

A Model of Synchronous Collaborative Information Visualization

Gloria Mark, Keri Carpenter, Alfred Kobsa
Department of Information and Computer Science
University of California, Irvine
{gmark, kcarpent, kobsa}@ics.uci.edu

Abstract

In this paper we describe a model of the process by which people solve problems using information visualization systems. The model was based on video analysis of forty dyads who performed information visualization tasks in an experiment. We examined the following variables: focused questions vs. free data discovery, remote vs. collocated collaboration, and systems judged to have high and low transparency. The model describes the stages of reasoning and generating solutions with visual data. We found the model to be fairly robust across task type, collaborative setting, and system type, though subtle differences were found. We propose that system transparency can support some stages of the process, and that support is needed in the last stage to help users translate their findings from visual to written representations.

1. Introduction

Recently there has been a surge of interest in developing systems to aid in visual data mining and discovery. As more and more real world tasks involve collaborative decision-making about data, the response has been the development of a large number of collaborative visualization environments (see e.g., CoVis [5], Cspray [16], CVD and Cave6D [11], TIDE [19], iScape [3], COVISA [24], and several proto-types of DIVA [20] and [4]). Yet despite the rise of development of systems to support visual data mining, so far user studies have been largely neglected. We maintain that user studies are essential to understand how groups can use visualization systems to make decisions about data.

In [12] we performed a user study and discovered that groups were more accurate than individuals in problem-solving with information visualization systems, but only when the system was transparent. (Transparency thereby refers to the system's quality to invoke an easy-to-understand system image in users [18].) Yet in this previous study we did not identify the processes by which groups made decisions. We feel this is important in helping to inform design. In this paper we report on new results that address more in-depth the processes underlying collaborative information visualization. This

study is part of a larger investigation to understand differences between collaborative and individual information visualization.

Current practice in data analysis and computer-supported cooperative work (CSCW) in organizations led us to consider examining several variables in our study. First, there are different purposes for using information visualization systems. In a survey of daily activities of data analysts, [6] found that answering focused questions is currently the most frequently occurring type of data analysis. Focused question tasks are when one has a very specific query in the data. For example, with census data, one might ask, "What proportion of the people in the U.S. are Catholic?" A second type of task is when people use information visualization systems for exploratory analysis. An example of such a free discovery task is when a researcher explores trends in the U.S. urban population using census data. We thus examined how decision-making processes are affected by the type of task.

Second, current collaborative practice in organizations involves collaboration both when people are collocated and also remote from each other. These distinctions are made in the field of CSCW ([8], [15]). We therefore investigated how different kinds of collaborative configurations affects decision-making processes using information visualization systems.

Third, the value of offering system transparency has been shown to provide benefits such as greater visibility for interactivity and feedback [10], satisfaction [14], and in general, an improved understanding of the relationships between system input and output [8]. We would therefore expect that an information visualization system with greater transparency would lead to different kinds of behavior. We also looked at how decision-making processes differ when systems offer different degrees of transparency.

1.1. Related work

[13] developed a methodology for collaborative data mining where participants work at different times. The methodology is neutral with regard to the data mining methods used (which may include information visualization). Participants' activities are far more

independent than is the case in our same-time setting. Based on this methodology, [2] conducted an experiment with teams in four different European countries who performed collaborative data analysis including information visualization. They found "that the added value of collaboration of different groups on this task was much smaller than hoped" and suggest that "information exchange should be made more efficient and synchronization should be improved". Our results reported in this paper caution that even if one makes information exchange much more efficient (by switching from different-time to same-time collaboration, and direct communication) this does not necessarily increase the added value of collaboration.

1.2 Visualization systems used in this study

In this experiment, we have chosen two different visualization systems: InfoZoom (formerly Focus [21][22]) and Spotfire (formerly IVEE [1]).¹ InfoZoom presents data in three different views. The *wide view* shows the current data set in a table format, with rows representing the attributes and columns the objects. The *compressed view* compresses the visualization horizontally to fit the window width. Numeric data values are thereby plotted as horizontal cell-wide bars whose distance from the row bottom corresponds to their values. A row may be sorted in ascending or descending order, with the values in the other rows being rearranged accordingly to make each column represent one object. This operation reveals dependencies between characteristics (like positive or negative correlations between numeric attributes). Hierarchical sorting of two or more attributes is possible as well and can, e.g., reveal differences in the distribution of numeric attributes dependent on one or more non-numeric attributes. In the *overview mode*, the values in the rows become detached from their objects. Rows here represent the value distributions of attributes in ascending or descending order, and are independent of each other. In all three views, values of (identical adjacent) attributes become textually, numerically or symbolically displayed whenever space permits this, which facilitates the comprehensibility of the data [9].

