

Capturing reasoning process through user interaction

Wenwen Dou, William Ribarsky and Remco Chang¹

¹Charlotte Visualization Center, UNC Charlotte, USA

Abstract

In recent years, visual analytics has taken an important role in solving many types of complex analytical problems that require deep and specific domain knowledge from users. While the analysis products generated by these expert users are of great importance, how these users apply their domain expertise in using the visualization to validate their hypotheses and arrive at conclusions is often just as invaluable. Recent research efforts in capturing an expert's reasoning process using a visualization have shown that some of a user's analysis process is indeed recoverable. However, there does not exist a generalizable principle that explains the success of these domain-specific systems in capturing the user's reasoning process. In this paper, we present a framework that examines two aspects of the capturing process. First, we inspect how a user's reasoning process can be captured by utilizing van Wijk's operational model of visualization. Second, we evaluate the likelihood of success in capturing a user's interactions in a visualization by introducing three criteria designed for disambiguating the meanings behind the interactions. Various visualization systems in the visualization and HCI communities are examined for the purpose of demonstrating the impact of the three criteria.

Categories and Subject Descriptors (according to ACM CCS): Information Interfaces And Presentation (e.g., HCI) [H.5.2]: User Interfaces—Graphical user interfaces (GUI)

1. Introduction

Much of the work in the field of visualization assumes a population of expert users who have knowledge and experience in analyzing problems in specific domains. In most cases, these expert users utilize visualization tools to explore data and solve domain specific tasks. The focus of the visualization, as well as the expert users, are typically on the “product” of the analysis in which the experts identify information that has not been previously discovered.

While these analysis products are of great value, we propose that the “process” of the analyses themselves also contain a great amount of knowledge. These “processes” often contain information on how the expert users identify the undiscovered, and why the expert users take the analysis steps that they do. In fact, we believe that deep and successful capturing of the reasoning process can be considered as the first step towards capturing the experts' domain knowledge and has numerous potential impacts that include training, communication, and better system designs [DJS*09, HMSA08, JKMG07].

The visualization community has become increasingly

aware of the concept that the “process” is often just as important as the “product” [GS06, SFC07]. Recently, numerous systems and applications have been published that aim to capture a user's interaction history (or sometimes referred to provenance) [SvW08, DJS*09, GZ08]. While these systems have all reported varying degrees of success, there still does not exist a set of fundamental explanations on why these systems are successful, or how others could learn from these successes and apply the techniques to their own domain. The questions we seek to answer in this paper are: how can one capture a user's reasoning process? And how can we improve the design of both interaction and visual interface to support better reasoning process capturing?

To answer the first question, we turn to van Wijk's operation of visualization model [vW05, SvW08] and examine how a user interacts with a visualization. Based on the model, we propose that there are in fact two separate modes of capturing: internal and external capturing to the visualization. Internal capturing refers to the methods of capturing within the visualization such as visualization state capturing [JKMG07] and interaction logging; whereas external capturing refers to the methods employed outside of visualization,

which includes the use of human observers or devices such as eye trackers, video camcorders, or advanced machineries (EEG (Electroencephalography) and fMRI (functional Magnetic Resonance Imaging)). We believe these two modes together represent all possible methods of capturing that are available today, but choosing the appropriate methods will depend on the goal and context in which the visualization is used.

The second question is rather broad and intricate. We recognize that the amount of reasoning that could be derived from artifacts through internal and external capturing depends on the interpreter, which could be either human coder or computer algorithms. But we argue even before the step of interpreting, making sure that much reasoning could be reflected in the artifacts is as important. Thus we further narrow the scope of the paper down to how to reflect more reasoning process in the artifacts obtained through internal capturing. The reason is two-fold: first, questions regarding using external capturing such as what to capture and how to analyze the artifacts to derive reasoning could be answered by literatures in the field of qualitative research, which describes meaning based on the collected artifacts rather than just making statistical inferences [IZCC08, Neu05]. Four basic types of data are commonly considered in the qualitative analysis: observations, interviews, documents (written artifacts) or audio-visual materials [Cre07]. The four types of data cover most of the artifacts that could be obtained through the external capturing methods, thus the qualitative research provides great guidance on external capturing. The second reason lies in the cost associated with external capturing, which we further elaborate in section 2.1.

