Perception of Skill in Visual Problem Solving: An Analysis of Interactive Behaviors, Personality Traits, and the Dunning-Kruger Effect

Mengyu Chen* Emily Wall†

Emory University

ABSTRACT

Domain expertise plays a significant role in visual data analysis. In the best case, domain expertise allows people to effectively guide analyses based on prior experiences. However, a lack of self-evaluation of domain expertise can sometimes lead people to fail to recognize their actual performance when making decisions.This reflects a well known psychological phenomenon Dunning-Kruger(DK) effect, wherein unskilled people in certain areas will overestimate their competence while the skilled tend to underestimate their performance accordingly. This paper reports on a within-subject study to (1) replicate the DK effect in a visual problem-solving task, (2) understand participants' interaction behavior while performing the task, and (3) examine the relationship between personality traits and the possibility of being subjected to DK effect. Results show that participants who ranked in the top group and bottom group did misjudge their competence.

In addition, we observe that participants in the two extreme groups employed different strategies to perform the task. This difference can also be found in people with different personalities. Our findings can serve to (1) extend the existing knowledge regarding cognitive biases already detected in visual analytics, (2) inspire ideas to prevent or alleviate its impact on users by analyzing their strategies, and (3) support the design of visualizations considering the individual factors such as expertise and personality.

Index Terms: Human-centered computing—Visualization—Empirical studies in visualization

1 INTRODUCTION

Think back to primary school. We all knew someone who thought they were smarter and more capable than they really were. Let's call this person 'Bob'. Despite the fact that Bob normally scored below average in school assessments, his unshakable confidence of prior knowledge drove him to constantly raise his hands, to answer questions in the classroom. However, for the teachers and fellow classmates, his irrelevant answers uncovered that he had no idea what he was talking about. Obviously, Bob did not recognize his incompetence and was overconfident with his knowledge within the areas, which reflects a known cognitive bias: Dunning-Kruger effect [22].

In the seminal paper titled "Unskilled and Unaware of It," Kruger and Dunning describe a phenomenon in which the people who perform the worst on various knowledge tests have an inflated perception of their abilities [21]. Bottom quartile performers believed that their performance was above average, while those in the top quartile underestimated their performance relative to their peers. This lack of realization about one's own skill reflects a metacognitive deficit, i.e., a lack of "knowing what we know" and "knowing what we don't know" [1].

*e-mail: mengyu.chen@emory.edu

†e-mail:emily.wall@emory.edu

DK effect can have numerous consequences. In the opening example, Bob's tendency to express uninformed views likely prevents other students with more fruitful perspectives from participating. This phenomenon may also affect organizations, in which the most capable people may not be the ones making decisions; instead, those with the greatest self-perceived ability take precedence. Hence, the social consequences of this bias can lead to larger systemic problems. Namely, DK effect can lead to situations wherein true expertise may not reach the decision-making table, dominated instead by those who may be unaware of their own lack of proficiency. Further, when subject to the DK effect, people's overconfidence in their abilities may lead to challenges in visual data analysis. For instance, people may confidently report on flawed analyses or draw incorrect conclusions.

In this paper, we attempt to identify novel measures of DK effect through interaction analysis in a visual problem-solving task. This can have implications in visual data analysis, by illuminating new ways to characterize when the bias is present, which can inform future mitigation strategies. While cognitive biases have been studied extensively in the Cognitive Science community [29], they have received relatively less attention in data visualization and visual analytics. Recent efforts have examined e.g., the attraction effect [7], the anchoring effect [5,32], and priming [30]; however, to our knowledge, the DK effect has not yet been studied in the context of visual problem-solving or visual data analysis.

This study is designed to detect behavioral patterns correlated with Dunning-Kruger effect in a visual problem-solving task: the 15 puzzle. The 15 puzzle is a sliding puzzle game, which has 15 square tiles and one empty space in a 4×4 grid. The goal of the puzzle is to arrange the tiles in numerical order.

We investigate users' interactive behaviors when performing the task and examine if their personality traits are indicative of this bias.

