

Investigating the Role of Locus of Control in Moderating Complex Analytic Workflows

R. Jordan Crouser¹ , Alvitta Ottley² , Kendra Swanson¹, and Ananda Montoly¹

¹Smith College, USA

²Washington University in St. Louis, USA

Abstract

Throughout the last decade, researchers have shown that the effectiveness of a visualization tool depends on the experience, personality, and cognitive abilities of the user. This work has also demonstrated that these individual traits can have significant implications for tools that support reasoning and decision-making with data. However, most studies in this area to date have involved only short-duration tasks performed by lay users. This short paper presents a preliminary analysis of a series of exercises with 22 trained intelligence analysts that seeks to deepen our understanding of how individual differences modulate expert behavior in complex analysis tasks.

CCS Concepts

• **Human-centered computing** → Visualization; User models;

1. Introduction

The development of data visualization research in the past decades enable data visualization systems to achieve greater general usability and usage in various domains. Such advancements improved not only the understanding of the data, but also the understanding of people and how they use data visualization systems. In particular, the visualization community has begun to consider the potential benefits of shifting away from one-size-fits-all data visualization interfaces, acknowledging that individual differences may play a key role in the use of visualization tools [Yi12].

Personality and cognitive abilities have been shown to correlate with task performance [GF10, GF12, ZCY*11], search behavior and other usage patterns [BOZ*14, OYC15], and even user satisfaction [Kob04] with various visualization designs. In some circumstances, these effects have critical impact in important decision-making processes. For example, prior work by Ottley et al. investigating the impact of visualization on medical decision-making showed that people with high spatial ability tended to derive more benefit from visual aides than their low spatial ability counterparts [OPH*15]. These experiments showed that participants with low spatial ability had difficulty interpreting and analyzing the underlying medical data when they were presented with visual representations, and that approximately 50% of the studied population were inadequately supported by the visualization tools when making a life-critical decision.

It is interesting to note that individuals with traits or abilities largely considered to be positive may also face problems when the data visualization tools they use do not match their characteristics.

A study by Conati & Maclaren found that people with high perceptual speed were less accurate in computing derived values using radar graphs compared with colored boxes [CM08]. A later study by Ottley et al. found that people with a more internal *locus of control* (abbreviated LOC, a measure of the extent to which a person believes they have control over the outcome of events occurring around them [Rot66, Rot75, Rot90]) were slower and less accurate using an indented tree compared with a dendrogram [OYC15]. Both high perceptual speed and more internal locus of control correlate with high intellectual ability, and these results suggest that performance declines with incompatible visualizations. For a comprehensive survey of research into the role of individual differences in visualization, please see [LCO20].

It has been hypothesized that we can use stable features such as LOC to inform personalized interface designs to better support individual users [Yi12, ZOC*12]. Unfortunately, existing work in this area falls short of enabling these critical advances because of the limited scope and duration of studies performed to date. Most studies of the effects of personality on visualization observe each participant's behavior for **only a single, brief session on highly-simplified tasks**. Moreover, many have utilized platforms such as Amazon Mechanical Turk to achieve a sufficiently large sample size, which limits the control researchers have over participants' background, training, and expertise. Because of the constraints imposed by this experimental paradigm, we have observed only the effect these features have in the early staged of simplified analysis tasks performed by untrained lay individuals, but not how they influence behavior over the course of a trained analyst's larger investigative strategy.

To expand upon these previous studies, we conducted a series of multi-day exercises with trained intelligence analysts to investigate their behavior during complex analysis tasks. During this study, participants completed a battery of personality and cognitive style assessments, and were then asked to analyze a large synthetic dataset using an instrumented interactive search and visualization tool. In our investigation, we looked for patterns in data visiting behavior of the participants and attempted to relate these patterns to measures of individual difference. Through reporting this analysis, we make the following contributions:

- We confirm that individual differences are correlated with expert behavior in a complex analytical task.
- Specifically, we demonstrate that individuals with a more internal *locus of control* tended to exhibit higher interaction volume, as well as more complete coverage of the data.
- Finally, we provide recommendations for future investigation in this emerging area.