Therefore, in this paper, we focus on investigating how to reflect a user's reasoning when employing the internal capturing method to record either the changes in the visualization state or the specific interactions by a user. Jankun-Kelly et al. have proposed the P-Set model [JKMG07] to systematically capture the visualization states during a user's exploration process, this paper therefore focus on interaction logging. Specifically, we propose that the effectiveness in capturing a user's reasoning process through interaction logging can be expressed using three criteria: the semantics encoded in the captured user interactions, information change caused by the interactions, and the visualization's interactivity. Collectively, these three criteria represent the degree of ambiguity in relating a user's interactions to the analysis process. Using these three criteria, we posit that visualizations with high interactivity, semantically rich interactions and low information change during interaction would tend to be more successful at capturing a user's reasoning process. Select visualization systems in the visualization and HCI communities are examined using these three criteria to demonstrate how these criteria relate to interaction logging and reasoning recovery.

2. A framework for reasoning process capturing

How much of an analysis process using a visualization can be captured? Clearly there is a theoretical upper bound of 100%, but intuitively that upper bound is not actually obtainable in practice. So the real questions are: how close can we get to that upper bound? And what do we have to do to get there?

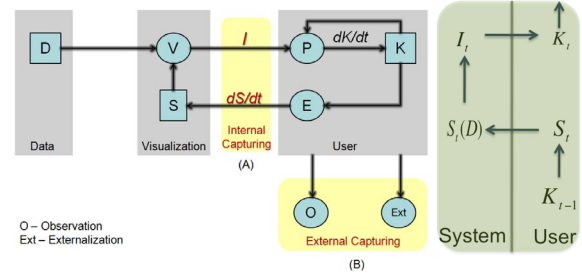


Figure 1: A framework for capturing user's reasoning process based on van Wijk's model of visualization (left), the yellow boxes (A) and (B) represent internal and external capturing methods respectively. Update of the original model of visualization (considering time) proposed by Shrinivasan and van Wijk (right).

To answer these questions, we turn to van Wijk's operational model of visualization to first understand how a user interacts with a visualization. The van Wijk operational model (Figure 1), although simple, distinctively depicts the flow and relationship between the user and the visualization. Specifically, there are two connections, I and dS/dt , between the user and the visualization. I stands for the images generated by the visualization that are perceived by the user. And the connection dS/dt represents the changes in the parameters of the visualization initiated by the user (through the use of a mouse, keyboard, or other input devices) that are applied to the visualization to generate the next sets of images I . Both of these connections can be captured directly within the visualization during user's exploration process by performing visualization state capturing and interaction logging respectively. We refer to these two methods collectively as "internal capturing" (Figure 1(A)).

In real life, however, solving a complex task is not restricted to only using a visualization. The user could jot down discoveries on a piece of paper, or watch the news on the web to gather up-to-date information. In order to fully capture a user's exploration process in solving a task, the user's activities outside of the visualization need to be captured and collected as well. We further categorize the capturing of these activities into two groups: externalization and observation. In externalization, the results that are explicitly externalized from the user of the reasoning process are collected and stored. These include the notes taken by the user during an investigation, or dictations taken using a voice

recorder. In observation, information around the user is captured through the use of additional hardware and machinery. For example, eye trackers can track the user’s focus, and a video camcorder can record the user’s activities in an environment. In addition, advanced technologies such as EEG and fMRI can be used to monitor the user’s neural activities. Together, externalization and observation are referred to as “external capturing”(Figure 1(B)).

We propose that both internal capturing (capturing I , dS/dt) and external capturing (externalization, and observation) represent a complete theoretical categorization of all reasoning process capturing mechanisms in visualization. In practice, however, the results and effectiveness will depend heavily on the implementation and accuracy of the methods as well as how the reasoning process gathered from the different methods are integrated into a cohesive story.

2.1. Internal Capturing vs. External Capturing

The advantage of external capturing lies in the numerous literatures on qualitative research to derive high-level meaning from collected artifacts. These research provide guidelines and methods for external capturing that range from what methods to use and how to analyze captured data to derive semantic meaning. However, the drawback is that such qualitative research can be very time-consuming if it requires large amounts of collected data to be analyzed and parsed [IZCC08]. Even more importantly, in most experiments involving external capturing, the fact that the analysts are externalizing their thoughts (e.g., via think-alouds), or are reminded of potential observers (e.g., in the case of being recorded on video) could change their behavior significantly. As noted by Shapiro, performing the think-alouds protocol may slow down a participant’s task performance and even alter the process of interest [Sha94]. Similarly, the use of observational tools could solicit an effect known as social facilitation and inhibition in which the participant would either over perform or under perform depending on their confidence in performing the task [ZUGH07]. Under such circumstances, internal capturing which unobtrusively captures user’s interaction and visualization states seem to be more economic and practical.