Our contributions include:

1. We extend the existing knowledge regarding cognitive biases already detected in a visual problem-solving task to examine DK effect.
2. We provide an analysis of how interactive behaviors and personality traits relate to DK effect.

2 RELATED WORK

2.1 Cognitive Bias in Visual Analytics

In Cognitive Science, the term **bias** refers to errors that occur when people make decisions using "rules of thumb" or heuristics [17,18,29]. Use of heuristics, typically, is an efficient method by which information can be processed and decisions made [13]. However, occasionally these biases may lead to ineffective or wrong decisions. For the DK effect, in particular, this bias may lead to a range of problematic outcomes. Even among highly educated communities, e.g., physicians [6], pilots [24], reviewers and editors [16], people exhibit a compromised ability to accurately assess their own skills.

While the DK effect has not been studied in visual data analysis thus far, other recent work in data visualization and visual analytics informs our efforts. For instance, Wall et al. defined metrics to quantify [33] and refine [31] signals of bias from interactive behavior. Other metrics have been innovated to similarly capture concepts such as analytic focus [35] and exploration pacing and uniqueness [11]. Some such metrics have been associated with, e.g., selection bias [14] or anchoring bias [32]. Other researchers have replicated a variety of other cognitive biases in visual analytics. For instance, Xiong et al. showed how prior knowledge or beliefs about data influence the way people interpret new charts and communicate with visualizations (the curse of knowledge) [34]. Cho et al. also demonstrated the anchoring effect in a visual analytic tool by priming [5]. To our knowledge, this work provides a first examination of the DK effect in interactive visual analytics.

2.2 DK Effect

In Kruger and Dunning's seminal work, they attributed the occurrence of DK effect to the lack of meta-cognitive abilities; that is, the knowledge about people's knowledge [21]. Three counter explanations, however, were put forward to explain why there might be such an effect.

Krueger et al. [20] point out a statistical artifact, regression toward the mean (RTM), which could explain why there would be a discrepancy between the actual performance and perceived performance for people at extreme levels of capability. What's more, they argue that it is a combination of RTM and a "better-than-average" (BTA) heuristic that helps to explain the asymmetry (the effect is greater at the low end than at the high end). Burson et al. [3] claim that task difficulty explains the reverse discrepancy when the task is really difficult; that is, poor performers were actually considerably better calibrated than high performers.

As a counter, Ehrlinger et al. [10] addressed these criticisms by using real-world settings and financial and social incentives. They investigated whether the performers in the bottom quartiles overestimated their relative and absolute ability after they controlled for measurement errors in real-world situations (in-class exams, debate tournaments) to confirm the original findings of Kruger and Dunning. In addition, their results suggest that (1) neither monetary nor social incentives affect the overestimation of performers in the bottom quartiles, and (2) the participants are inaccurate due to mistaken beliefs about their own performance, rather than due to a misconception about the performance of others.

Although there is a some work focusing on discontents [3, 19, 20], a large body of active and varied research work on DK effect confirmed the existence of the effect. The effect has been uncovered in many settings involving medical residents training [25], gun owners by quizzing on their knowledge of firearms [10], tournament players by using field surveys to test players' beliefs about their relative performance in "Texas Hold'em" poker and chess tournaments [9], debate teams by estimating how well they thought they had done in debate tournament [10], and beginning aviators [26]. Additionally, recent work examines DK effect in the context of nuclear weapons, English grammar and logical reasoning, taking personality and cognitive characteristics into consideration [27]. Here, we expand on the work to determine the presence of DK effect in a visual problem-solving task.

3 Experimental Setup

The goal of our experiment is to (1) examine if DK effect exists in a visual problem-solving task, (2) investigate users' interactive behaviors when performing the task, and (3) examine if their personality traits are indicative of this bias. To realize the three goals, we designed a pre-registered[1] experiment involving the 15 puzzle task. Experimental details can be found below.

Interface. We used a 15 puzzle board for our experiment (shown in Figure 1). The primary view is a 4 tiles high and 4 tiles wide frame which contains 15 square tiles numbered 1–15, leaving one unoccupied tile position. Tiles in the same row or column of the open position can be moved by dragging and dropping them horizontally or vertically, respectively. To guarantee comparability, all participants start the game with the same initial configuration. The goal of the puzzle is to rearrange tiles in ascending numerical order (1, 2, ..., 15).

