Evaluating User Behavior and Strategy During Visual Exploration

Khairi Reda
Argonne National Laboratory
kreda@anl.gov

Andrew E. Johnson
University of Illinois at Chicago
ajohnson@uic.edu

Jason Leigh
University of Hawai'i at Mānoa
leighj@hawaii.edu

Michael E. Papka
Argonne National Laboratory
& Northern Illinois University
papka@anl.gov

ABSTRACT

Visualization practitioners have traditionally focused on evaluating the outcome of the visual analytic process, as opposed to studying how that process unfolds. Since user strategy would likely influence the outcome of visual analysis and the nature of insights acquired, it is important to understand how the analytic behavior of users is shaped by variations in the design of the visualization interface. This paper presents a technique for evaluating user behavior in exploratory visual analysis scenarios. We characterize visual exploration as a fluid activity involving transitions between mental and interaction states. We show how micro-patterns in these transitions can be captured and analyzed quantitatively to reveal differences in the exploratory behavior of users, given variations in the visualization interface.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology*

Keywords

Exploratory visual analysis; insight-based evaluation.

1. INTRODUCTION

Exploratory visual analysis (visual exploration, for short) represents one of the main use cases for interactive visualizations. Often, the primary reason for employing visualizations is to explore the data from a flexible point of view in order to recognize unexpected relationships and patterns [2, 20]. Visual exploration can be driven by the undirected observations of patterns, outliers, and salient visual features in the visualization [19], or it can be guided by hypotheses, intuition, or prior exploratory goals [7]. These

bottom-up and top-down processes often interact in complex ways [5], making visual exploration a highly fluid and emergent activity. Naturally, the manner in which this activity unfolds is likely to impact the outcome of the exploratory process, the nature of insights acquired, and ultimately the success of the visualization tool. It is therefore crucial to capture and evaluate user behavior and strategy during visual exploration, and understand how that strategy is influenced by design variations.

The visualization community has mostly focused on evaluating outcomes, by recording and quantifying insights acquired by users [9, 16, 18], for instance. However, outcome evaluation provides designers with little information on *how* users acquired such insights. Recently, there has been renewed interest in techniques for evaluating the problem-solving behavior of users engaged in visual analysis. These techniques often involve recording interaction events followed by analysis of the resulting log files in attempt to reconstruct some of the dynamics between the user and the visualization tool [11]. Although log file analysis come sometimes reveal quantitative patterns in interaction, it is often difficult to uncover reasoning processes behind these interactions. To overcome this limitation, the evaluator could consider additional sources of data, such as verbal statements collected during interviews and think-aloud experiments, as well as by observing users' viewing behavior. Qualitative analysis here could provide a richer understanding of user strategy and ultimately lead to insights on how to improve the visualization tool [8]. In ill-defined exploratory tasks, however, there can be enormous variation in the high-level strategy of users, which can be sensitive to the dataset as well as differences in analytical abilities and problem-solving styles of individuals. Ultimately, it may not be clear what strategy constitutes an optimal exploratory behavior.

One potential alternative to studying visual exploration is to look at the user's 'micro-analytic' strategy, by simultaneously examining his/her mental processes and interaction with the visualization, at a finer temporal grain. This could enable the evaluator to quantitatively characterize patterns in the reasoning processes of users, and how these mental processes are affected by variations in the interactions and visual representations employed. This short paper proposes such a methodology by extending Saraiya et al's and North's insight-based evaluation [9, 16] to capture a wider range of mental processes and interactions states collected from verbal protocols, video recordings, and log files.

2. METHODOLOGY

Our goal in this evaluation methodology is to characterize the exploratory behavior of users, and how this behavior is potentially influenced by variations in the design of the visualization tool. Ultimately, this could enable us to evaluate the effectiveness of a visualization tool on insight acquisition. However, rather than simply measuring the outcome, we want to understand behavioral patterns in how insights are acquired during visual exploration. We believe this question can be addressed at a 'micro-analytic' level. In this sense, a micro characterization of user strategy comprises patterns of transitions between 'mental' and 'interaction' states. That is, moments when the user is performing mental computation in his/her head, and moments when he/she is interacting with the visualization, hence offloading some of the information processing onto the visualization tool [10]. This way of looking at the analytic activity reflects a somewhat higher-level characterization compared to the low-level operators found by Amar et al. [1]. However, this characterization also constitute a finer-grained analysis compared to existing qualitative evaluation methods (e.g., [8]).