2. Case Study

Task design is critical to the success of an evaluation [Mun09], and researchers have created taxonomies for the types of tasks and interactions that are feasible for a given visualization (e.g., [AES05] and [YaKS07]). Our experiment focuses on exploratory data analysis, but we recognize that "exploration" as a task carries several different meanings [BH19]. In this paper, we want to distinguish between bottom-up exploration and top-down exploration. Bottom-up explorations "are driven in reaction to the data" [AZL*18] or "may be triggered by salient visual cues" [LH14]. This is open-ended and the user's instincts largely drives the interactions. Top-down explorations, on the other hand, are based on a high-level goals or hypothesis [BH19, LH14]. Much of the existing work on individual differences focus on the latter and have studied goal-driven search tasks. Grounded by existing literature, our experiment examines how individual traits mediate exploration patterns during an open-ended visual analytics task.

2.1. Participants

We recruited 32 Navy Reservists who chose this study from among several potential activities, all of whom had some training or experience in intelligence analysis. Of these, 22 were able to complete the assessment battery, as well as both days of the task. The remainder of this paper will refer only to those participants for whom we were able to collect complete data. Because this sample population differs substantially from the general population, we report a detailed demographic distribution for this subgroup below:

- **Age:** 4 participants were between the ages of 18-24, 10 between 25-34, 5 between 35-44, and 3 were over the age of 45.
- **Race/Ethnicity:** Our sample was predominantly white (16). 3 participants self-identified as multiracial, 2 as asian, and 1 as black. Across all racial groups, 3 participants self-identified as hispanic/Latinx.
- **Sex/Gender:** 13 participants self-identified as male, 2 as female, and 7 preferred not to have their sex/gender recorded.
- **Education:** Participants were highly educated: 7 held a graduate degree, 5 held a bachelor's degree, 7 held an associate's degree

or had at least some college, and the remaining 3 had high school diplomas. Education level correlated predictably with age.

- **Comfort with computers:** Participants were asked to rate their comfort using computers in both work and casual settings on a 4-point, forced-choice Likert scale. All but one participant reported feeling comfortable or extremely comfortable using computers.
- **Locus of Control:** Participants' LOC was scored using an online version of Nowicki and Strickland's 1973 questionnaire [NS73]. Using this instrument, participants scoring 8 or below (approximately 33% of the general population) are classified as having an *internal* LOC, while those scoring 17 or above (approximately 15% of the general population) are classified as having an *external* LOC. Those scoring 9-16 are categorized as *intermediate*. The LOC scores for our participants deviated substantially from the distribution over the general population: the median LOC score for our participants was 6, and 15 of the 22 participants scored 8 or below.

2.2. Task and Data

After completing a demographic survey and battery of cognitive style and personality factor assessments, participants received a short briefing outlining relevant background information regarding their specific task. The task used in this case study was adapted from the 2014 VAST Challenge [WCG*14], an annual contest with synthetic data and challenges designed to reflect real-world tasks in realistic conditions. Participants were presented with a hypothetical scenario describing a gas company on the fictitious island nation of Kronos. There has been a kidnapping involving some of the company's employees, and participants were asked to assist the investigators in the case. Their objective was to uncover the organizational structure of the group responsible for the kidnapping by analyzing several related synthetic data sources: *email headers* from 1170 internal company emails spanning two consecutive weeks, *resumes and short biographies* of 35 of the company's employees, *employee records* for 54 employees (with some overlap with the previous resumes and bios), *historical reports* and descriptions of the countries involved, and 458 current and historical *news reports* from multiple domestic and translated foreign sources.

2.3. Interface Design

Each participant was given a laptop and access to a web-based application through which they were able to explore the data. A series of collaborative design conversations with analysts in advance of these exercises surfaced a collection of common tasks performed during the early stages of analysis and information foraging:

- **Establishing Baseline:** During this phase, an analyst attempts to build her understanding of the typical or "normal" condition of the phenomenon which is being analyzed.
- **Entity Detection:** The process of identifying distinct entities that are regarded as objectively or subjectively significant.
- **Connection Detection:** The process of discovering if and how two or more entities are related.
- **Change Detection:** During this task, an analyst tries to identify inflection points which result in a change in some factor of interest that may result, or has already resulted, in a new baseline.

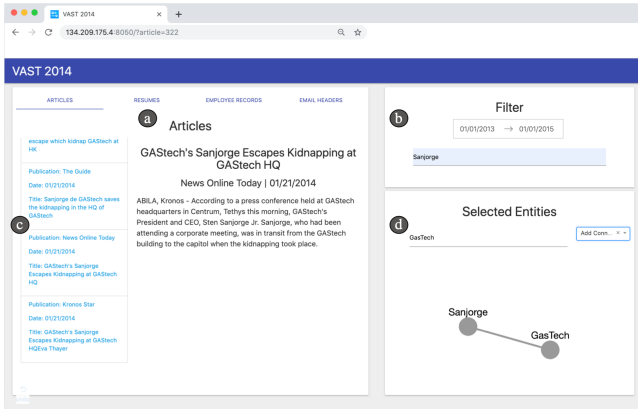


Figure 1: The information retrieval interface used in this study, containing four coordinated views: (a) details, (b) filter widget, (c) list of matching results, and (d) entity/connection scratchpad.