Fig. 1 shows portions of InfoZoom's compressed mode, with data from a web-based dating service containing users' self-descriptions. The weights of the individuals are hierarchically sorted in ascending order by their response to the question "Do you participate in sports?" The upper line shows the weight distributions of these groups (there is a small third group that did not answer this question). The lower line shows the average

weight per group. Those who engage in sports are heavier than those who do not.

InfoZoom's central operation is "zooming" into information subspaces by double-clicking on attribute values, or sets/ranges of values. InfoZoom thereupon shows records only that contain the specific attribute value(s).

Spotfire offers several types of mostly familiar



Figure 1: InfoZoom's compressed view

visualizations, including scatterplots, bar charts, pie charts, graphs, parallel coordinates, trellises, etc. Unlike in InfoZoom, they are independent stand-alone visualizations. Two variables can be selected for display in the x and y coordinates, and a few additional variables can be selected for coding by color, shape, etc. Fig. 2. shows a bar chart visualization of the average weights by response to the sports question. Focusing on information subspaces is performed by excluding or including attribute values using sliders, checkboxes and radio buttons in the so-called "query device" of the system.

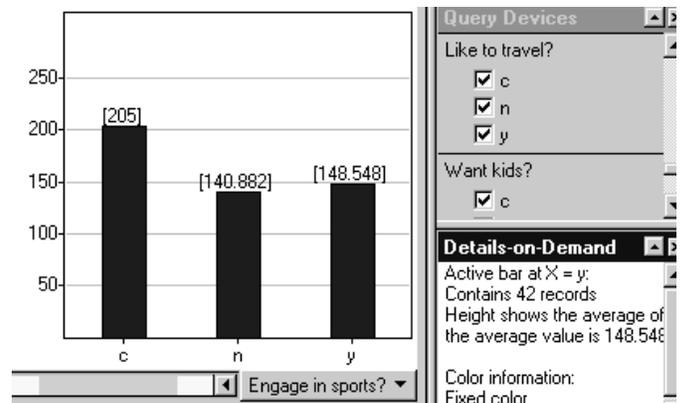


Figure 2: Bar chart visualization in Spotfire

1.2. System transparency

The design of the interface of information visualization systems uses a variety of approaches. A system is generally regarded as "transparent" if it evokes an easy-to-understand system image in users [18]. Conspicuity of the system's functionality, for instance through toolbars with well-designed icons, contributes towards system transparency. For information visualization systems, we posit that the ready visibility of

¹ The software versions used were InfoZoom 3.40 EN Professional from humanIT AG (www.humanIT.com) and Spotfire DecisionSite 6.3 from Spotfire, Inc. (www.spotfire.com).

all data dimensions also contributes to system transparency since this enables users to construct more easily a model of the visualized data and the underlying domain.

The two systems seem to rank differently in terms of system transparency. In Spotfire, a significant portion of the system's functionality is not immediately visible and directly accessible. Instead, a "Properties" menu must be accessed which contains controls for numerous parameters that can influence the shape of the visualization. On the data visualization side, Spotfire displays relatively few variables that must be deliberately selected. [9] found that Spotfire therefore imposes fairly high "cognitive setup costs" on users: they needed considerable time to choose among the different visualizations that Spotfire offers and to set them up correctly, particularly when solutions included several variables and therefore several steps. Users must plan ahead what variables to use and how to visualize them. Selecting a visualization for the first one or two variables imposes severe restrictions on the display of the remaining variables. The upfront planning must therefore be very thorough and comprehensive so as to avoid dead ends. This planning moreover must be performed without assistance from a visualization and takes considerable time.