To capture and analyze the micro-analytic strategies of users, we propose two steps: First, we identify and code for relevant cognitive processes from verbal protocols as well as instances of user-initiated interactions from log files and/or video recordings. Here, the evaluator has latitude in defining what mental processes are relevant or interesting, though in visual analytics one is usually interested in capturing processes that indicate insight acquisition, such as making observations and generating hypotheses. Once coding is finished, the second step is to analyze the flow between mental and interactions states using transition diagrams. We describe these two steps in detail.

2.1 Coding for interactions and mental processes

It is relatively straightforward to infer instances of interaction, by directly recording those into a log file as events, or by recognizing them from video recordings or screen grabs. Since different interactions incur varying costs [6] and possibly imply different mental processes [17], we can define separate states for distinct types of interaction (e.g., brushing-and-linking, filtering, etc...). Similar interactions, however, maybe coalesced under a single state (e.g., pan/zoom and 3D rotation operations can be coalesced into a 'navigation' state).

Mental processes, on the other hand, can be inferred from verbal statements uttered by subjects during visual exploration by instructing and training participants to think-aloud during exploration [3]. Here, in addition to coding for insights (observations and hypotheses) as proposed by Saraiya et al. [16], we can also code for other mental activities that are relevant to visual exploration. For instance, *goal formation* represent instances when users set objectives for themselves intended to steer the exploratory process.

2.2 Analysis of transitions

Upon coding for insights and mental operations from the verbal protocol and interaction events from video recordings and/or log files, we combine those events into a single sequence sorted by timestamp. This sequences can be thought of as an 'extended log' that captures mental and computational processes within the combined human-computer cognitive system. From this sequence, we can analyze patterns in the transitions from interactions to insight-generating mental states. That is, we can discover which interactions are likely to lead to insight and how variations in the design of the visualization influence insight acquisition.

To visually study transitions, we construct a state transition diagram from the extended log for each participant in the study. The states in this diagram reflect the mental and interaction states established in the previous step. For each pair of consecutive events in the extended log, we record one transition between the two states that correspond to these events. For instance, if the participant formulated a hypothesis after performing a pan/zoom operation, we record a transition from the 'navigation' state to 'formulate hypothesis' state. Transitions can also be reflexive (starting and ending at the same state). For example, two distinct observations reported successively by the user are recorded as a reflexive transition from/to 'make observation'.

The resulting state transition diagram, which can be visualized using a graph representation (or an adjacency matrix), depicts transition frequency between mental and interaction states during visual exploration. Edge weights represent the number of times a transition has occurred between two states (see figure 2). This diagrammatic representation provide us with a convenient way of capturing micro-level patterns in the analytic behavior of users. Our technique produces a separate diagram for each participant in a user study. However, one may construct an 'average' diagram by averaging transition frequencies from multiple individuals and comparing them across different experimental conditions to potentially reveal systematic effects for design variations on user strategy. In this case, edge weights can also be normalized by the length of the exploratory activity to account for variations in experiment length during open-ended scenarios.

3. EXAMPLE STUDY

We illustrate this methodology with an example user study aimed at characterizing the effects of increasing the resolution and size of the display on user behavior and insight acquisition during visual exploration [14, 15]. This study revolved around the visual analysis of crime patterns in the city of Chicago. The visualization interface was based around a small-multiples design, enabling participants to explore crime patterns along multiple years, for distinct crime types, and in different parts of the city. To study the effect of screen size and resolution on user exploratory behavior, we varied the size of the visualization display.