For the purposes of these exercises, we built a simple, baseline interface that supports these primary tasks (see Fig. 1) which consists of four coordinated components:

- a A central area in which detailed information about a single data element could be displayed, with tabs corresponding to the four information types: Articles, Resumes, Employee Records, and Email Headers;
- b A search widget which accepted both date ranges and generic text entry, and could also be used to specify more complex filter commands via a simplified query language;
- c A list view displaying the metadata for items matching the current filter parameters;
- d An interactive scratchpad where participants can record user-generated entities and the connections between them.

2.4. Data Collection

The study took place over the course of two half-day sessions held at North Carolina State University. Participants were colocated in a large conference room with up to 15 other individuals engaged in the same task, and communication between participants was not restricted. As participants analyzed the data, their interactions with both the interface and Google Docs (provided for note-taking) were logged and categorized as follows:

- *Search* actions were recorded whenever the participant clicked the Search button. The search parameters were also recorded.
- *GetDetail* actions were recorded whenever a participant clicked on an article, email header, resume, or employee data element. The data type and item ID were also recorded.
- *EditNotes* actions were recorded whenever a participant typed in their Google doc; a new *EditNotes* action was recorded each time the document auto-saved for as long as the participant was actively editing. These sequential actions were later condensed into a single action.
- *AddElement* and *AddConnection* actions were recorded whenever a participant interacted with the scratchpad.

3. Results

Figure 2 shows a temporal view of participants’ actions across the two half-day sessions. As expected, we observe that *Search* and *GetDetail* actions dominate the majority of analysts’ early interactions as they work to establish a baseline. These are punctuated by varying degrees of documentation, predominantly through *EditNotes* actions recorded by Google Docs, and more rarely through the addition of named entities and connections in the scratchpad.

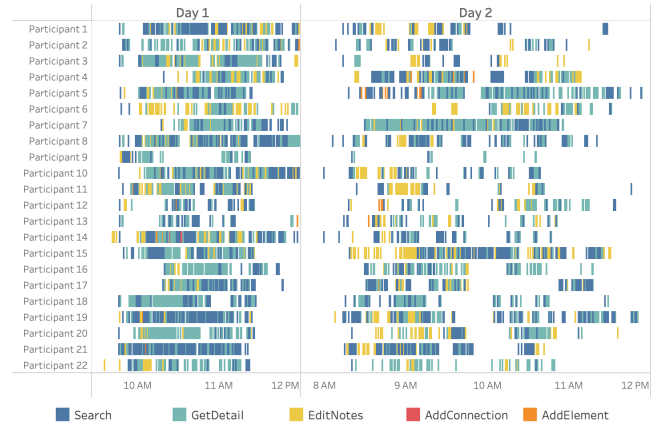


Figure 2: The temporal distribution of five distinct actions, by participant. We observe that *Search* and *GetDetail* actions dominate early in the analysis, and that interaction volume is not uniform.

3.1. Interaction Volume

We first observe that the overall volume of interaction is not uniform across all participants. We began our analysis by examining how the number of distinct actions performed varies with a participant’s LOC score. We observed that participants with a more internal LOC tended to perform more distinct actions with both the interface and in editing their notes than those whose LOC was more external. Figure 3 breaks the data down by the three most frequent actions that we observed: *EditNotes*, *GetDetail*, and *Search*. Though our sample size is too small to validate these observations statistically, our analysis revealed the same correlation between the number of interactions and participants’ LOC across all three primary action types.

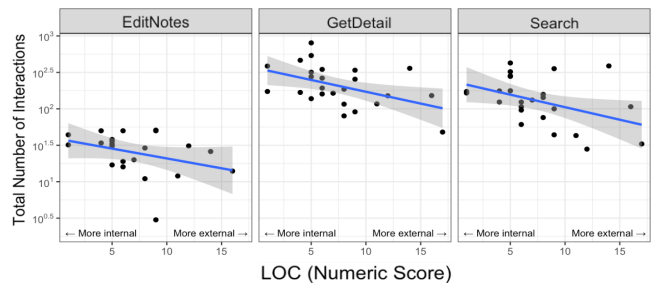


Figure 3: The distribution of total number of interactions across participants, ordered by LOC and broken down by interaction type.