Users of InfoZoom, in contrast, can rely much more often on visual cues when accessing both data and system functionality. (The only major exceptions are the dialog windows for defining new "derived" variables and for the charting function, both of which violate the otherwise relatively straightforward "click on what you want" paradigm.) [9] found that users can interact very effortlessly with the system. Users can also plan incrementally, i.e. perform a few steps, see how far they have come, and proceed or switch to a different view if the partial plan turned out to be wrong (some users even developed a "click first - think afterwards" problem solving behavior). InfoZoom can therefore be regarded as more transparent than Spotfire.

2. Methodology

2.1. The Experiment

In this paper we report on a video analysis of subjects performing information visualization tasks in an experiment. One hundred undergraduate students participated who had majors in computer science or engineering at the University of California, Irvine. All subjects had at least one year of computer usage. Subjects were paid \$25 for their participation.

Subjects were randomly assigned in a 2 x 3 between-subjects design, to use either InfoZoom or Spotfire, and to work in three types of configurations:

1. Alone (N=20 subjects),
2. Remotely where two subjects sat at workstations in adjacent rooms using an application-sharing program, Netmeeting, and a speaker phone (N=20 pairs; see Figure 3), or
3. Collocated where two subjects worked side-by-side in front of a SmartBoard, a large 60" diagonal touch-sensitive electronic whiteboard (N=20 pairs; see Figure 4).

The experiment lasted two hours: 20 minutes of training, 30 minutes of focused questions, and 40 minutes of exploratory data discovery. The remaining time was used for setting equipment and transitions between tasks.

2.2. Analysis Methods

Subjects were audio and videotaped. We also recorded their screen activity using a screen capture utility. The Alone and Netmeeting subjects were each videotaped using an IBM computer camera mounted on the top of their computer. The Smartboard subjects were videotaped using a regular videocamera set up approximately 5 ft. from the Smartboard. In neither condition could subjects view their images on the video.

Two different coders analyzed the videotapes that were projected onto a wall. The screen recordings and videotapes were viewed jointly. Grounded theory [23], a research methodology that involves the systematic generation of theory from data, was used to identify concepts unique and common to the conditions. 75 hours of videotapes were coded.

2.3. Tasks performed

Subjects performed two types of tasks using the visualization systems. The first task was a focused question task in which subjects viewed anonymized data from an online dating service to answer ten specific questions, such as: "Did males cheat more on their girlfriends than females on their boyfriends?" (other examples are given in Figs. 1 and 2). This was an objective task where correct answers could be measured. The second task was an open-ended discovery task in which subjects were instructed to discover as many findings in the data of a population survey as they could. The data was of general relevance such as Gender, Wage, Years of Experience, etc. and no special background knowledge was required to comprehend it.



Figure 3. Subjects working in Netmeeting condition

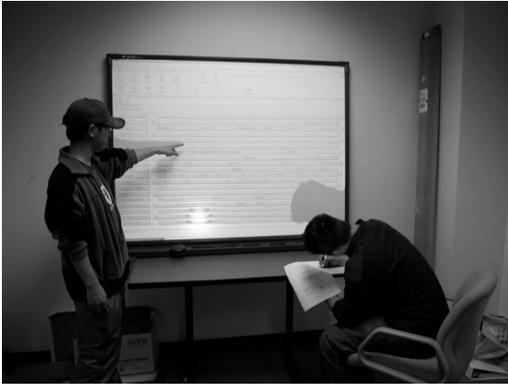


Figure 4. Subjects working in the Smartboard condition using InfoZoom

3. Results

We first report the speed of task completion for the focused question task.

3.1. Time to completion

Table 1 shows the average completion time per problem for focused questions. An ANOVA indicates that subjects using InfoZoom were significantly faster in solving the problems than subjects using Spotfire: $F(1,58)=6.8, p<.01$. Netmeeting and Smartboard conditions showed no difference, and when combined, an ANOVA showed that groups were slower than individuals: $F(2,58)=5.6, p<.02$. There was no significant interaction.

Interaction	System		Average
	InfoZoom	Spotfire	
Alone	116(103)	123(100)	120(101)
Netmeeting	134(101)	176(156)	154(131)
SmartBoard	115(114)	169(145)	141(132)
Group Avg.	125(107)	172(150)	149(129)

Table 1: Average completion time per problem (and standard deviations) in seconds for focused questions

3.2. A model of the process of collaborative information visualization

In this section, we present a model of the process that we found to hold in the group conditions for solving information visualization tasks. We describe how the model can take into account the type of task, the kind of tool used, and the nature of the collaborative setting.