Researchers in the visual analytics community have long realized the importance role of user interactions. Pike et al. stated that in visual analytics, the development of human insight is aided by interaction with a visual interface [PSCO09]. Therefore large amount of human reasoning is embedded within user interactions which serves as a media for the dialogue between user and interface. But unlike the rich resource on the topic of qualitative research, there is less research done regarding how to embed high-level reasoning information into captured user interactions. In the next section, we mainly focus on examining how to improve internal capturing for the purpose of deriving user’s reasoning process.

3. Criteria for accessing effectiveness of user interaction capturing

As mentioned before, we certainly recognize that how well the reasoning can be derived from user interactions depends on the interpreters (either human coders or computer algorithms). However, intrinsically, the amount of a user’s reasoning that is encoded within the user’s interaction log is not affected by interpretation, but instead by the degree of ambiguity in the interaction log itself.

To better understand how a user’s reasoning relates to interaction logging, we re-examine the internal capturing aspect of our provenance model of visualization (Figure 1(A)). As shown in the model, the captured images (I) connect the visualization to the user’s perception (P), while the captured user interactions (dS/dt) relate a user’s exploration (E) to the visualization system. In this model, there is no direct link between the visualization and the user’s reasoning process (or knowledge, K). In order to reconstruct a user’s analysis process using only dS/dt and I , one must first make sure that the captured dS/dt and I relate as closely as possible to the user’s exploration and perceptual processes.

Based on this observation, we propose that visualizations that are more effective at capturing a user’s reasoning processes are in fact collecting dS/dt and I in such a way that the captured dS/dt and I can describe the user’s intention behind the interaction and the focus and interest of the user’s perception with minimal ambiguity.

By examining each access to the processes directly related to user’s knowledge in the van Wijk’s visualization model for the purpose of disambiguating the meaning of interactions, we identified 3 criteria for evaluating the effectiveness of interaction capturing in a visualization environment, namely **Semantics of user interactions**, **Information change caused by user interactions** and **Degree of interactivity**(Figure 2).

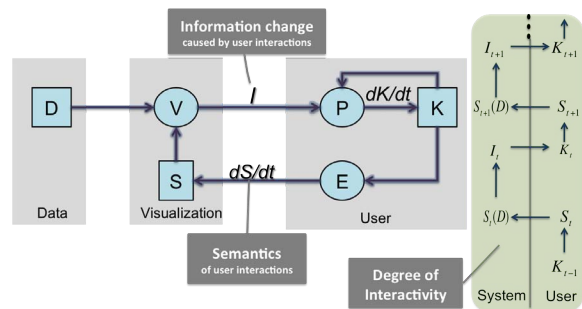


Figure 2: Three criteria identified based on van Wijk’s model of visualization. Gray boxes denote the criteria.

3.1. Criterion1: Disambiguating dS/dt – Semantics of User Interactions

In the van Wijk’s model of visualization, dS/dt is the only output from a user’s interactive exploration process (P). In order to interpret dS/dt from the user to the visualization, we need to understand the *semantics* within the interactions. The capturing of a user’s reasoning can be thought of as either capturing low-level or semantic-level user interactions. Low-level interactions can be considered as user interface events which are generated as a natural product of the normal operation of window-based user interface systems [HR00]. Semantic-level interaction reflects a user’s intention when performing interface actions [HMSA08, YKSJ07]. It has been accepted in the visualization community that in reconstructing the user’s reasoning process, low-level interaction logging is insufficient [HR00].

While the distinction between low-level and high-level interactions has been defined, Gotz and Zhou [GZ08] proposed that there exist additional categorizations of interaction types. Specifically, they characterized the user’s activities into four tiers based on their semantic richness: Tasks, Sub-Tasks, Actions and Events. The Events tier corresponds to low-level user interaction. The Actions tier relates to high-level interactions as “atomic analytic steps” such as explore, filter and zoom. The Sub-Tasks tier refers to concrete analytical goals that are tightly coupled with domain specific problems and the available features within the visualization (such as identifying trends in the financial markets). The Tasks tier categorizes the highest level of the user’s analytical goals that are often open-ended or ambiguous (such as generating financial investment recommendations).

