Logging Interactions to Learn About Visual Data Coverage

Tanja Blascheck*, Steffen Koch*, and Thomas Ertl*

*University of Stuttgart, Germany

Abstract

Evaluating interactive visualizations requires examining the complex interaction techniques used in the visualization, but also makes it necessary to investigate which part of the data participants are exploring. In this paper, we discuss how logging of these complex interaction techniques may help people to understand the data explored in order to improve a visualization technique. By answering questions about how much, which part, or how often a part of the data was inspected, we can infer valuable information about the usefulness and effectiveness of an interactive visualization technique. This visual data coverage allows us to also make inferences about traceability and accountability of a visualization technique. We present experiences we made during our studies, discuss challenges, and point out future directions of this work.

1 Introduction

To increase insight generation based on interaction logs in the context of modern, interactive visualizations, it is valuable to also log the extent and kind of data surveyed with continuous interactions. This information about *visual data coverage* brings new potential to the *evaluation* of a visualization and its interaction design, *tracking* which portion of data a user inspected visually at a specific level of detail and suggestions on what to investigate next, as well as *accountability* of data analysis and insights. Therefore, we focus on these three topics and in particular on closely related questions regarding data coverage during visual analysis: how much of the data was explored, how often was the same data item inspected, which parts of the data were examined at which level of detail, and was the right data investigated?

In order to analyze visual data coverage, logging of this information is required. There is extensive work on analyzing interactions [18, 23, 4] and many approaches exist on how to categorize interactions [4, 23] or interaction costs [14]. We discuss the benefits and limitations of logging visual data coverage when using complex interaction techniques in context of evaluation, traceability, and accountability. We illustrate our experiences, insights, and challenges when collecting information about the data explored with a focus-and-context lens and share some preliminary ideas about how to analyze this data.

2 Evaluation

Interactive visualizations are developed to explore data to find insights in the represented data. When developing an interactive visualization method, a designer makes design decisions. Evaluating a novel visualization technique requires finding evidence if users can find insights in the data while applying a technique. Thus, the goal of an evaluation is to analyze if the design decisions made were the right ones and can be understood by users to generate insights from the data.

Analyzing a visualization with complex interaction techniques, like focus-and-context techniques, requires looking at both how users interact with the technique and how much data is explored visually at which level of detail. Recording the visual data coverage during a study, for example, by logging each data item investigated can be used to answer our questions. If the amount of data (how much data has been seen) is low, this can indicate issues with the efficiency of chosen interaction and visualization techniques. For example, some parts of the data might not be accessible with a certain visualization technique. Which part or which data items participants inspected can give insights about accessibility of the data. If an attention guiding method is used, for example, a heatmap indicating dense regions that are in particular interesting for close inspection, analyzing which part of the data was examined might also give insights about how well this attention guiding worked. If the heatmap led users to explore interesting parts of the data, this information can be inferred from visual data coverage. If a count is recorded each time a participant explored a data item in detail, an analyst can see which data items got more attention and which received less. This may indicate that some parts of the data were more interesting, more in focus, or more available. Another important question is whether or not a user examined the right data. Depending on the task, some part of the data may be more valuable or even necessary to complete a task successfully. This requires defining a ground truth from data items that need to be investigated for each task if this is to be assessed during evaluation. With this information it is at least possible to determine whether a user explored data important for solving a task or not.

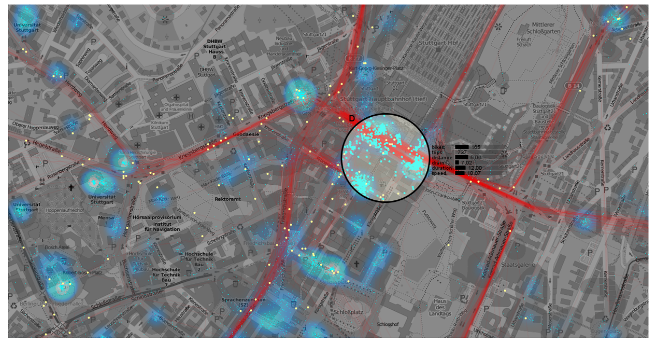

Figure 1: TrajectoryLenses is a Visual Analytics system allowing a user to explore trajectories displayed on a map using a focus-and-context lens. The heatmap shown depicts the end points of movement trajectories to guide the user in exploring the data.