Our model is based on an analysis of the verbal interaction and system behavior in the group conditions. This observation of the human-computer interaction supplied numerous clues regarding the sequence and intentions behind the process of solving information visualization tasks. In the Alone condition, we could not infer the reasoning processes, as there was no verbal interaction and the screen activity did not provide enough clues to inform us of the individual subjects' processes. The model therefore only applies to collaborative settings.

We discovered that the process of collaborative information visualization occurs in a series of stages (figure 5). Though some differences existed according to task, tool, and collaborative configuration, the model can still explain the basic processes.

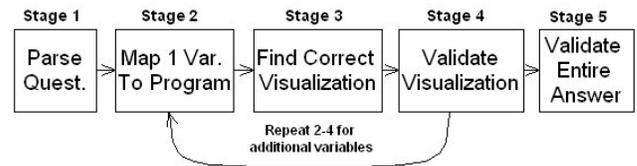


Figure 5. Stages of information visualization

In the first stage of the process for the focused question task, subjects read the question and parsed it into distinct variables. The subjects identified variables in the question that were represented in the visualization programs. For example, the question: *Are the people who do not participate in sports heavier than the rest?*, was

parsed into phrases containing variables, i.e. “participate in sports” and “heavier”.

In stage two, subjects then mapped one variable from the question (usually the first variable that appeared) to the corresponding variable representation in the visualization system. For example, the variable, “participate in sports” was mapped onto the variable in the data set that was labeled “engage in sports?” (see figure 1).

In stage three, users manipulated the system to find the appropriate visualization so that the variable could be visualized in a meaningful way. For the variable “participate in sports”, subjects needed to compare people who participated in sports with those who did not. They therefore viewed visualizations that displayed each of these representations.

In stage four, users validated the results from the variable that they had just examined, reassessing whether the results helped them progress towards the overall goal of solving the problem.

If there was agreement about the correctness of the visualization and an idea of where to go for the next step, there was little verbal interaction. However, if one subject felt the display was incorrect or the system operator did not know what to do next, an extensive interaction ensued and often a change of cursor control. In the free discovery task, validation was somewhat different: subjects made sure that they were making progress towards interesting findings. If so, they proceeded to the next variable. If not, they discussed how to rectify the situation (which was often associated with a change of cursor control between subjects).

Once the validation was completed for the current variable, subjects cycled back to Stage 2 in the process, choosing the next variable to be represented. Following the same example above, the variable “heavier” was mapped to the variable “weight” in the visualization system.

The process of trying to discover how to represent a current variable sometimes revealed incorrect reasoning about a previous variable. For example, reasoning about the variable “heavier” sometimes revealed incorrect thinking about the previous variable “participate in sports”. In such a case, subjects would return to re-examine the variable “participate in sports” to include both proportions of those who participated in sports and those who did not. Subjects then continued to examine the variable “weight”. Once an average of weight was determined, they proceeded on to Stage 5, an overall validation.

In the last stage, the users validated the entire answer and recorded their results. Subjects explicitly conferred with each other to make sure that their answer satisfied every aspect of the question. This stage primarily

involved verbal interaction and little system use. Subjects then translated their numerical results into written prose, recording it on paper. The main challenge for subjects at this stage was to translate the visual representation into a written representation of the results.

3.2.1 Explaining task types with the model

The model is slightly modified to explain the process of discovering answers in the free discovery task. In the focused question tasks, the goal in stage 1 was to parse the question into variables. In the free discovery task, subjects instead first had to choose variables to search for. The search space consisted of the entire dataset and thus subjects needed to bound the search space in order to discover “interesting” results. For example, one subject might say, “*Let’s look at female wages.*” The groups usually articulated a two-variable question of interest as a starting point to focus on distinct variables. The subjects then parsed this proposal into “female” and “wages”.

Stages 2 through 4 proceeded the same with both task types. In the free discovery task, stages 2 through 4 were cycled through until a potential “interesting” result was found. Groups did not always stick to their original propositions from stage 1. They often added more variables or abandoned previously “interesting” variables in order to find a result that passed their threshold of “interesting”.