From the perspective of capturing a user’s reasoning process, more semantic information encoded within the user’s interactions would lead to less ambiguity during interpretation. Unfortunately, as noted by Gotz and Zhou, user activities above the Actions tier are often domain specific and not easily generalizable. Most existing visualizations that provide frameworks for high-level interaction logging therefore rely on capturing activities in the Actions tier and are subsequently limited in the encoding of semantics in the user’s interactions [GS06, GZ08, HMSA08, SvW08]. When the interaction logging is more specifically coupled with clearly defined domain problem, researchers have demonstrated that high-level semantics can both be encoded in the interaction as well as extracted during interpretation [DJS*09].

3.2. Criterion2: Disambiguating I – Information Change Caused By User Interactions

We consider the effect of a user’s interaction that changes a visualization from generating an image $I(S_0)$ to $I(S_1)$ as the “information change caused by user interactions.” Intuitively, for the purpose of disambiguating user interactions, a high amount of information change is not desirable. If a

user interaction results in large amounts of information being communicated to the user all at once, it is difficult to interpret what part of the information change is perceived by the user as relevant.

We examine a few existing visualization and interaction designs based on the amount of information change. Highlighting is a common interaction technique that is used to reveal additional information about a visual object. In most cases, highlighting causes minor and specific information change that can be easily interpreted. Zooming, on the other hand, has the potential of changing the overall image I in a drastic way, but the amount of information is specific and localized. In interpreting an interaction that results in zooming, the intention behind the interaction is clear.

There are some interactions that cause high amount of information change. Animation, for example, displays a series of temporal frames given a single user interaction (such as a mouse click). Since the viewers need to keep the changing visual objects in memory for association [Lam08], visual objects with rather complex movement over a long period of time would result in high amounts of information change since the specific information relevant to the user would be lost. In this regard, complex animation such as that in Gapminder [Gap] would cause a higher amount of information change than simpler animations that are used to depict the transitioning between statistical states [HR07]. Another example of potentially high information change are the interactions within Coordinated Multiple Views (CMV). Many notable visualization systems apply the CMV interface, including Xmdv [War94], Spotfire [Ah196], etc. However, as Roberts noted [Rob07], as the number of coordinated views increases, it becomes harder for the user to keep track of the contexts and relationships between the views. In terms of information change, this means that interacting with more coordinated views will result in higher information change as the simultaneous updates in all views make it difficult to isolate the meaning and intent behind the interaction.

3.3. Criterion3 – Degree of Interactivity

The basic assumption made in the previous two criteria is that there are in fact some capturable user interactions within a visualization. However, not all visualizations incorporate the same degree of interactivity. From the perspective of capturing a user’s reasoning process, a high degree of interactivity within the visualization is preferred. Ideally, the analysis process should be driven by the user’s interactions so that there’s sufficient amount of information regarding reasoning for every step of the user’s analysis. If the visualization is more static in nature, the user’s analysis process would not manifest itself as recordable interactions, and will remain internal to the user.

Therefore we present the degree of interactivity of a visualization as the third and final criterion to capturing a user’s

reasoning process. Examples of visualization systems with low interactivity include systems in casual and ambient info-vis where no user interaction is required; whereas the other end of the spectrum is exemplified by systems that rely on the user's interactions to *drive* the visualization. For example, in the data visualization software Tableau, the user's interactions are part of the process of constructing a query in VizQL. Similarly, in ScatterDice [EDF08], the interaction controls the transition between dimensions of a scatter plot. During the transition, the animation gives rise to the user's understanding between the data and the dimensions. Without the interaction, the visualization cannot express the relationships in the data effectively.

3.4. Heuristics for efficient interaction capturing based on the 3 criteria

The three proposed criteria can be considered as independent dimensions in evaluating the effectiveness of visualizations in capturing user interactions for the purpose of reconstructing one's reasoning process. We propose that for a visualization to be effective in capturing a user's reasoning, it needs to rank highly in all three dimensions. In other words, visualizations with high interactivity, semantically rich interactions and low information change during interaction would tend to be more effective at capturing a user's analysis process. It is important to note that the three criteria do not compensate for each other in that scoring highly on two dimensions and receiving a low score on one will still render the visualization ineffective in capturing. For example, a system with high interactivity, low information change but which captures user interaction with little semantic information would only result in gathering low-level Events tier interactions [GZ08] that could not be used towards reconstructing a user's analysis process.

In the visualization community, many systems are designed with high interactivity as a core feature and would therefore rank highly under the criterion of *degree of interactivity*. However, many of them are also designed to be broadly applicable to multiple domains which limits their ability in capturing semantic-level interactions beyond the Actions tier and would therefore receive an average grade in *semantics of user interaction*. Finally, if these systems further employ interaction techniques that cause high *information change* such as multiple coordinated views or complex animation, it would further reduce their ability to capture a user's analysis process.

