

Characterizing Exploratory Behaviors on a Personal Visualization Interface Using Interaction Logs

Poorna Talkad Sukumar¹, Gonzalo J. Martinez¹, Ted Grover², Gloria Mark², Sidney K. D'Mello³, Nitesh V. Chawla¹, Stephen M. Mattingly¹, and Aaron D. Striegel¹

¹University of Notre Dame; ²University of California, Irvine; ³University of Colorado Boulder




Figure 1: Screen captures of the personal visualization interface. The left image presents the original layout designed for viewing on large screens. The right image presents an adapted (scrollable) design for viewing on mobile devices.

Abstract

Personal visualizations present a separate class of visualizations where users interact with their own data to draw inferences about themselves. In this paper, we study how a realistic understanding of personal visualizations can be gained from analyzing user interactions. We designed an interface presenting visualizations of the personal data gathered in a prior study and logged interactions from 369 participants as they each explored their own data. We found that the participants spent different amounts of time in exploring their data and used a variety of physical devices which could have affected their engagement with the visualizations. Our findings also suggest that the participants made more comparisons between their data instances than with the provided baselines and certain interface design choices, such as the ordering of options, influenced their exploratory behaviors.

1. Introduction

Interaction logs provide a valuable means for studying user behaviors with visualizations, especially in scenarios where it is not feasible to gather any other form of input from users. The visualization literature reveals a diverse usage of interaction logs. User interactions can be analyzed, for example, to understand problem-solving behaviors in visual analytics [DJS*09] and to capture finer facets of user exploration with interactive visualizations on the

web [FPH18]. However, the utility of interaction logs in personal data contexts has not received attention thus far. In this paper, we look at how low-level interaction metrics can be used to characterize exploratory behaviors with personal visualizations where users interact with their own data and on their own devices, potentially to self-reflect and draw insights [HTA*14]. Specifically, we aim to define interaction metrics and patterns that are meaningful in personal data contexts and outline user behaviors which can contribute to an empirical understanding of personal visualizations.

We leveraged the *Tesserae* study [MGA*19], a large-scale, year-long study where various personal data attributes, such as physical activity and sleep, of 757 information workers across five cohorts/organizations were tracked through wearables and Bluetooth beacons. We designed *The Tesserae Personal Data Explorer*, shown in Figure 1, to present visualizations of personal data corresponding to each participant, as well as aggregated data corresponding to the participant's cohort, and aggregates of all the participants in the study to serve as suitable baselines for comparisons. We sent unique, anonymized links to non-local participants (564) to view their respective personal visualizations and logged their interactions over a period of 18 days.

From those 564 invitees, 369 participants accessed their visualizations during this time. The participants spent different amounts of time in exploring their visualizations and many also returned to their visualizations. The participants used a variety of physical devices which could have affected their engagement with the visualizations. We extracted interaction patterns corresponding to potential comparisons made which suggest that the participants made more comparisons between instances of their personal data than with the provided baselines.

The contributions of this work are the following: (1) results from a longitudinal, realistic study of personal visualizations using interaction logs; and (2) interaction metrics and patterns useful for empirical evaluations of personal visualization.

2. Related Work

Personal Visualizations. Interactive visual representations of tracked personal data fall under the class of personal visualizations. Personal visualizations enable people to reflect upon their data and make meaningful inferences about themselves, potentially leading to changes in behavior [HTA*14, CL*15, PTT*15]. Research on the use of such visualizations and visual tools exist across many disciplines, such as ubiquitous computing and human-computer interaction. Huang et al. [HTA*14] coalesce this previous work to formally define and delineate personal visualizations within the InfoVis domain. To study personal visualizations, we should consider their distinctive goals and ensure that we study people interacting with their own data and in realistic settings [TLCC17].

To design personal visualizations, it is important to understand how people gather insights from visual representations of their data. A common insight-gaining pattern involves users forming questions or hypotheses about their data and then finding answers in the visualizations [CLZ*17, LDF11]. Choe et al. [CL*15] present types of visualization insights that users identify as they reflect on their personal data by analyzing quantified selfers' video presentations. These insight types are further refined by Choe et al. [CLZ*17] through an evaluation of their personal visualization application, *Visualized Self*. We referred to the gathered insight types to design our personal visualization interface.