Stage 5 functioned slightly differently in the free discovery task. Once a result was judged “interesting”, validation and translation of the representation occurred. As subjects did not begin with a pre-defined question, there were often multiple ways of describing the results. For instance in the example above, groups could either report the wages of union vs. non-union females or make an explicit comparison, as in:

“Females in unions earn, on average, \$10.50 per hour while females not in unions earn \$8.79” or

“Females in unions earn, on average, more than those not in unions.”

This process of translating the visual representation into a meaningful written answer often required much discussion.

There were task type differences observed at the end of stage 5. In the focused questions, subjects moved on to the next question. In the free discovery tasks, subjects did not start over with new variables but instead tended to build upon previous results. Using the same example, groups usually kept the representation already displayed in the system and searched for a variable such as “live in

south?” to discover what effect living in the south had on female union wages. If this was not interesting enough, they would then investigate the effect of being married on female union wages, and so forth. Subjects usually began a new round of searching for variables at stage 2 rather than stage 1. Returning to stage 1 was only done if groups felt that they had exhausted all the “interesting” findings working with their current variables. This happened seldom: on the average of three times during the 40-minute discovery task.

3.2.2 *Explaining system differences with the model*

System differences were most evident in stages 2, 3, and 4 of the model. In stage 2, InfoZoom users made fewer errors mapping the variables to the program. All variable names in this system are in a list on the left-hand side of the screen and almost all are visible. In contrast, with Spotfire, relevant variables can appear in either or both of two screen spaces. Often relevant variables must be scrolled through on the top right hand section of the screen. By not having all the variables visible on one screen, errors in mapping the variables can occur.

In stage 3, finding the correct visualization, Spotfire users tended to make more errors. Spotfire has more options with which to manipulate the information. Manipulation could happen using the query devices or the XY chart. With InfoZoom, users had only one view of the data, and therefore they tended to have fewer problems manipulating the system to find a correct view.

We observed no differences in stage 4, i.e. in validating the variables, with Spotfire and InfoZoom. The validation involves mostly human interaction, and little system interaction, and thus groups seemed to be able to interpret the results from InfoZoom and Spotfire equally.

3.2.3 *Explaining collaboration type with the model*

The social dynamics of the different collaborative configurations created subtle differences in the process. In the collocated and remote conditions, subjects adopted quite different roles. In the Smartboard (collocated) condition, one subject primarily worked alone through all the stages until stage 5, when both subjects participated in validating the whole answer. In Netmeeting (remote), both subjects participated collaboratively throughout all stages of the process.

With Netmeeting use, subjects tended to divide their labor into distinct roles of “system operator” and “system director”. The system operator manipulated the system

while the system director gave instructions to the partner on how to interact with the system and to oversee overall progress towards the goal. If the system operator became “lost”, cursor control was then switched.

In contrast, with Smartboard use, subjects tended to divide their labor into distinct roles of “system user” and “observer”. The person who stood on the left-hand side of the screen controlled the application, as the most controls for manipulating the visualization were located on the left side of the screen. After a few minutes, the person who initially stood on the right hand side of the screen soon sat down and became the scribe for the pair. Negotiation of one’s position in front of the whiteboard was done during the first few questions. The system user stood by the Smartboard, operated the system, and made decisions throughout all the stages of the process. The observer watched and generally only interacted during stage 5, in validating the entire answer. For both collocated and remote groups, once the role of the controller of the application was set in place, hands were changed only if the controller became “lost”.

4. DISCUSSION

We have introduced a model derived from empirical data which shows stages in reasoning about problems using information visualization systems. We found the model to be fairly robust with respect to different systems, different tasks, and different collaborative settings. Another stage model to describe collaborative information visualization was presented by Park et al. [17]. They also found similar stages as we did with problem interpretation (stage 1 in this paper), agreeing on visualization tools (stage 3), and negotiating a conclusion (stage 4). Their study differs from ours in that they provided private data views and as a result they also found stages where partners searched independently and reported results back to the partner. In our study, subjects rarely worked independently, even in the Netmeeting condition.