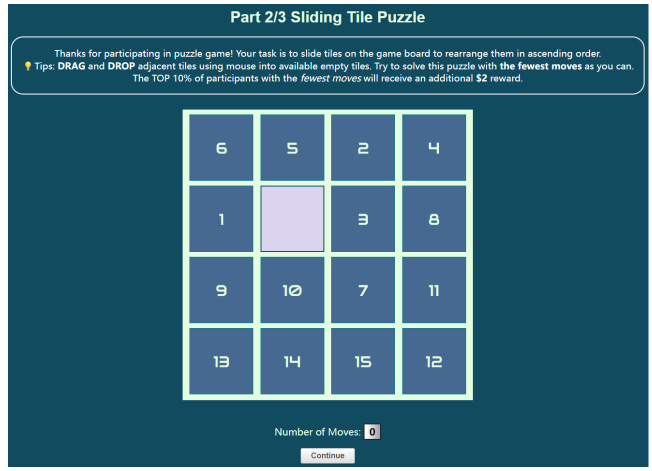

Figure 1: 15 Puzzle Board.

Participants. We recruited 48 participants through the Prolific crowdsourcing platform based on a power analysis of pilot data with 18 participants who completed the 15-puzzle using a similar experimental setup. Our minimum target sample size is 36 participants to obtain .8 power to detect a medium effect size of .25 at the standard .05 alpha. We attempted to recruit up to 48, assuming not all participants would complete the task. Participants were paid at a rate of $10 / hour for an estimated 15-minute task. Participants were also incentivized with an additional $1 performance bonus if they completed the puzzle with the fewest number of moves in the top 10%. Data of individuals who failed the attention check were discarded.

48 participants finished the whole task. We collected complete data of 39 individuals The data of 9 individuals were discarded due to (1) five people's logs were missing, and (2) four people refreshed the browser to solve the puzzle during the task, leading to erroneous move counts in subsequent logs.

This may lead to a skew in the data (e.g., poor performers with high move counts may be the ones more likely to refresh); however, we must use completed trials only for comparable analysis of this task. 23 out of the 39 individuals solved the puzzle in the optimal move count- 10 moves. The optimal solution was computed by A* algorithm based on our initial layout [15].

Procedure. Participants were first asked to complete an initial survey related to personality traits (including 20-item Big Five Personality [8] and 5-item Locus of Control [12]) and two attention check questions. Participants were then asked to solve a 15 puzzle game online using click and drag interactions. All participants started with a same initial layout. They were required to solve the puzzle game in the fewest movements possible. The 15-minute task estimate on the Prolific platform ensured a max time of 56 minutes on the task; however, users spent on average 7.5 minutes to complete the task. The interaction behaviors (including tiles in the board they clicked, positions they moved from/to and time stamps of each movement) while performing the task were recorded and later analyzed. Finally, after completing the puzzle, participants were asked to indicate their perceived task performance relative to their peers (recruited from Prolific crowdsourcing platform) with respect to (1) time spent, (2) number of moves and (3) reasoning

[1] https://osf.io/hqp6w

ability using percentile. Higher perceived percentile means that the participant perceived that they performed better than more of their peers.

Hypotheses. We hypothesize that:

H1. Less competent individuals, compared with their more competent peers, will dramatically overestimate their ability and performance.

H2. There will be detectable differences in strategies used by individuals who are more and less competent

H3. People with different personality traits will have different interactive movements/strategies.

H4. There will be correlations between personality traits (such as openness, conscientiousness, extraversion, agreeableness and neuroticism) and their perceived task performance.

4 Results

4.1 H1: D-K Effect in the Visual Task

4.1.1 Actual and Perceived Performance

To test **H1**, we assigned a percentile ranking for each participant based on their actual performance, as measured by actual **move count**. When performing the task, participants were informed that their performance was based on the number of moves to solve the puzzle. This analysis focuses on the low-skilled ($n = 9$) and high-skilled ($n = 9$) participants, whose actual performance fell in the bottom quartile and top quartile.