Characterizing User Behaviors from Their Interactions. User interactions with visualizations come packed with a wide variety of information which can serve numerous purposes. For example, they can help uncover the strategies of and insights gained by analysts [DJS*09, GGZL15, BH19]. They can also be good indicators of task performance and personality traits of users [BOZ*14, OYC15].

Interaction metrics describing user behaviors can be defined depending on the contexts of use. Boy et al. [BDF15] define low-level metrics based on user intents for a comparative evaluation of two visualization designs. Interaction metrics have also been defined to measure cognitive bias in visual analytics [WBF17] and to capture more fine-grained exploration behaviors [FPH18]. While these metrics have been applied to characterize user interactions with the same data visualization, we look at personal visualizations where each user interacts with the same type of visualization but with different data.

3. Design of The Tesserae Personal Data Explorer

We present below the key considerations in designing our large-scale personal visualization interface (shown in Figure 1).

Data. For each participant, the interface presents weekly and monthly summaries of six different features corresponding to the participant (labeled as "You") and those of the baselines. The features are activity, sleep, heart-rate variability (HRV), hours spent at home, hours spent at office, and number of steps. Weekly summaries for each feature consist of data points representing average values for each week in the study. Similarly, monthly summaries consist of values averaged over each of the enrolled months. The amount of data visualized for a participant depended upon their compliance with protocols and their duration in the study.

Baselines. People generally require an appropriate reference or baseline for interpreting their personal data [HTA*14]. We provide two baselines for each participant - the aggregated data corresponding to the participant's cohort (labeled as "Cohort") and all the participants in the study (labeled as "All"). Since the number of participants in the Cohort and All groups varied during the course of the study, to enable informed comparisons, we also visually encoded the number of participants in the respective groups by coloring the data points using sequential color scales.

Line graphs. The tracked time-series personal data were visualized using line graphs and the visualizations were developed using D3.js [D3.20]. Line graphs are commonly used to visualize self-tracking data and hence many of the insight types found in personal data contexts [CL*15, CLZ*17] are also visual insights associated with line graphs [YLZ14], such as, identifying values, identifying extremes, comparing values, and finding trends over time.

Explanatory text. Personal visualizations are very likely to be used by people who are not visualization experts or have less experience with visualizations [HTA*14, CLZ*17]. Hence we included explanatory text alongside each type of visualization describing the visual encodings, the data points, and how they were computed.

Interaction affordances. Participants could hover over the data points to see details including x and y values, and also the number of participants in the case of Cohort and All data. Participants could select between the two types of summaries (weekly [default] vs. monthly) and among the six feature types. We logged all the hover and click interactions (or corresponding touch interactions in the case of touchscreen devices) of participants. We used AJAX to send details of the mouse/touch events as JSON objects to a PHP script which then logged them on the server.

Adapting design to mobile devices. The interface was adapted to

Table 1: A summary of the user interaction sessions categorized by the device types used in the sessions.

Device types used in sessions	Number of sessions	Total time of sessions	Median time per session	Number of abstract/elaborate operations	Number of explore operations	Number of identify operations	Number of inspect operations
1. Mobile devices (form factors $\leq 414 \times 833$ px)	83	7,713.28s	66.58s	99	562	21	26
2. Large touchscreen displays (form factors $\geq 1280 \times 529$ px)	171	19,354.98s	70.12s	152	1,080	1,495	2,520
3. Large displays with mouse interaction (form factors $\geq 1095 \times 533$ px)	335	43,983.12s	91.07s	258	1,608	3,417	4,460
4. Mixed (A combination of the above types)	5	1,089.53s	145.34s	7	39	43	27

be accessed on mobile devices with varying form factors (see Figure 1). Nevertheless, the participants were informed that the visualizations would be best browsed on a laptop/desktop device.