4.1. **The role of task type**

We chose to study both focused questions and free discovery of data, as we felt that these are common tasks that information visualization systems are used for. The model explained the process with both kinds of tasks well. However, the model helped to uncover some important process differences involved in the different tasks that should be considered.

For focused questions, subjects followed a deterministic path to the solution in the sense that there was one correct solution and clear goal. The choice of the representation was therefore crucial to determining

whether the answer was found. Therefore, Stage 3 is critical, in that if the appropriate visualization was not found, it affected the likelihood of finding the correct answer. Solving focused question tasks is therefore a top-down process, and the choice of the visualization is critical to finding the correct answer.

In contrast, for free discovery tasks, the process is rather opportunistic. People are generally not given specified instructions of what to search for, though they may have a general goal. We found that groups defined variables of interest as they interacted with the system. In this sense, stage 3 is also very critical, as the visualization chosen will affect people's ability to discover relationships, and thus it affects the quality of the results.

It is in stage 5 where groups have the most opportunity to find errors as they are validating the entire answer. For the focused question tasks, the participants confirmed that all aspects of the question were satisfied. Conversation for the free discovery tasks was much more extensive. It is important to consider that in a free discovery task, it is much more difficult to retrace one's steps to validate an answer against the visualization as the process is opportunistic. We observed that there was more interaction in stage 5 for free discovery tasks compared to focused questions, and this may indicate that validation may require extra effort.

4.2. The role of collaborative setting

We found that the type of collaborative setting affects the division of labor. Contrary to our expectation, working collocated did not lead to equal participation among the partners.

Overall, using Netmeeting led to more active participation among both subjects in the dyad. Both participants had equivalent views of the data (each on their own display in separate rooms) and only an audio channel for communication. In contrast, both subjects in the Smartboard condition could see the activity on the same physical display, and much interaction occurred via nonverbal communication.

Different coordination devices were thus used by the groups. With Netmeeting, confirmations of how to proceed were done explicitly. We observed that subjects frequently spoke their plans out loud (e.g., choosing an attribute) and waited to proceed until their remote partner agreed with this. In contrast, subjects using the Smartboard rarely explicitly confirmed their plans to their partners. People in the collocated condition referred to the board to emphasize their comments, directing the partner's attention to specific points on the screen. In the remote condition, people used the cursor to also refer to areas, but this was understood less clearly. Often, the

person controlling the cursor had to ask their partner more than once to attend to where the cursor was pointing. A cursor's position was not always evident that it meant "referring" to part of the data. Without explicit mention, the cursor could just have been resting at that position. In contrast, when a person pointed or gestured on the Smartboard, it was quite clear they were "referring" to the data, and were requesting the attention of their partner to that area.

We interpret these observations to mean that the collocated partners assumed that the other was aware of the context, i.e. the context was shared. However, the remote partners did not share this assumption. They needed to explicitly and continually establish the shared context. Thus, coordination in the remote condition required more effort as system use needed to be more explicit.

With Netmeeting use, both participants are aware that the process slows down if the system operator makes an incorrect choice. We surmise that the equal participation of both partners minimizes this possibility. Both participants track the process and the system director can take over the system manipulation at any time with as little reorienting overhead as possible.

With Netmeeting use, the system director could devote all of her attention to reasoning about the problem, while the other partner controlled the system. In the Smartboard condition, the main time that the observer usually contributed reasoning was in helping validate the results in Stage 4 and 5.

We did not find differences in accuracy [12] or time to completion between the remote and collocated conditions. One reason for this could be that in both conditions, both subjects participated in validating the answer. Thus, with two people checking answers, there was a greater likelihood of finding errors.

4.3. The role of the system

In [12], it was found that InfoZoom users answered more focused questions correctly and produced significantly more findings than Spotfire users. In this paper, we discovered that InfoZoom users were significantly faster in solving focused questions.

We explain these results according to the model. First, InfoZoom offers users a full view of all the variables, which results in fewer errors in mapping the variables to the representations in the system. Second, Spotfire offers more options (i.e. more data views) for visualizations. This particularly impacted Stage 3, as more choices will take users longer to find the appropriate visualization. Also, we argue that more data views can potentially lead to more errors, as users can

pick the wrong data view for making decisions (e.g. a pie chart as opposed to a scatter plot).