As Fig 2 illustrates, bottom quartile participants whose actual move counts ranked in the 18^{th} percentile on average, overestimated their performance to be around the 40^{th} percentile (blue dotted line, left) on average. In the top quartile, however, participants whose actual performance fell in the 90^{th} percentile grossly underestimated their move count compared to their peers to be in the 60^{th} percentile (blue dotted line, right) on average. When success is measured by time spent, similar trends can also be found. Complete figures are included in the supplemental materials.

4.1.2 Additional Analyses

We also consider that participants who performed the fewest movements might not spend the least time, as they are likely to spend more time strategizing before moving. Thus, we conducted additional analyses based on **time spent** to see if there are visible differences in results. We also examined perceived **reasoning ability**, gauged by a similar self-evaluation. The results are shown in Figure 2. The same trend can be observed here. Participants who ranked in the bottom quartile overestimated the efficiency of their time spent and reasoning ability compared to their peers to be around in the 40^{th} (purple dotted line, left) and 50^{th} (gray dotted line, left) percentiles, respectively. Participants in the top quartile also underestimated their time spent and ability to be in the 50^{th} (purple dotted line, right) and 62^{nd} (gray dotted line, right) percentiles, respectively. Overall, the trends for all three measures appear similar.

Collectively, we find support for **H1**.

4.2 H2: Interactive Strategies and Task Performance

To test **H2**, we visually examine the interactive behaviors of participants in four quartiles using lines overlaid on the puzzle grid, with thickness encoding number of moves to see in which quadrant each participant tended to move. The width of lines is normalized by each individual's own actual counts. Figure 3 depicts the movement path followed by participants in four groups. Among 39 participants, 23 of them got the optimal solution, thus, the movement path of the third ($n_{Q_3} = 10$) and top ($n_{Q_4} = 9$) quartile groups is the same. However, obvious differences in bottom and top quartile movement paths can be observed. Low-skilled people tended to randomly explore the board to find a solution, compared to their higher-skilled peers.

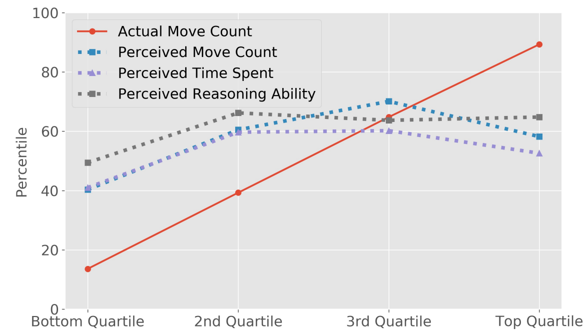

Figure 2: Perceived performance as a function of actual performance measured by movement counts. *X-axis* represents four quartiles divided by move counts; *Y-axis* represents percentile ranking.

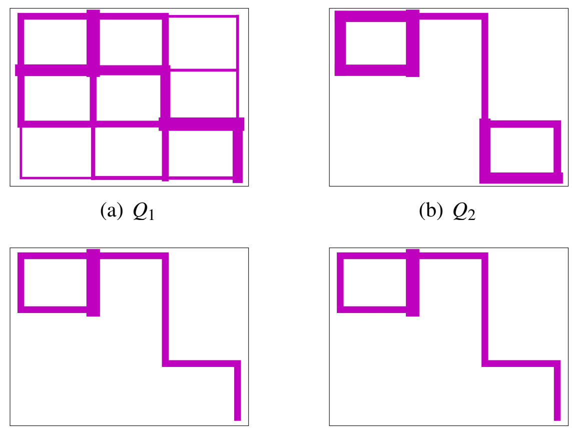

(a) Q_1 (b) Q_2

(c) Q_3 (d) Q_4

Figure 3: Movement path triggered by four group participants.