4. Metrics

To apply low-level interaction metrics to describe user behaviors on the interface, we categorize the interactions by user intents in performing them. We refer to the interaction categories presented by Yi et al. [YaKS07] and Boy et al. [BDF15] to create four categories to describe the semantic operations users could perform with their personal visualizations:


- **identify:** show me specific values of my data [*hover/touch data points on the "You" line graph*]
- **inspect:** show me specific values of baseline data [*hover/touch data points on the "Cohort" or "All" line graphs*]
- **explore:** show different set of data [*click/touch a feature type*]
- **elaborate/abstract:** show more or less detail [*click/touch a summary type*]

We define *sessions* of interaction for each participant and determine the time spent and analyze interactions within each session. We determine the start of an additional session for a participant when two consecutive timestamps in their interaction log are separated by more than five minutes. We also recorded the types of devices on which the visualizations were accessed for each user session.

5. Results and Discussion

We analyzed participant interactions logged for 18 days from when they received the links to their visualizations. 369 out of the 564 participants accessed their visualizations during this time. Our key findings are presented below.

5.1 Exploration Time, Revisits, and Reading Time. There was quite a bit of variation in the total amount of time people spent on the interface. While the median exploration time was 90.61 seconds, the exploration time ranged from less than 1 second to 1473.9 seconds among the participants as shown in Figure 2. There could be numerous factors contributing to the exploration-time differences including differing goals and expectations of participants [HTA*14], the devices on which the visualizations were accessed (discussed next), and the variation in the amount of data visualized for each participant.

**Figure 2:** Histogram of the total exploration time of participants.

We also investigated revisits in the interaction logs. Personal visualizations are characterized by higher engagement levels [CL*15] and revisits provide a good means to gauge user engagement [BDF15, RHF10]. In the 594 sessions recorded from the 369 participants, 89 participants had more than two sessions and 32 participants had more than three sessions. Exploration time and revisits could also be good indicators of "recall", a key insight type in personal data contexts where people remember background information to understand their data [CLZ*17].

On devices using mouse interaction, we recorded the time when the mouse was present in the description section to infer if the participants read the explanatory text [BDF15]. In the 335 sessions with mouse interaction, participants spent a median time of 0.39 seconds per session in the description section which contained an average of 240 words per visualization. Assuming an average user reading score of 200-250 words per minute [BDF15], it is plausible that many participants did not read the explanatory text for any visualization in its entirety. It has been found that people generally look to other textual descriptions in visualizations when they find the titles to be insufficient [BBK*15].

5.2 Device Types. While traditional visualizations and visual analytics are generally studied in environments with workstations and large screens, personal visualizations may be accessed by people on a variety of devices [HTA*14]. This aspect is reflected in our study where, despite nudging the participants towards using large screens, they used various devices to access the visualizations. The display type and form factors of the devices could have contributed to the participants' engagement with the visualizations [MKK15].

A summary of the user sessions with the different device types is presented in Table 1. We can see that the numbers of *identify* and *inspect* operations performed in the mobile-device sessions are distinctly lower than those in sessions using large screens. These differences could be attributed to the affordances of the different device types. Large screens present a complete view of the layout and enable users to perform more freeform exploration. Touch interactions can be cumbersome and slow, especially on mobile devices, compared with mouse hover interactions.

37 participants used different devices in separate sessions and there were five "mixed" sessions (see Table 1) where the participants used multiple devices within the same session. In four of the mixed sessions, the participants simultaneously accessed their visualizations using a mobile device and a desktop. These participants also spent more time exploring the visualizations in these sessions. Mobile devices were almost always used in portrait orientation and large screens were almost always used in landscape orientation. Our findings reaffirm that people are very likely to use mobile devices for viewing personal visualizations and hence future work should also consider designing such visualizations specifically for use on mobile devices [BLIC18, SH14].

5.3 Comparative Insight Patterns. Comparisons are an essential means to gather insights from data [HTA*14]. We extracted patterns of semantic operations from user interactions to make inferences about the possible comparisons made by participants [CLZ*17]. Similar to Guo et al. [GGZL15], we define a pattern as a short sequence of consecutive semantic operations commonly occurring across all participant interactions.