Validation of the answer mostly involves human interaction (checking with the system and with each other). However, we argue that the ability to retrace one's steps through the system to validate an answer, is an important factor affecting the results.

4.4. Design recommendations

Information visualization systems should be designed in such a way as to support the different stages in the information visualization process. We noted two system characteristics that seem to provide such support.

- *System transparency to support stages 2 through 4:* As described in more detail in Section 1.2, a previous experiment with individual subjects revealed that the more transparent system (InfoZoom) made it easier for subjects to select variables (since nearly all of them were directly visible at the interface), and to loop through stages 2 – 4 (since they could do this incrementally while Spotfire subjects had to plan the loops in advance). We noted similar effects in collaborative settings, which contribute to the faster task solution times with InfoZoom.
- *Summarization to support stage 5:* In this stage, subjects “retranslate” the visualization into natural language. We assume that system support in this retranslation will help analysts discover errors in the selected visualization more easily. InfoZoom features a status line that summarizes the selections made in zooming operations, which seems to be a first rudimentary step in this direction.
- Finally, we also noticed that subjects who were standing on the left-hand side in the Smartboard conditions had much greater control over the visualization than subjects on the right-hand side, due to the fact that important control devices of the visualization were located on the the left side. It remains to be seen whether roles become more equal in face-to-face collaborative visualization when interface controls are more equally distributed.

5. Conclusions

We developed a model of the information visualization process based on an empirical study. The model showed that people follow stages of reasoning that are fairly robust across different task types, systems, and collaborative configurations. The model helped us understand that strategies for interacting with the information visualization environment appear to be influenced both by system design and social boundaries.

We believe an understanding of this process is useful for designers of systems in terms of thinking about how their tool can assist the user with respect to different stages in the process. Designers should also keep in mind the different configurations (collocated or remote) and task types that their system will likely be deployed in. Working remotely and collocated each have their own social conventions which affect system use.

We believe this data will contribute to developing information visualization systems that are truly designed with the user and the processes that they go through in mind.

6. Acknowledgements

This research was supported by the National Science Foundation under grant no. 0093496, by the Center for Research on Information Technology and Organizations (CRITO), and by instrument grants from Smart Technologies Inc, humanIT AG and Spotfire Inc. We would like to thank Victor Gonzalez, Richard Hyunh, Jeffrey Cheng, Cristina Gena and David Lim for helping with the experiments and their evaluation.

7. References

- [1] C. Ahlberg and E. Wistrand, IVEE: An Information Visualization and Exploration Environment, InfoVis'95, New York, NY, 1995, pp. 66-73.
- [2] Blockeel, H. and S. Moyle (2002). Centralized Model Evaluation for Collaborative Data Mining. Conference on Data Mining and Warehouses (SiKDD 2002), Ljubljana, Slovenia, <http://www-ai.ijs.si/DunjaMladenic/SiKDD02/papers/BlockeelSep02.pdf>
- [3] K. Börner, A Collaborative Memory Palace for Digital Library Search Results, in *Usability Evaluation and Interface Design. Proceedings of the 2001 International Conference on Human-Computer Interaction.*, vol. 1, M. J. Smith, G. Salvendy, D. Harris, and R. J. Koubek, Eds. London: Lawrence Erlbaum, 2001, pp. 1160-1164.
- [4] I. Brewer, A. M. MacEachren, H. Abdo, J. Gundrum, and G. Otto, Collaborative Geographic Visualization: Enabling Shared Understanding of Environmental Processes", InfoVis 2000: IEEE Symposium on Information Visualization, Salt Lake City, UT, 2000, pp. 137-141.
- [5] D. Edelson, R. Pea, and L. Gomez, Constructivism in the Collaboratory, in *Constructivist Learning Environments: Case Studies in Instructional Design*, B. G. Wilson, Ed. Englewood Cliffs, NJ: Educational Technology Publications, 1996.
- [6] Gonzalez, V. and A. Kobsa. User Adoption of Information Visualization Systems. *I-KNOW'03 Workshop on Knowledge and Information Visualisation 2003 (KIV2003)*, Graz, Austria, 2003.