3.1 Study design

We used a between-subjects design with a single dependent variable: The display size (Small, Large). Half of the participants undertook the study using a small display while the other half experienced the visualization environment on a large display. The study took place in the *CAVE2* environment which consists of a cylindrically shaped 18 x 4 tiled display [13]. The small display condition utilized 3 of the 18 columns available, giving participants a resolution of 4,098 x 3,072 (12.5 Megapixels) and approximately a 40-degree field-of-view. The large display condition utilized 13 columns giving participants a resolution of 17,758 x 3,072 (54.5 Megapixels) and approximately a 190-degree field-of-view. The two experimental conditions provided an
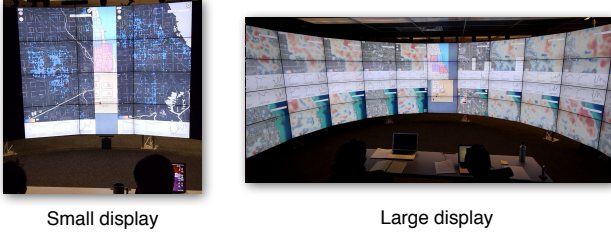
Small display   Large display

Figure 1: Two experimental conditions provided identical interactions and visual representations, but two different display sizes, which served to modulate the number of views that can be juxtaposed by users during visual exploration.

identical visualization interface. However, the size of the display served to implicitly modulate the number of views that can be viewed simultaneously; while the small display requires the participant to frequently switch between views to explore the dataset, the large display affords the juxtaposition of views side-by-side. Figure 1 illustrates these two conditions.

3.2 Procedure

Ten volunteers participated in the study (4 female). The experiment began with a 15 minutes training session with the experimenter explaining the task and demonstrating the visualization environment and its various interactive features. Participants were then given 150 minutes (2.5 hours) to explore the Chicago crime dataset, and instructed to think aloud during the activity and report interesting observations, salient patterns and outliers, correlations, trends, as well as hypotheses that explicate their observations. Participants, however, were free to end the experiment earlier if they felt that they had exhaustively explored the dataset. The session was video and audio recorded.

3.3 Analysis

Our main focus in this analysis was on the video and audio data, which contained a record of the participant's verbal protocol as well as the state of the visualization environment at the time. Since our main goal in this experiment was to study the effects of increasing the size and resolution of the display on insight acquisition, we focused primarily on capturing insight generation during visual exploration. Additionally, we were also interested in goal formation as an important part of exploratory analysis. Our coding scheme for the verbal protocol consisted of the following three categories:

- *Observation*: A unit of knowledge acquired from looking at and interacting with the visualization.

- *Hypothesis*: A conjuncture made by the participant, usually as a result of making a series of observations.

- *Goal*: A statement reflecting the formation of an exploratory objective.

These above codes comprised three mental states we are interested in: *Make observations*, *Formulate hypothesis*, and *Form goal*, respectively. As for interaction, we were interested in two general types of user-initiated actions:

- *Layout-preserving interaction*: Comprises actions that do not result in major changes to the state and layout of the visualization environment. Interaction events coded here include brushing-and-linking and panning the map.

- *Layout-changing interaction*: Comprises action that result in major changes to the visualization state, potentially requiring the participant to rebuild his/her 'mental map' [12]. This category comprised the following interactions: Creating, closing, and positioning views, changing the year or the crime type in one or more views.

3.4 Results

Combining the coded verbal protocol with the coded interaction stream into one sequences (sorted by timestamp) enabled us to create an extended activity log for each participant. This was used to create a transition diagram for each participant. For example, Figure 2 illustrates the state transition diagrams for participants S5 (who undertook the experiment with the small display) and L5 (who utilized the large display). We can see that S5 had to perform an extensive amount of layout-disruptive operations on the small display compared to L5. Moreover, we notice more transitions to the goal and hypothesis formulation states in L5's diagram, which suggests that the participant was able to devise and follow up on a larger number of exploratory goals.

participant S5   *participant L5*
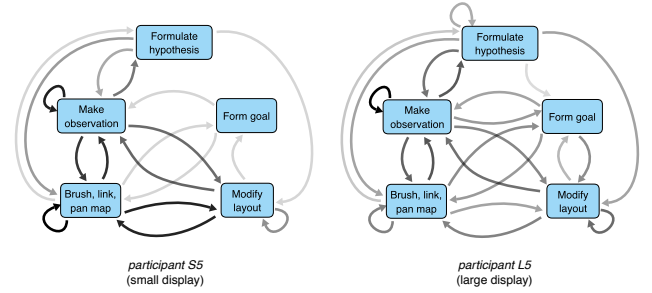(small display)   (large display)

Figure 2: Two state transition diagrams illustrating differences in strategy between participant S5 who used the small display to undertake the exploratory task (left) and participant L5, who utilized the large display (right). The weight of edges represent transition probability between two states (with log transformation applied). Thus, darker arrows represent more likely transitions.