Despite this, only taking into account the visual data coverage might not be enough to make clear if participants actually perceived or understood all the data inspected during exploration. So far, we have only looked at the visual data coverage logged during a study. For example, a high visual data coverage or a thorough investigation of one data item might be the result of the logging mechanism. Thus, just looking at this metric might not be enough and other parameters or data sources have to be combined to make better inferences. Eye tracking could be one means to investigate if participants actually explored a specific data item for a sufficient amount of time, or a user's attention passed over a visual detail without realizing its importance. Interaction logs have been correlated with other data sources such eye tracking [2, 3, 1, 19] or think-aloud

*e-mail: firstname.lastname@vis.uni-stuttgart.de

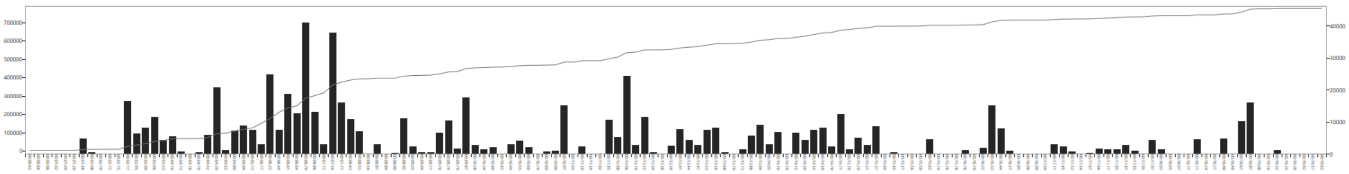

Figure 2: Visual data coverage graph showing the amount of data a user has visually explored over time (bar chart) as well as aggregated over the complete task duration (line chart). Note that the graph shows two different granularities: the overall visual data coverage (about 7,000,000 data items) and the visual data coverage of one time step (maximum about 42,000 data items).

protocols [7, 16, 10]. However, these approaches do not consider the actual data and visual level of detail inspected.

3 Traceability

Taking our idea one step further, visual data coverage analysis can also be used for tracing user interactions and guiding the user. Analyzing the traceability of the insight process has been proposed by Gotz and Zhou [9]. Their approach focuses on low-level versus high-level actions, where low-level actions are basic interactions. They do not consider the data that a user has investigated. Based on our questions, we can ponder on how visual data coverage information can be used in applications that continuously log user interaction. If the recorded data information is available immediately, it can be analyzed to guide the user to new data items or parts of the data that was not inspected visually yet. The information about how much of the data a user has explored can be used to make the user aware that s/he has only examined a small part or already investigated all available data. The part of the data a user has explored can be used to guide a user to parts of the data not analyzed yet. Additionally, if some ground truth is available for a task, a user can be guided to inspect the right data when solving a task.

For example, a visualization as shown in Figure 3 is useful to let users understand which parts of the data they explored in detail and which might deserve additional attention. The time spent for inspecting particular data items may be plausible if the interactive visualization is well designed and helped to steer users attention to exactly those details in the data that are interesting in the context of the task to be solved. With often underspecified tasks and explorative approaches, as supported by visual analytics solutions, this perfect match might not be guaranteed. Tracing the level of detail, the frequency in which data items were inspected as well as the quantity of explored data items and reflecting this to users, might help to prevent severe errors in assessing a situation. In the given example, the bar graph shows that some elements have not been investigated at all, yet. Extending this visualization to show all data items might help the user to see which parts were not inspected. Adding brushing methods and linking the visualization with the visual analytics tool that was used for exploration, for example, by clicking on an item to highlight the corresponding elements, could be a valuable means to reduce analysis faults from overlooking important information.