- **Comparing measured values segmented by time.** People generally compare their personal data measured at different times [CLZ*17]. To make inferences about such comparisons, we searched for sequences of two *identify* operations performed on the same type of data but with distinct time values. We found a total of 2,343 such sequences across all participant interactions.
- **Comparing against external data.** People also tend to compare their data with that of other people or their peers [CL*15]. To infer possible comparisons where participants compared their data with both of the baselines, we searched for sequences containing any of the six combinations of an *identify* operation, an *inspect* (Cohort) operation and an *inspect* (All) operation performed on the same type of data and with the same time values. To infer possible comparisons with either type of baseline data, we looked for sequences with an *identify* operation and an *inspect* operation, in either order, performed on the same type of data and with the same time values. We found a total of 678 sequences of which 182 pertained to comparisons with both the Cohort and All data and 496 to comparisons with either the Cohort or All data (but not both).

Consistent with prior findings [CL*15, CLZ*17], we found that participants engaged more in making comparisons among their own data instances than in comparing their data with the baselines. However, we also observed that the participants made several comparisons involving only the baselines, such as, comparing measured values of Cohort data at different times (1661) and comparing Cohort and All data for the same time instances (719). We suspect




Figure 3: Number of identify and inspect operations across all participant interactions for each type of visualization presented in the interface. (The counts for monthly summaries are generally lower because they contained fewer data points.)

that the additional visual encoding of the number of participants in the Cohort and All data points could have contributed to these sequences of *inspect* operations. While people are generally more interested in their own data, additional visual encodings and attractive elements could also engage users in personal visualizations.

5.4 Influences of Interface Design. We found that some of the design choices made inadvertently influenced user behaviors on the interface. Weekly summaries of the activity feature was the default visualization which appeared on page load for all the participants. We found that participants interacted with this visualization more than any of the other five weekly-summary visualizations (having the same number of data points) as shown in Figure 3.

We also observed that many participants explored the features in a linear fashion, that is, in the order in which the options were listed. We investigated sequences of *explore* operations within participant interactions and found that 261 out of the 369 participants explored the six features in the order in which the choices were presented. Activity was the default feature appearing on page load following which these participants explored the remaining features, namely, sleep, HRV, hours spent at home, hours spent at office, and number of steps, sequentially.

We relate these findings to "structural" interactions [BVV*18] in visualizations where users make use of visually-apparent structures when exploring an unfamiliar interface. While structuring can be leveraged to guide people through their visualizations, it could also be important to give users more control over how they explore their data [HTA*14], for example, by using a dashboard layout. Future work should also look into different structuring techniques and how they affect user exploration with their personal data.

6. Conclusion

We illustrate how interaction logs can serve as a meaningful means to study personal visualizations. The exploratory behaviors identified in our study can be used to inform the design of personal visualizations. Our proposed method and metrics could be used in empirical evaluations of personal visualization and also supplement suitable qualitative methods [TLCC17].