To understand systematic variations in user behavior between the two experimental conditions (small vs. large display), we averaged the transition frequencies for participants under the same condition, giving us two 'average diagrams' corresponding to the small and large displays. To guarantee equal contributions from each participant to the average diagrams, we normalized transition frequencies by the time it took a participant to complete the activity. Figure 3 shows the two average transition diagrams side-by-side, illustrating important differences in the overall exploratory behavior of participants under the two conditions. Figure 4 highlights these variations with a 'difference' diagram and transition matrix, showing the relative changes in transition probabilities as a result of increasing the display size.
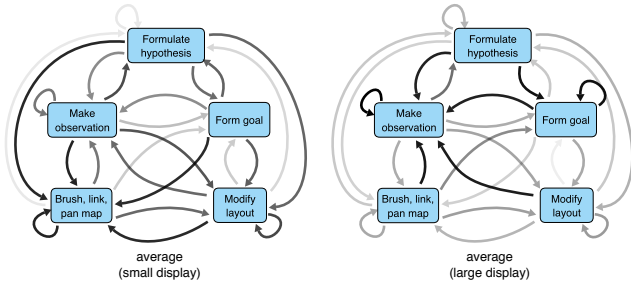
Figure 3: Two transition diagrams illustrating 'average' behavior of participants under the small (left) and large (right) display conditions.
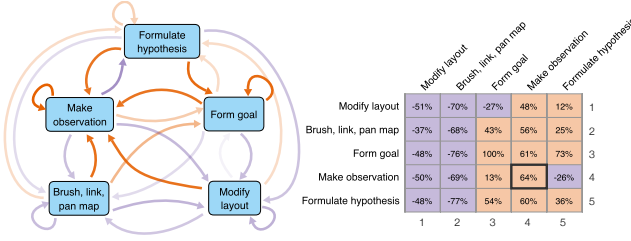

Figure 4: A transition diagram illustrating the effect of increasing the display size on the behavior of participants (left). The weight of an edge represent difference in the transition probability between the large and the small display ($P(a,b) = P_{large}(a,b)$-$P_{small}(a,b)$). Edges that have been strengthened are color coded with orange whereas weakened edges are color coded with purple. The adjacency matrix on the right shows percentage changes in transition probabilities, with outlined cells indicating significant differences between the two displays conditions ($p < .05$, Bonferroni corrected for 25 tests).

The large display diagram is marked by a decrease in the transition to the *Modify layout* state (column 1 of the transition matrix in Figure 4), indicating that participants were less likely to initiate layout-disruptive epistemic actions on the large display. We also see decreased transition probability to the *Brush, link, pan map state*, indicating that participants were also less likely to initiate brushing-and-linking and map panning operations (column 2 of the transition matrix). However, generally, we see an increased tendency for participants to transition from interaction states to insight-generating mental states with the large display (columns 4 and 5). Furthermore, we see an increased likelihood for participants to remain in these states with the large display (represented by elements (4,4), and (5,5) in the adjacency matrix). Indeed, post-hoc analysis indicates a significant increase in the probability of remaining in the *Make observation* state ($t(8) = 4.995, p < .001$).

Overall, the above variations suggest that the large display was more effective at eliciting insights and keeping participants in the 'cognitive zone' [4], where they are likely to continue to generate additional insights. Furthermore, we see a slight tendency for users to transition more frequently to the *Form goal* state, which suggests that they are able to set more ambitious exploratory goals for themselves. One pos-

sible explanation for this is that cognitive costs involved in pursuing such goals is reduced with the large display, given the ability to juxtapose and simultaneously consult a larger number of views.

4. CONCLUSIONS AND FUTURE WORK

In this paper we proposed a methodology for evaluating user behavior and strategy in open-ended visual exploration scenarios. Our techniques, which can be considered an extension of insight-based evaluation methodologies, characterizes the micro-analytic strategy of users by considering transitions between mental and interaction states during visual exploration. We demonstrated this methodology with an example study on the effects of increasing the size and resolution of the visualization interface. We believe the techniques is effective at revealing quantitative differences in user behavior given variations in the physical properties (e.g., form factor, size, and resolution) or the design of the visualization tool. We also illustrated how coded event sequences from verbal protocols and interaction logs can be combined, modeled and visualized using state transition diagrams.