4 Accountability

Accountability has been discussed so far in different context. It has been used to describe 'truthfulness' of visualization approaches [21] which is inevitably hampered by data uncertainties and the impossibility of reflecting the full complexity of the real world. Reducing this problem certainly includes a critical reflection of a users analysis process which can be supported by improving traceability. The term accountability is also used with respect to security and privacy, for example, by Butin [6] and Weitzner [22] who suggest to make business transactions verifiable and dishonest or even illegal use of information transparent. While these aspects are not directly covered by our discussion, it would be interesting to indicate them if information foraging is supported through interactive visual interfaces. Many other semantic meanings of accountability exist in the context of HCI and other research fields [8]. Accountability is of special importance in decision-making situations like visual analytics accountability [15, 20]. If an expert using a tool is misled by the data and comes to a wrong decision, the question of who is responsible becomes an issue. Logging interaction data and visual data coverage might be one option to overcome such an issue. It may become clear from the logging data that an expert has not inspected the data thoroughly enough. For example, how much data an expert investigated can indicate if the expert has examined all data or just parts. If we can infer which parts of the data s/he explored it may demonstrate if just obvious parts, or all parts of the data were inspected. Knowing which data must be explored at least to make valid inferences can again help to know if the right part of data was examined. However, all of this information may also help to investigate if an expert was not able to make a valid decision because s/he was misled by the visualization or could not get all appropriate data necessary.

5 Example of Evaluation using Visual Data Coverage

To show how visual data coverage information can be used when evaluating an interactive visualization, we give an example from a recent study we conducted. In this study, we collected interaction data as well as data that a participant inspected while using a focus-and-context lens. The interaction data was collected by instrumenting the analyzed system. We analyzed the visual analytics system called TrajectoryLenses [13] consisting of a focus-and-context lens which can be used to explore trajectory data displayed on a map (cf. Figure 1). We recorded the data with a frequency of 60 Hz to achieve a sufficient sampling rate of the data and to be able to synchronize it with eye movement data we recorded as well. In our case, we collected an ID, timestamp, and the position of the lens, all in addition to the data that was currently depicted underneath the lens. We achieved this through a hit test with the data items' positions in order to save its ID. We saved all IDs as a string attached to the current interaction. Following this, we calculated how many data items were investigated at each time step and accumulated the data for each unique data item that was examined over time into an overall visual data coverage amount. Additionally, we calculated how often a participant explored each data item by counting how often the ID was recorded.

Since we do not focus on analyzing and visualizing interaction logs, many of the proposed methods are only partially useful. Typically, a timeline is used which depicts interactions either as thumbnails [11, 9] or as color-coded glyphs [7, 12]. Transition matrices or transition charts depicting the transitions between different interaction categories [17, 10, 5] are a different approach, however, we are more interested in how to represent the visual data coverage, rather than individual interactions. Thus, we present two ideas on how to depict this data.

Figure 2 shows the visual data coverage graph for one participant and data visualized at a specific level of detail. On the x-axes we depict the time and on the two y-axis we depict the overall data amount as well as the visual data coverage for each time step indi-