We further explore the differences in movement path among people who did not get the optimal solution. To have an apparent difference in the movement path, we categorized all participants who did not achieve the optimal solution into three groups, based on move counts equal to 12 ($n = 6$), between 12 - 100 ($n = 7$), and above 100 ($n = 3$). As depicted by Figure 4, people with greater move counts went over all grids in the board relatively evenly, while those with lower move counts appeared to interact less with the top right and bottom left. This could be due to strategies that involved placement of some fixed tiles in those areas of the board. Collectively, we find support for **H2**.

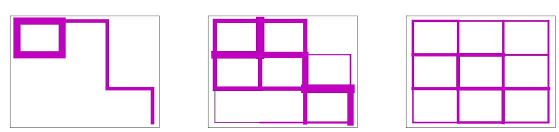

(a) 12 Counts (b) 12 - 100 Counts (c) Above 100 Counts

Figure 4: Movement path by people who did not get optimal solution.

4.3 H3: Interactive Strategies and Personality Traits

To test **H3**, we visualize participants' movement patterns similarly to the analysis for **H2**, but divided by low and high values for personality traits. Participants' personality traits scores are interpreted as "average" if scores are within one-half standard deviation of the mean. Scores outside that range can be interpreted as "low" or "high." Figure 5 illustrates movement paths by varied personality groups. For four personality traits (conscientiousness, extraversion,

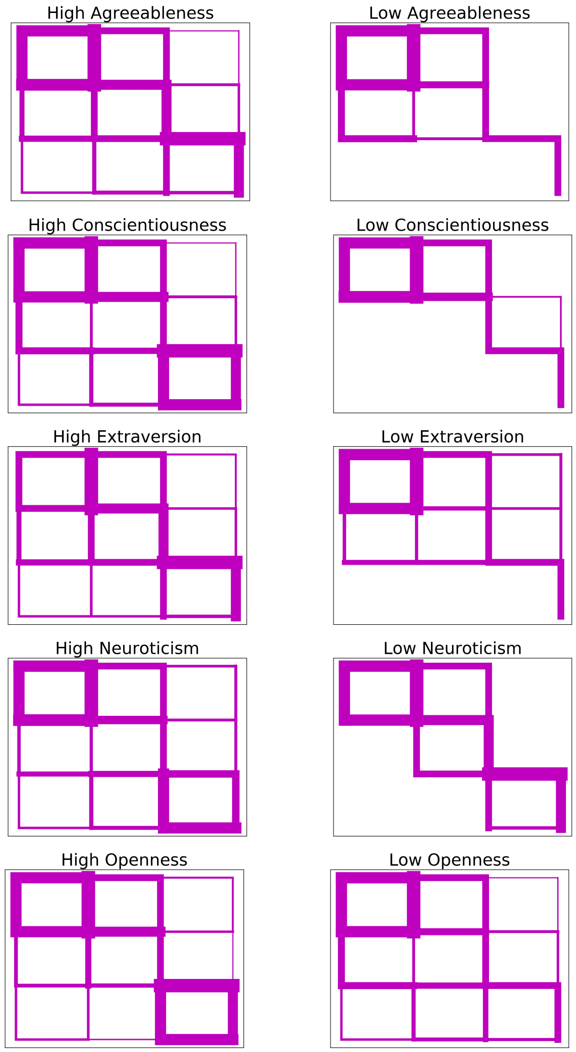

Figure 5: Movement path triggered by different personality traits.

agreeableness and neuroticism), participants with a higher score tended to explore each grid in order to find a solution. However, participants with a lower score left blank areas in certain grids. This could be due to strategies they employed, e.g., thinking about the reasoning before moving the tile. Regarding openness, there is less detectable difference, but participants with a higher score tended to move in the two diagonal fields. Thus we find mixed support for **H3**.

4.4 H4: Personality Traits and DK Effect

To test **H4**, we used a scatter plot to investigate the correlation between personality traits (i.e., scores for individual personality traits, x-axis) and their corresponding actual (Figure 6) or perceived (Figure 7) task performance (y-axis).