Case Study 1 – WireVis We specifically examine the visualization system WireVis 3 [CGK*07]. Based on the three criteria, WireVis scores highly in *degree of interactivity* as well as *semantics of user interaction* since the visual analytics system is highly interactive and the purpose of the views are clearly defined so that each interaction can be associated with specific semantics. However, WireVis employs a multiple coordinated views interface which would make it dif-

ficult for the interpreters to disambiguate the intent behind user interactions that cause *information change* in the coordinated views. Although the negative effect of *high information change* is likely to be limited in practice since WireVis only uses three coordinated views [Rob07], we nonetheless believe that certain amount of user's reasoning process may not be captured due to the high information change in WireVis system.

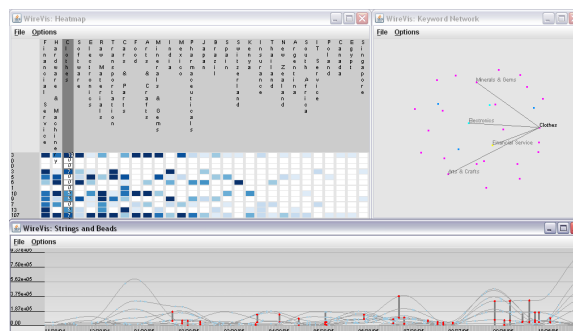


Figure 3: An overview of the WireVis system showing the heatmap (top left), keyword graph (top right), and time series view (bottom).

We compare our analysis of WireVis with the results of the study by Dou et al. [DJS*09] in which the authors reported a 60%-80% correlation between the interpretation of their captured semantic-level user interactions and the expert users' original reasoning process. Our evaluation of WireVis is consistent with the authors' report that many of the mis-correlations stem from not knowing which part of the visual change the experts were focusing on [DJS*09] due to all views getting updated at the same time, thus in support of our hypothesis that minimizing information change is beneficial to capturing a user's reasoning.

Case Study 2 – Two-stage dimension reduction testbed

The two-stage dimension reduction testbed by Choo et al. [CB09] is an efficient tool for visualizing clustered high dimensional data. Based on our criteria, it scores low in *degree of interactivity* since the system provides very limited interactions except for pull-down menus for choosing dimension reduction methods and buttons that turn on and off labels. Although the *semantics* associated with such user interactions are clear, most of the reasoning process remain in the user's head rather than carried out by interacting with the visual interface. In this case, the reasoning process that could be captured through internal capturing is limited, therefore we suggest that external capturing methods such as think-aloud or writing down notes is necessary for the purpose of deriving user's reasoning process.

4. Conclusion and Future Work

In this paper, we first propose a framework based on van Wijk's operational model of visualization to inspect how a

user's reasoning process could be captured when using a visual analytics system. Various methods for capturing a user's analysis process are categorized into internal and external capturing. Our illustration of the two categories serves as a general classification of existing capturing methods that record a user's exploration in a visualization environment. Furthermore, we present three criteria for evaluating the effectiveness of visual analytics systems in capturing a user's reasoning process. The three criteria are developed based on the characteristics of existing visualization systems along with a close inspection of the relationship between a visualization and its user as described in van Wijk's visualization model. We conclude that highly interactive visualization systems with semantically rich interactions and low information change caused by a user's interactions would likely be more effective in capturing a user's analysis process.

We discovered some interesting findings in the process of proposing our framework and the general heuristic. Through the demonstration of the available capturing methods, we are aware that even applying all capturing methods could still not record 100% of a user's analysis process. Regardless of the sophistication of the capturing techniques, a portion of the user's reasoning would always remain internal and therefore not capturable. In our future goal to reconstruct a user's analysis process through examining various captured artifacts, we are aware of this limitation but we nonetheless seek to identify new methods that would improve the accuracy of the reconstruction.