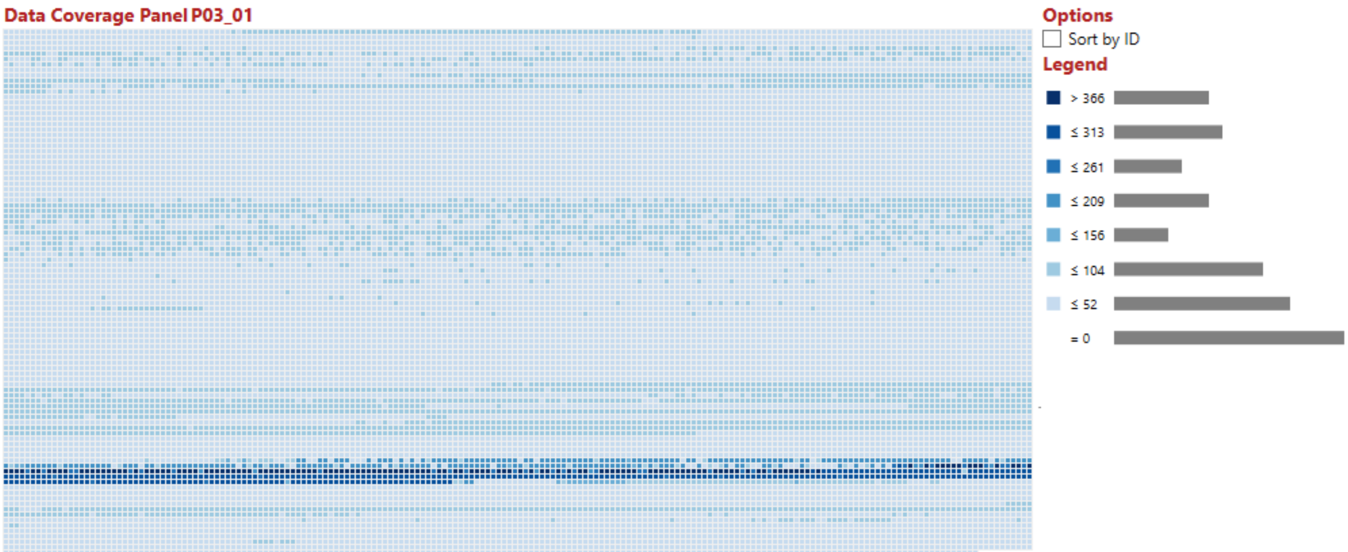

Figure 3: The visual data coverage of one participant where each rectangle represents a data item. The data items are sorted based on the time they were examined for the first time. The color corresponds to the number of times a data item was inspected, the darker the more it was explored. In the legend a bar chart depicts how may data items belong to each range.

vidually. The line in the chart represents the overall amount of data a participant investigated over time. The bar charts indicate the amount of detail data explored interactively over time during the study. We can see that at the beginning of the study, this participant examined a lot of data indicated by the bars on the left of Figure 2. The number of overall data items the participant inspected in detail was 774,141 data items and the maximum number of data items investigated at one time step was 41,070. This provides us with information about how much of the data this participant analyzed.

Additionally, we created a simple chart for showing how often each data item was inspected. Figure 3 shows the data from the same participant. The color of each rectangle represents how often a data item was explored with the lens-based technique using a binning technique. We can see that a few data items have been examined thoroughly, as the dark blue rectangles at the bottom of Figure 3 show. The bar chart on the right next to the ranges indicate how many data points have been investigated in each range. In this case, there is also a large amount of data that has not been inspected in detail (see Figure 3; lowest bar next to = 0). If this part of the data, which a participant did not explore, contains valuable information, a hypothesis could be that the visualization system was not developed appropriately to guide the user. However, if the system was developed trying to guide the user to parts of the data which are of high interest, having a large amount of data not being examined, can be a positive result as well if a user did not inspect unnecessary information. With this chart, we can infer which part of the data and how often each data item was investigated.

6 Conclusion

In this paper we have explored the idea of analyzing the visual data coverage while using an interactive visualization. We have shown how analyzing the amount and parts of the data represented can give valuable insights into how well a visualization technique was designed, and what we can infer from this data regarding traceability and accountability. We indicated which insights can be gained from logging information as well as the visually represented data as an effect of logged interactions. Despite this, there are still many open issues on how to track and analyze this kind of data. We believe that analyzing a user's visual data coverage may help to guide to exploring parts of a visualization more closely in the future. Traceability may help in real time to inspect important and interesting data more appropriately, and from the interaction logs we can get insights with respect to accountability if questionable decisions are made based on visual analysis.