We further conducted a regression analysis using Least Squares method [23]. The results suggest that there are not detectable correlations between personality traits and performance. However, we detected a statistically non-significant positive correlation between conscientiousness traits and actual move count ($R^2 = 0.106$, $p = 0.129 > .05$), shown in Figure 6(b). There also exists a weak negative correlation between neuroticism traits and perceived move counts ($R^2 = 0.114$, $p = 0.086 > .05$), shown in Figure 7(d). Complete regression results are included in supplemental materials. Collectively, we find no support for **H4**.

5 Discussion

This paper provides an analysis of interactive behavior patterns and personality traits related to DK effect. Here we discuss the strengths of our analysis, limitations of our approach and future work.

Interactive Strategies Reveal Differences. By visualizing participants' movement paths, we observed differences in strategies employed by various groups, by **skill** and by **personality**. Low-skilled people tended to randomly explore the board to achieve a solution compared to their high-skilled peers. Furthermore, we observe different interactive strategies in individuals who possess low or high characteristics in some personality traits. This provides additional support for prior work indicating that personality traits can be correlated with different interactive behaviors [2].

This sheds light on behavioral patterns of different groups who are subjected to DK effect and thus provides ideas for subsequent studies. For instance, if we are able to identify low and high skilled individuals, or individuals with specific personality traits based on interactive behaviors, can we introduce personalized interventions that promote more effective decision-making processes?

Limitations and Future Work. Our work provides an analysis of how interactive behaviors and personality traits relate to DK effect, extending the existing knowledge regarding cognitive biases already detected in visual analytics; however, our approach has at least three primary limitations.

We were able to detect a curve that is similar to the original DK effect paper; however, due to the relative ease of our task, many participants achieved the optimal solution. This limits our ability to make deeper inferences about strategies. Future studies could elucidate additional differences in observed DK effect and behavioral patterns stratified by task difficulty. One of them could be examining the impact of initial configurations on performance and strategies, as different layouts result in different levels of task difficulty. Although we found some weak trends in certain personality traits and performance, this could be studied further in future experiments with larger sample sizes to detect stronger correlations. Furthermore, we can envision additional studies that assess different visual data analysis tasks (e.g., clustering, ranking, etc).

Second, our analysis of strategies is solely based on the holistic movement paths of participants. Future work could include additional analysis, e.g., sequence analysis to examine if there are earlier patterns that could predict strategy or lack thereof. Similarly, for those who strategized prior to interacting with the puzzle, measuring gaze movements with eye tracking could be beneficial to differentiate alternative strategies. Additional work is needed to assess the benefits and drawbacks of alternative visual representations and interaction designs that might capture behaviors associated with bias.

Lastly, another limitation relates to participant engagement in the task. Due to the COVID-19 pandemic, studies were conducted remotely and asynchronously without a study administrator present. If participants rapidly click through the task without careful thought about the task, this could also impact the veracity of our findings. Future work could include additional higher-powered studies to isolate potential confounding factors and determine whether the effect leads to discernible patterns in interactive behavior. Another beneficial avenue of future work could be to explore: are there ways in which visualizations may help *improve* people's metacognitive abilities and thus their ability to accurately assess their own skills and shortcomings? We believe that work on guidance [4] in visualization and reflective design [28] in HCI may be fruitful starting points for designing future visualization systems that may help users become more capable of accurate self-assessment.

6 Conclusion

In this paper, we designed a study to detect behavioral patterns correlated with Dunning-Kruger effect in a visual problem-solving task,

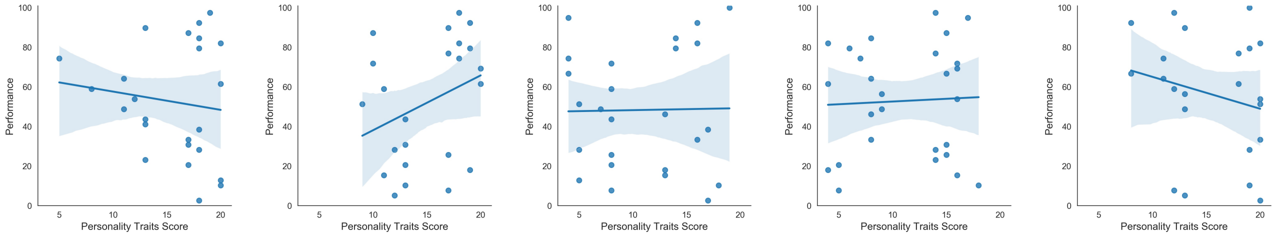

(a) Agreeableness (b) Conscientiousness (c) Extraversion (d) Neuroticism (e) Openness

Figure 6: Correlation between personality traits and actual move count.