- [7] R. Johansen, R. (1988). *Groupware: Computer Support for Business Teams*. New York, NY, Free Press.
- [8] Johnson, J. and Johnson, P. (1993). Explanation facilities and interactive systems. *Proceedings of Intelligent User Interfaces'93*, 159-166.
- [9] A. Kobsa, An Empirical Comparison of Three Commercial Information Visualization Systems, *IEEE Symposium on Information Visualization*, San Diego, CA, , pp. 123-130. 2001. <http://www.ics.uci.edu/~kobsa/papers/2001-INFOVIS-kobsa.pdf>
- [10] Koenemann, J. and Belkin, N. (1996). A case for interaction: A study of information retrieval behavior and effectiveness. *Proceedings of CHI'96*. New York: ACM Press.
- [11] C. Lascara, G. Wheless, D. Cox, R. Patterson, S. Levy, A. E. Johnson, and J. Leigh, TeleImmersive Virtual Environments for Collaborative Knowledge Discovery, *Advanced Simulation Technologies Conference*, San Diego, CA, 1999 <http://evlweb.eecs.uic.edu/aej/papers/astc99.pdf>
- [12] Mark, G., Kobsa, A. and Gonzalez, V. (2002). Do four eyes see better than two? Collaborative versus individual discovery in data visualization systems. *Proceedings of IEEE Sixth International Conference on Information Visualization (IV'02)*, London, July 10-12, 2002. IEEE Press, pp. 249-255.
- [13] Moyle, S. and A. Jorge (2001). RAMSYS-A methodology for supporting rapid remote collaborative data mining projects. *ECML/PKDD01 Workshop: Integrating Aspects of Data Mining, Decision Support and Meta-learning (IDDM-2001)*, Freiburg, Germany, September 2001, <http://ai.ijs.si/branax/iddm-2001-proceedings/workshop/Moyle.pdf>.
- [14] Muramatsu, J. and Pratt, W (2001). Transparent search queries: Investigating users' mental models of search engines. *Proceedings of SIGIR*. New York: ACM Press.
- [15] G. M. Olson and J. S. Olson, Distance matters, *Human-Computer Interaction*, vol. 15, pp. 139-179, 2000.
- [16] A. Pang and C. M. Wittenbrink, Collaborative 3D Visualization with CSpray, *IEEE Computer Graphics and Applications*, vol. 17, pp. 32-41, 1997.
- [17] Park, K. S., Kapoor, A. and Leigh, J., Lessons learned from employing multiple perspectives in a collaborative virtual environment for visualizing scientific data. *Proceedings of ACM CVE 2000*, San Francisco, pp. 73-82.
- [18] J. Preece, Y. Rogers, and H. Sharp, *Interaction Design: Beyond Human-Computer Interaction*. New York, NY: Wiley, 2002.
- [19] N. Sawant, C. Scharver, J. Leigh, A. Johnson, G. Reinhart, E. Creel, S. Batchu, S. Bailey, and R. Grossman, The Tele-Immersive Data Explorer: A Distributed Architecture for Collaborative Interactive Visualization of Large Data-sets, *4th International Immersive Projection Technology Workshop*, Ames, Iowa, 2000 http://evlweb.eecs.uic.edu/cavern/TIDE/tide_ipt2000.pdf
- [20] B. Schönage, *DIVA: Architectural Perspectives for Information Visualization*. Dissertation, Vrije Universiteit, Amsterdam, Netherlands, 2000.
- [21] M. Spence and C. Beilken, (1999), Discovery Challenge: Visual, Interactive Data Mining with InfoZoom—the Financial Data Set, *Workshop Notes on Discovery Challenge, 3rd European Conference on Principles and Practice of Knowledge Discovery in Databases, PKDD'99*, pp. 33-38. <http://fit.gmd.de/~cici/InfoZoom/DiscoveryChallenge/Financial.ps>
- [22] M. Spence, C. Beilken, and T. Berlage, (1996). The Interactive Table for Product Comparison and Selection, *UIST 96 Ninth Annual Symposium on User Interface Software and Technology*, Seattle, 1996, pp. 41-50. <http://fit.gmd.de/~cici/Focus/Paper/uist96.htm>
- [23] Strauss, A. and Corbin (1998). *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Thousand Oaks, CA: Sage Publications.
- [24] J. Wood, *Collaborative Visualization*, School of Computer Studies, The University of Leeds, Leeds, England, 1998.