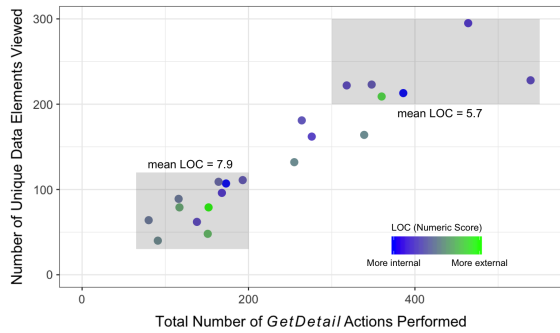


Figure 4: Each participants’ total volume of *GetDetail* interactions plotted against the number of unique data elements viewed. Participants are color-coded by LOC.

3.2. Data Coverage and Revisiting

The previous observation begs the question: are participants who perform more overall interactions engaging with a larger portion of the dataset during the course of their analysis, or are they interacting with the same data in a different way? The near-perfect linear trend in Figure 4 implies that across all participants and regardless of LOC, roughly half of all *GetDetail* actions involve a previously-unvisited data element. However, participants with a more internal LOC tend to visit more unique data elements, covering a larger area within the dataset in the same amount of time.

3.3. Anchoring

In addition to these overarching trends, we also observed a common pattern of data visiting behavior with respect to specific data elements: 9/22 participants performed the same sequence of *GetDetail* operations on articles 406, 121, 265, 227, and finally 56. These correspond to the first five articles that appear the top of the results list when the system is initially accessed, or when all filters are cleared. Similar patterns were observed for other data types, suggesting that participants use “clicking down” the list of results in order to get their bearings. While anchoring effects related to results listings have been documented elsewhere, a relationship between this behavior and LOC has not been previously observed (see Fig. 5). Subjects with a more internal LOC were much more likely to perform a click-down sequence of actions on the unfiltered data early in their investigation, whereas subjects with a more external LOC tended to perform this sequence later, and more often. This suggests that those with a more external LOC may be using the interface as a way to reorient their analysis after taking a break or hitting a dead end.

4. Discussion

The results of these exercises demonstrate that there is a relationship between LOC and expert behavior on complex analytical tasks. Specifically, we observed that LOC score was negatively correlated with interaction volume: the more internal a participant’s score, the more actions they performed and the more of the available data they were able to cover in the same amount of time. These

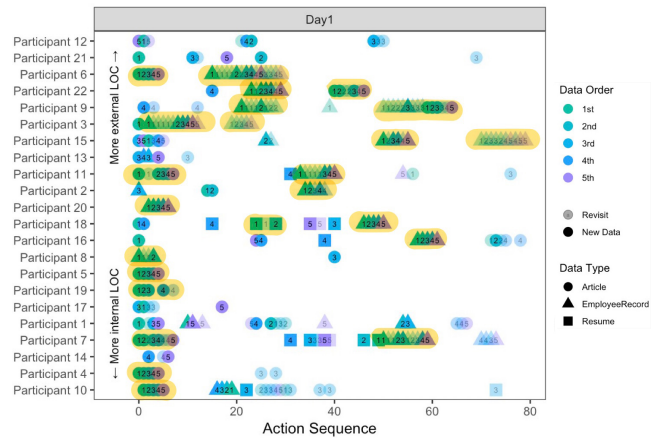


Figure 5: Participants’ interactions with the first 5 filtered results as they appear in panel (c) of the interface. “Click-down” actions (highlighted in yellow) are an emergent meta-action in which a participant progresses sequentially through these data items. Rows are ordered from most external (top) to most internal (bottom) LOC.

findings are in line with prior work in the visualization community showing that internal LOC generally maps to longer interaction times. However, to the best of our knowledge this is the first study to investigate and quantify interactions at this level of granularity. Analyzing low-level interactions, especially for complex analytical tasks gives us a window into *why* we tend to observe performance differences for individuals with different LOC scores.

In addition, we observed that those with a more external LOC were more likely to use features such as the interface’s ordering of the unfiltered data to reorient their analysis. This last phenomenon underscores the importance of understanding how individual differences influence analytical behavior: features like the ordering of unfiltered data items are not inherently meaningful, but in the absence of other cues some participants may be more inclined to treat it as a means to guide their analysis.