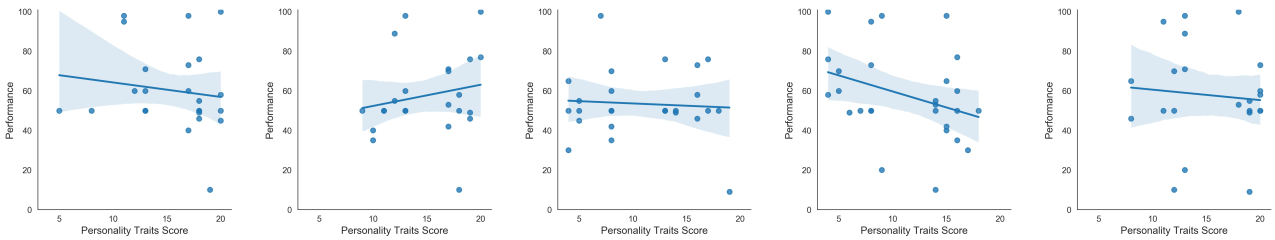

(a) Agreeableness (b) Conscientiousness (c) Extraversion (d) Neuroticism (e) Openness

Figure 7: Correlation between personality traits and perceived move count.

as well as to examine if personality traits are indicative of this bias. We detected the DK effect in our study; that is, incompetent people think they perform much better than they actually do while competent people underestimate their actual performance. In addition, our analysis suggests interactive strategies can look different for individuals of high and low skill and those with varying personality traits, although we found no statistically significant support to suggest that personality traits correlate to susceptibility to DK effect.

Our work builds on the knowledge of the presence of cognitive biases in visual data analysis and provides additional support for the power of inference based on interactive behaviors.

Acknowledgments

The authors wish to thank anonymous reviewers, Emory CAV Lab members, Emily Qiao, and Jie Mei for their help.

References

[1] E. Blakey and S. Spence. *Developing metacognition*. ERIC Clearinghouse on Information and Technology, 1990.

[2] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding waldo: Learning about users from their interactions. *IEEE Transactions on visualization and computer graphics*, 20(12):1663–1672, 2014.

[3] K. A. Burson, R. P. Larrick, and J. Klayman. Skilled or unskilled, but still unaware of it: how perceptions of difficulty drive miscalibration in relative comparisons. *Journal of personality and social psychology*, 90(1):60, 2006.

[4] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):111–120, 2016.

[5] I. Cho, R. Wesslen, A. Karduni, S. Santhanam, S. Shaikh, and W. Dou. The anchoring effect in decision-making with visual analytics. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 116–126. IEEE, 2017.

[6] D. A. Davis, P. E. Mazmanian, M. Fordis, R. Van Harrison, K. E. Thorpe, and L. Perrier. Accuracy of physician self-assessment compared with observed measures of competence: a systematic review. *Jama*, 296(9):1094–1102, 2006.

[7] E. Dimara, A. Bezerianos, and P. Dragicevic. The attraction effect in information visualization. *IEEE transactions on visualization and computer graphics*, 23(1):471–480, 2016.

[8] M. B. Donnellan, F. L. Oswald, B. M. Baird, and R. E. Lucas. The mini-ipip scales: tiny-yet-effective measures of the big five factors of personality. *Psychological assessment*, 18(2):192, 2006.

[9] D. Dunning. The dunning–kruger effect: On being ignorant of one's own ignorance. In *Advances in experimental social psychology*, vol. 44, pp. 247–296. Elsevier, 2011.

[10] J. Ehrlinger, K. Johnson, M. Banner, D. Dunning, and J. Kruger. Why the unskilled are unaware: Further explorations of (absent) self-insight among the incompetent. *Organizational behavior and human decision processes*, 105(1):98–121, 2008.