References

- [Ahl96] AHLBERG C.: Spotfire: an information exploration environment. *SIGMOD Rec.* 25, 4 (1996), 25–29. 4
- [CB09] CHOO J., BOHN S., PARK H.: Two-stage framework for visualization of clustered high dimensional data. *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on* (30 2009–Nov. 1 2009), 67–74. 5
- [CGK*07] CHANG R., GHONIEM M., KOSARA R., RIBARSKY W., YANG J., SUMA E., ZIEMKIEWICZ C., KERN D., SUDJANTO A.: Wirevis: Visualization of categorical, time-varying data from financial transactions. *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on* (30 2007–Nov. 1 2007), 155–162. 5
- [Cre07] CRESWELL J. W.: *Qualitative Inquiry and Research Design: Choosing Among Five Approaches*, second. ed. Sage Publ Inc, January 2007. 2
- [DJS*09] DOU W., JEONG D. H., STUKES F., RIBARSKY W., LIPFORD H. R., CHANG R.: Recovering reasoning processes from user interactions. *IEEE Computer Graphics and Applications* 29 (2009), 52–61. 1, 4, 5
- [EDF08] ELMQVIST N., DRAGICEVIC P., FEKETE J.-D.: Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *Visualization and Computer Graphics, IEEE Transactions on* 14, 6 (Nov.-Dec. 2008), 1539–1148. 5
- [Gap] GAPMINDER.: <http://www.gapminder.org>. 4
- [GS06] GROTH D. P., STREFKERK K.: Provenance and annotation for visual exploration systems. *IEEE Transactions on Visualization and Computer Graphics* 12, 6 (2006), 1500–1510. 1, 4
- [GZ08] GOTZ D., ZHOU M.: Characterizing users' visual analytic activity for insight provenance. *Visual Analytics Science and Technology, 2008. VAST '08. IEEE Symposium on* (Oct. 2008), 123–130. 1, 4, 5
- [HMSA08] HEER J., MACKINLAY J., STOLTE C., AGRAWALA M.: Graphical histories for visualization: Supporting analysis, communication, and evaluation. *Visualization and Computer Graphics, IEEE Transactions on* 14, 6 (Nov.-Dec. 2008), 1189–1196. 1, 4
- [HR00] HILBERT D. M., REDMILES D. F.: Extracting usability information from user interface events. *ACM Comput. Surv.* 32, 4 (2000), 384–421. 4
- [HR07] HEER J., ROBERTSON G.: Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1240–1247. 4
- [IZCC08] ISENBERG P., ZUK T., COLLINS C., CARPENDALE S.: Grounded evaluation of information visualizations. In *BELIV '08: Proceedings of the 2008 conference on BEyond time and errors* (New York, NY, USA, 2008), ACM, pp. 1–8. 2, 3
- [JKMG07] JANKUN-KELLY T. J., MA K.-L., GERTZ M.: A model and framework for visualization exploration. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (2007), 357–369. 1, 2
- [Lam08] LAM H.: A framework of interaction costs in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1149–1156. 4
- [Neu05] NEUMAN L. W.: *Social Research Methods: Quantitative and Qualitative Approaches*, 6 ed. Allyn & Bacon, Boston, September 2005. 2
- [PSCO09] PIKE W. A., STASKO J., CHANG R., OCONNELL T. A.: The science of interaction. *Information Visualization* 8, 4 (2009), 263–274. 3
- [Rob07] ROBERTS J.: State of the art: Coordinated & multiple views in exploratory visualization. pp. 61–71. 4, 5
- [SFC07] SILVA C., FREIRE J., CALLAHAN S.: Provenance for visualizations: Reproducibility and beyond. *Computing in Science & Engineering* 9, 5 (Sept.-Oct. 2007), 82–89. 1
- [Sha94] SHAPIRO M. A.: *Measuring Psychological Responses To Media Messages*. Lawrence Erlbaum Associates, 1994, ch. Think-Aloud and Thought-List Procedures in Investigating Mental Processes, pp. 1–14. 3
- [SvW08] SHRINIVASAN Y. B., VAN WIJK J.: Supporting the analytical reasoning process in information visualization. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2008), ACM, pp. 1237–1246. 1, 4
- [vW05] VAN WIJK J.: The value of visualization. *Visualization, 2005. VIS 05. IEEE* (Oct. 2005), 79–86. 1
- [War94] WARD M. O.: Xmdvtool: integrating multiple methods for visualizing multivariate data. In *VIS '94: Proceedings of the conference on Visualization '94* (Los Alamitos, CA, USA, 1994), IEEE Computer Society Press, pp. 326–333. 4
- [YKSJ07] YI J. S., KANG Y. A., STASKO J., JACKO J.: Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1224–1231. 4
- [ZUGH07] ZANBAKA C. A., ULINSKI A. C., GOOLKASIAN P., HODGES L. F.: Social responses to virtual humans: implications for future interface design. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2007), ACM, pp. 1561–1570. 3