References

- [BBK*15] BORKIN M. A., BYLINSKII Z., KIM N. W., BAINBRIDGE C. M., YEH C. S., BORKIN D., PFISTER H., OLIVA A.: Beyond memorability: Visualization recognition and recall. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 519–528. 3
- [BDF15] BOY J., DETIENNE F., FEKETE J.-D.: Storytelling in information visualizations: Does it engage users to explore data? In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), pp. 1449–1458. 2, 3
- [BH19] BATTLE L., HEER J.: Characterizing exploratory visual analysis: A literature review and evaluation of analytic provenance in tableau. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 145–159. 2
- [BLIC18] BREHMER M., LEE B., ISENBERG P., CHOE E. K.: Visualizing ranges over time on mobile phones: a task-based crowdsourced evaluation. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 619–629. 4
- [BOZ*14] BROWN E. T., OTTLEY A., ZHAO H., LIN Q., SOUVENIR R., ENDERT A., CHANG R.: Finding waldo: Learning about users from their interactions. *IEEE Transactions on visualization and computer graphics* 20, 12 (2014), 1663–1672. 2
- [BVV*18] BLASCHECK T., VERMEULEN L. M., VERMEULEN J., PERIN C., WILLETT W., ERTL T., CARPENDALE S.: Exploration strategies for discovery of interactivity in visualizations. *IEEE transactions on visualization and computer graphics* 25, 2 (2018), 1407–1420. 4
- [CL*15] CHOE E. K., LEE B., ET AL.: Characterizing visualization insights from quantified selfers’ personal data presentations. *IEEE computer graphics and applications* 35, 4 (2015), 28–37. 2, 3, 4
- [CLZ*17] CHOE E. K., LEE B., ZHU H., RICHEL N. H., BAUR D.: Understanding self-reflection: how people reflect on personal data through visual data exploration. In *Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare* (2017), pp. 173–182. 2, 3, 4
- [D3.20] D3.js - Data-Driven Documents. Accessed: 02-07-2020. URL: <https://d3js.org/>. 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, 3 (2009), 52–61. 1, 2
- [FPH18] FENG M., PECK E., HARRISON L.: Patterns and pace: Quantifying diverse exploration behavior with visualizations on the web. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 501–511. 1, 2
- [GGZL15] GUO H., GOMEZ S. R., ZIEMKIEWICZ C., LAIDLAW D. H.: 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, 1 (2015), 51–60. 2, 4
- [HTA*14] HUANG D., TORY M., ASENIERO B. A., BARTRAM L., BATEMAN S., CARPENDALE S., TANG A., WOODBURY R.: Personal visualization and personal visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 21, 3 (2014), 420–433. 1, 2, 3, 4
- [LDF11] LI I., DEY A. K., FORLIZZI J.: Understanding my data, myself: supporting self-reflection with ubicomp technologies. In *Proceedings of the 13th international conference on Ubiquitous computing* (2011), pp. 405–414. 2
- [MGA*19] MATTINGLY S. M., GREGG J. M., AUDIA P., BAYRAK-TAROGLU A. E., CAMPBELL A. T., CHAWLA N. V., DAS SWAIN V., DE CHOUDHURY M., D’MELLO S. K., DEY A. K., ET AL.: The Tesserae project: Large-scale, longitudinal, in situ, multimodal sensing of information workers. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), pp. 1–8. 2
- [MKK15] MAHYAR N., KIM S.-H., KWON B. C.: Towards a taxonomy for evaluating user engagement in information visualization. In *Workshop on Personal Visualization: Exploring Everyday Life* (2015), vol. 3, p. 2. 3
- [OYC15] OTTLEY A., YANG H., CHANG R.: Personality as a predictor of user strategy: How locus of control affects search strategies on tree visualizations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), pp. 3251–3254. 2
- [PTT*15] PERIN C., THUDT A., TORY M., WILLETT W., CARPENDALE S.: Personal Visualization: Exploring Data in Everyday Life, 2015. 2
- [RHF10] RODDEN K., HUTCHINSON H., FU X.: Measuring the user experience on a large scale: user-centered metrics for web applications. In *Proceedings of the SIGCHI conference on human factors in computing systems* (2010), pp. 2395–2398. 3
- [SH14] SADOWSKI S., HEIDMANN F.: A Visual Survey of Information Visualizations on Smartphones. 4
- [TLCC17] THUDT A., LEE B., CHOE E. K., CARPENDALE S.: Expanding research methods for a realistic understanding of personal visualization. *IEEE computer graphics and applications* 37, 2 (2017), 12–18. 2, 4
- [WBF17] WALL E., BLAHA L. M., FRANKLIN L., ENDERT A.: Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2017), IEEE, pp. 104–115. 2
- [YaKS07] YI J. S., AH KANG Y., STASKO 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. 3
- [YLZ14] YANG H., LI Y., ZHOU M. X.: Understand users’ comprehension and preferences for composing information visualizations. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 1 (2014), 1–30. 2