[11] M. Feng, E. Peck, and L. Harrison. Patterns and pace: Quantifying diverse exploration behavior with visualizations on the web. *IEEE transactions on visualization and computer graphics*, 25(1):501–511, 2018.

[12] L. R. Goldberg, J. A. Johnson, H. W. Eber, R. Hogan, M. C. Ashton, C. R. Cloninger, and H. G. Gough. The international personality item pool and the future of public-domain personality measures. *Journal of Research in personality*, 40(1):84–96, 2006.

[13] D. G. Goldstein and G. Gigerenzer. Models of ecological rationality: the recognition heuristic. *Psychological review*, 109(1):75, 2002.

[14] D. Gotz, S. Sun, and N. Cao. Adaptive contextualization: Combating bias during high-dimensional visualization and data selection. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*, pp. 85–95, 2016.

[15] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[16] S. Huang. When peers are not peers and don't know it: The dunning-kruger effect and self-fulfilling prophecy in peer-review. *Bioessays*, 35(5):414–416, 2013.

[17] D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

[18] D. Kahneman and S. Frederick. A model of heuristic judgment. the cambridge handbook of thinking and reasoning, eds holyoak kj, morrison rg, 2005.

[19] M. Krajc and A. Ortmann. Are the unskilled really that unaware? an alternative explanation. *Journal of Economic Psychology*, 29(5):724–738, 2008.

[20] J. Krueger and R. A. Mueller. Unskilled, unaware, or both? the better-than-average heuristic and statistical regression predict errors in estimates of own performance. *Journal of personality and social psychology*, 82(2):180, 2002.

[21] J. Kruger and D. Dunning. Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-

assessments. *Journal of personality and social psychology*, 77(6):1121, 1999.

[22] T. D. Lab. Why can we not perceive our own abilities?, January 2021.

[23] M. Merriman. *A List of Writings Relating to the Method of Least Squares: With Historical and Critical Notes*, vol. 4. Academy, 1877.

[24] S. R. Pavel, M. F. Robertson, and B. T. Harrison. The dunning-kruger effect and siuc university's aviation students. *Journal of Aviation Technology and Engineering*, 2(1):6, 2012.

[25] M. Rahmani. Medical trainees and the dunning–kruger effect: when they don't know what they don't know. *Journal of Graduate Medical Education*, 12(5):532–534, 2020.

[26] C. Sanchez and D. Dunning. Overconfidence among beginners: Is a little learning a dangerous thing? *Journal of personality and Social Psychology*, 114(1):10, 2018.

[27] D. J. Sanchez and A. Speed. The impact of individual traits on domain task performance: Exploring the dunning-kruger effect. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2020.

[28] P. Sengers, K. Boehner, S. David, and J. Kaye. Reflective design. In *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*, pp. 49–58, 2005.

[29] A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131, 1974.

[30] A. C. Valdez, M. Ziefle, and M. Sedlmair. Priming and anchoring effects in visualization. *IEEE transactions on visualization and computer graphics*, 24(1):584–594, 2017.

[31] E. Wall, A. Arcalgud, K. Gupta, and A. Jo. A markov model of users' interactive behavior in scatterplots. In *2019 IEEE Visualization Conference (VIS)*, pp. 81–85. IEEE, 2019.

[32] E. Wall, L. Blaha, C. Paul, and A. Endert. A formative study of interactive bias metrics in visual analytics using anchoring bias. In *IFIP Conference on Human-Computer Interaction*, pp. 555–575. Springer, 2019.

[33] E. Wall, L. M. Blaha, L. Franklin, and A. Endert. 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)*, pp. 104–115. IEEE, 2017.

[34] C. Xiong, L. Van Weelden, and S. Franconeri. The curse of knowledge in visual data communication. *IEEE transactions on visualization and computer graphics*, 26(10):3051–3062, 2019.

[35] Z. Zhou, X. Wen, Y. Wang, and D. Gotz. Modeling and leveraging analytic focus during exploratory visual analysis. *arXiv preprint arXiv:2101.08856*, 2021.