References

[1] F. Beck, T. Blascheck, T. Ertl, and D. Weiskopf. Exploring word-sized graphics for visualizing eye tracking data within transcribed experiment recordings. In *ETVIS*, 2015.

[2] T. Blascheck, M. John, S. Koch, L. Bruder, and T. Ertl. Triangulating user behavior using eye movement, interaction, and think aloud data. In *Proceedings of the Symposium on Eye Tracking Research & Applications*, pages 175–182. ACM, 2016.

[3] T. Blascheck, M. John, K. Kurzhals, S. Koch, and T. Ertl. VA^2: A visual analytics approach for evaluating visual analytics applications. *IEEE Transactions on Visualization and Computer Graphics*, 22(01):61–70, 2016.

[4] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013.

[5] E. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding waldo: Learning about users from thier interactions. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1663–1672, 2014.

[6] D. Butin, M. Chicote, and D. L. Mtayer. Log design for accountability. In *Security and Privacy Workshops (SPW)*, pages 1–7. IEEE, 2013.

[7] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. Lipford, and R. Chang. Recovering reasoning processes from user interactions. *IEEE Transactions on Computer Graphics and Applications*, 29(3):52–61, 2009.

[8] S. Eriksén. Designing for accountability. In *Proceedings of the Second Nordic Conference on Human-computer Interaction*, pages 177–186. ACM, 2002.

[9] D. Gotz and M. X. Zhou. Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009.

[10] H. Guo, S. Gomez, C. Ziemkiewicz, and D. Laidlaw. A case study using visualization interaction logs and insight metrics to understand how analysts arrive at insights. *IEEE Transactions on Visualization and Computer Graphics*, 22(01):51–60, 2016.

[11] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala. Graphical histories for visualization: Supporting analysis, communication, and evalua-

tion. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1189–1196, 2008.

[12] M. R. Jakobsen and K. Hornbæk. Fisheyes in the field: Using method triangulation to study the adoption and use of a source code visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1579–1588. ACM, 2009.

[13] R. Krüger, D. Thom, M. Wörner, H. Bosch, and T. Ertl. Trajectorylenses - a set-based filtering and exploration technique for long-term trajectory data. *Computer Graphics Forum*, 2013(3):451–460, 2013.

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

[15] V. L. Lemieux and T. Dang. *Building Accountability for Decision-Making into Cognitive Systems*, pages 575–586. Springer, 2013.

[16] H. Lipford, F. Stukes, W. Dou, M. Hawkins, and R. Chang. Helping users recall their reasoning process. In *IEEE Symposium on Visual Analytics Science and Technology*, pages 187–194. IEEE, 2010.

[17] Z. Liu and J. Heer. The effect of interactive latency on exploratory visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2122–2131, 2014.

[18] W. Pike, J. Stasko, R. Chang, and T. O'Connell. The science of interaction. *Information Visualization*, 8(4):263–274, 2009.

[19] M. Pohl and F. Scholz. *How to Investigate Interaction with Information Visualisation: An Overview of Methodologies*, pages 17–29. Springer, 2014.

[20] A. C. Robinson and C. Weaver. Re-visualization: Interactive visualization of the process of visual analysis. In *Workshop on Visualization, Analytics & Spatial Decision Support at the GIScience conference*, 2006.

[21] T. R. Wallace and C. van den Heuvel. Truth and accountability in geographic and historical visualizations. *The Cartographic Journal*, 42(2):173–181, 2005.

[22] D. J. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, and G. J. Sussman. Information accountability. *Commun. ACM*, 51(6):82–87, 2008.

[23] J. S. Yi, Y. ah Kang, J. Stasko, and J. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, 2007.