5. Conclusion

This short paper documents a preliminary case study to investigate whether relationships previously observed between individual differences such as *locus of control* and user behavior persist when studied in the context of more complex analytical tasks, and presents preliminary evidence affirming that they do. Moreover, the data collected through this case study presents a unique opportunity to observe analyst behavior on realistic tasks. To support continued inquiry in this emerging area, the anonymized dataset has been approved for unrestricted public release: github.com/SmithCollegeHCV/EuroVis2020-Data

Acknowledgements

The authors wish to thank Brent Younce (NCSU) for his technical expertise, as well as the 32 Naval Reservists who participated in this study for their time and service. This project was supported

in part by the Laboratory for Analytic Sciences at North Carolina State University.

References

- [AES05] AMAR R., EAGAN J., STASKO J.: Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.* (2005), IEEE, pp. 111–117. [2](#)
- [AZL*18] ALSPAUGH S., ZOKAEI N., LIU A., JIN C., HEARST M. A.: Futzing and moseying: Interviews with professional data analysts on exploration practices. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 22–31. [2](#)
- [BH19] BATTLE L., HEER J.: Characterizing Exploratory Visual Analysis: A Literature Review and Evaluation of Analytic Provenance in Tableau. In *Computer Graphics Forum* (2019), Wiley Online Library. [2](#)
- [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. [1](#)
- [CM08] CONATI C., MACLAREN H.: Exploring the role of individual differences in information visualization. In *Proceedings of the working conference on Advanced visual interfaces* (2008), ACM, pp. 199–206. [1](#)
- [GF10] GREEN T. M., FISHER B.: Towards the personal equation of interaction: The impact of personality factors on visual analytics interface interaction. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on* (2010), IEEE, pp. 203–210. [1](#)
- [GF12] GREEN T. M., FISHER B.: Impact of personality factors on interface interaction and the development of user profiles: Next steps in the personal equation of interaction. *Information Visualization* 11, 3 (2012), 205–221. [1](#)
- [Kob04] KOBSA A.: User experiments with tree visualization systems. In *IEEE Symposium on Information Visualization* (2004), IEEE, pp. 9–16. [1](#)
- [LCO20] LIU Z., CROUSER R. J., OTTLEY A.: Survey on individual differences in visualization. *ArXiv* (2020). [1](#)
- [LH14] LIU Z., HEER J.: The effects of interactive latency on exploratory visual analysis. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2122–2131. [2](#)
- [Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics* 15, 6 (2009), 921–928. [2](#)
- [NS73] NOWICKI S., STRICKLAND B. R.: A locus of control scale for children. *Journal of consulting and clinical psychology* 40, 1 (1973), 148. [2](#)
- [OPH*15] OTTLEY A., PECK E. M., HARRISON L. T., AFERGAN D., ZIEMKIEWICZ C., TAYLOR H. A., HAN P. K., CHANG R.: Improving bayesian reasoning: The effects of phrasing, visualization, and spatial ability. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 529–538. [1](#)
- [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), ACM, pp. 3251–3254. [1](#)
- [Rot66] ROTTER J. B.: Generalized expectancies for internal versus external control of reinforcement. *Psychological monographs: General and applied* 80, 1 (1966), 1. [1](#)
- [Rot75] ROTTER J. B.: Some problems and misconceptions related to the construct of internal versus external control of reinforcement. *Journal of consulting and clinical psychology* 43, 1 (1975), 56. [1](#)
- [Rot90] ROTTER J. B.: Internal versus external control of reinforcement: A case history of a variable. *American psychologist* 45, 4 (1990), 489. [1](#)
- [WCG*14] WHITING M., COOK K., GRINSTEIN G., LIGGETT K., COOPER M., FALLON J., MORIN M.: Vast challenge 2014: The kronos incident. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014), IEEE, pp. 295–300. [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. [2](#)
- [Yi12] YI J. S.: Implications of individual differences on evaluating information visualization techniques. *International Journal of Human-Computer Studies* 45, 6 (2012), 619–637. [1](#)
- [ZCY*11] ZIEMKIEWICZ C., CROUSER R. J., YAUILLA A. R., SU S. L., RIBARSKY W., CHANG R.: How locus of control influences compatibility with visualization style. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2011), IEEE, pp. 81–90. [1](#)
- [ZOC*12] ZIEMKIEWICZ C., OTTLEY A., CROUSER R. J., CHAUNCEY K., SU S. L., CHANG R.: Understanding visualization by understanding individual users. *IEEE computer graphics and applications* 32, 6 (2012), 88–94. [1](#)