science*, *18*(1), 87-122.