In the future, we would like to further investigate the possibility of analyzing high-level patterns in the resulting state transition systems. In particular, we would like to pose and investigate the following questions:

- What components of the high-level visual analytic activity can we (automatically) deduce from the combined analysis of verbal protocols and interaction logs?

- Can we objectively analyze, compare, and score the high-level strategy of users given this data?

Acknowledgments

References

[1] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *Information Visualization, 2005. INFOVIS 2005. IEEE Symposium on*, pages 111–117. IEEE, 2005.

[2] W. S. Cleveland. *The Elements of Graphing Data*. Wadsworth Advanced Books and Software, 1985.

[3] K. Ericsson and H. Simon. *Protocol analysis: Verbal reports as data*. MIT Press, revised edition, 1993.

[4] T. Green, W. Ribarsky, and B. Fisher. Building and applying a human cognition model for visual analytics. *Information Visualization*, 8(1):1–13, 2009.

[5] C. G. Healey and J. T. Enns. Attention and visual memory in visualization and computer graphics. *Visualization and Computer Graphics, IEEE Transactions on*, 18(7):1170–1188, 2012.

[6] H. Lam. A framework of interaction costs in information visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 14(6):1149–1156, 2008.

[7] Z. Liu and J. T. Stasko. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):999–1008, 2010.

[8] E. Mayr, M. Smuc, and H. Risku. Many roads lead to rome: Mapping users' problem-solving strategies. *Information Visualization*, 10(3):232–247, 2011.

[9] C. North. Toward measuring visualization insight. *Computer Graphics and Applications, IEEE*, 26(3):6–9, 2006.

[10] P. Parsons and K. Sedig. Distribution of information processing while performing complex cognitive activities with visualization tools. In *Handbook of Human Centric Visualization*, pages 693–715. Springer, 2014.

[11] M. Pohl, S. Wiltner, and S. Miksch. Exploring information visualization: describing different interaction patterns. In *Proceedings of the 3rd BELIV'10 Workshop: BEyond time and errors: novel evaLuation methods for Information Visualization*, pages 16–23. ACM, 2010.

[12] H. C. Purchase, E. Hoggan, and C. Görg. How important is the "mental map"?–an empirical investigation of a dynamic graph layout algorithm. In *Graph Drawing*, pages 184–195. Springer, 2007.

[13] K. Reda, A. Febretti, A. Knoll, J. Aurisano, J. Leigh, A. Johnson, M. Papka, and M. Hereld. Visualizing large, heterogeneous data in hybrid-reality environments. *Computer Graphics and Applications*, 33(4):38–48, 2013.

[14] K. Reda, A. Johnson, V. Mateevitsi, C. Offord, and J. Leigh. Scalable visual queries for data exploration on large, high-resolution 3d displays. In *SC Companion: High Performance Computing, Networking, Storage and Analysis*, pages 196–205. IEEE, 2012.

[15] K. Reda, C. Offord, A. Johnson, and J. Leigh. Expanding the porthole: leveraging large, high-resolution displays in exploratory visual analysis. In *CHI'14 EA on Human Factors in Computing Systems*, pages 2047–2052. ACM, 2014.

[16] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *Visualization and Computer Graphics, IEEE Transactions on*, 11(4):443–456, 2005.

[17] K. Sedig and P. Parsons. Interaction design for complex cognitive activities with visual representations: A pattern-based approach. *AIS Transactions on Human-Computer Interaction*, 5(2):84–133, 2013.

[18] L. Shupp, C. Andrews, M. Dickey-Kurdziolek, B. Yost, and C. North. Shaping the display of the future: The effects of display size and curvature on user performance and insights. *Human–Computer Interaction*, 24(1-2):230–272, 2009.

[19] A. Treisman. Features and objects in visual processing. *Scientific American*, 255(5):114–125, 1986.

[20] J. W. Tukey. We need both exploratory and confirmatory. *The American Statistician*, 34(1):23–25, 1980.