Toward a Design Space for Mitigating Cognitive Bias in Vis

Emily Wall* Georgia Tech

John Stasko† Georgia Tech

Alex Endert‡ Georgia Tech

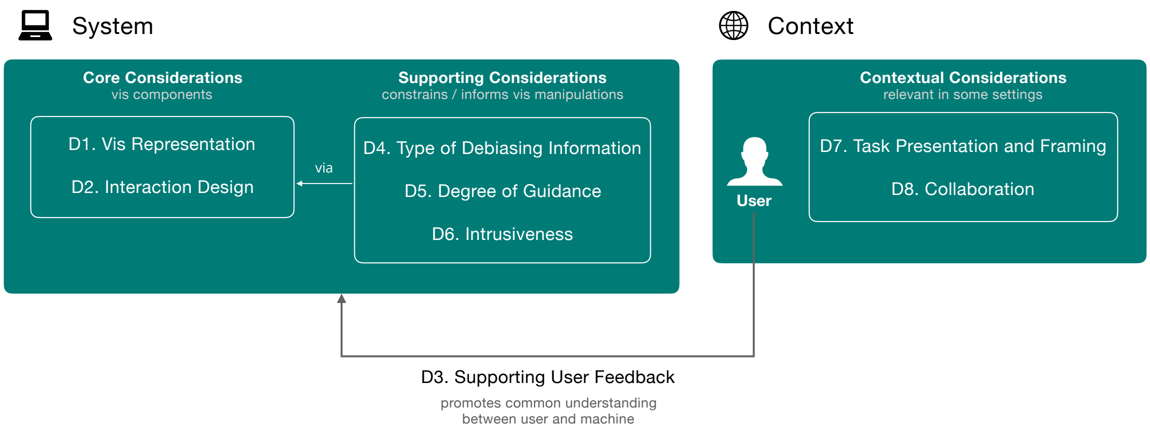

Figure 1: The design space is comprised of 8 dimensions, described in Section 3. D1 (VISUAL REPRESENTATION) and D2 (INTERACTION DESIGN) are the two *core* components of a visualization [13] that can be manipulated to mitigate biased decision making processes. How these components are manipulated is informed and constrained by *supporting* considerations, including D4 (TYPE OF DEBIASING INFORMATION), D5 (DEGREE OF GUIDANCE) and D6 (INTRUSIVENESS). Some *contextual* considerations may only be relevant in specific settings, including D7 (TASK PRESENTATION AND FRAMING) and D8 (COLLABORATION). Finally, D3 (SUPPORTING USER FEEDBACK) connects the user and contextual setting to the system by promoting a common understanding between user and machine.

Abstract

The use of cognitive heuristics often leads to fast and effective decisions. However, they can also systematically and predictably lead to errors known as *cognitive biases*. Strategies for minimizing or mitigating these biases, however, remain largely non-technological (e.g., training courses). The growing use of visual analytic (VA) tools for analysis and decision making enables a new class of bias mitigation strategies. In this work, we explore the ways in which the design of visualizations (vis) may be used to mitigate cognitive biases. We derive a design space comprised of 8 dimensions that can be manipulated to impact a user's cognitive and analytic processes and describe them through an example hiring scenario. This design space can be used to guide and inform future vis systems that may integrate cognitive processes more closely.

Index Terms: Human-centered computing—Human Computer Interaction (HCI); Human-centered computing—Visualization

1 Introduction

Visual perception has been an ongoing focus in visualization research (e.g., [8, 18, 37]); however, cognition has received relatively less attention (e.g., [15]). Nonetheless, cognition forms a vital component of visual data analysis, encompassing sensemaking, decision making, and so on. We argue that increasing attention to cognition has the potential for profound impact. As interactive visualizations are increasingly used for data analysis and decision making in widespread domains, these processes can be improved by designing systems that can both leverage analysts' cognitive strengths and guard against cognitive limitations and weaknesses. One potential weakness is *cognitive bias*, errors arising from the use of heuristics in decision making [26]. In this work, we focus on **deriving a design space for visualization systems that can mitigate bias**.

Prior work detailing bias mitigation, or debiasing techniques, has largely relied on non-technological strategies, like training courses [17, 20]. However, as data analysis increasingly takes place through technological media, particularly using visualization, we are motivated to consider ways in which vis design can improve decision making processes. While some prior work has provided guidelines toward mitigating one type of bias in a particular context [9, 14], we take a more general approach aimed at increasing real-time awareness of bias abstracted from a specific scenario. Given the recent emergence of bias mitigation in vis, our design space is derived from (1) prior work describing bias mitigation strategies outside of the vis community, as well as (2) potential areas of vis research that may inform the design of systems that mitigate bias.

Toward this goal, we must make a key assumption: that systems have information about bias in the user's decision or analytic process. Prior work has developed techniques that make this assumption reasonable. For example, computational methods exist for quantifying bias in the analytic process [14, 41, 42]. Given this information, or other forms of de-biasing information, the goal is then to design systems that can help people make better decisions by compensating for the ways in which people are likely to make cognitive errors.

Many different types of biases can have common observed behavioral effects [27, 41]. For example, an analyst subject to vividness criterion [20] (over-reliance on information that is vivid or personal) may interact with a particularly vivid data point repeatedly. The same behavior is likely to be observed if the analyst is subject to a different type of cognitive bias, like the continued influence ef-

*e-mail: emilywall@gatech.edu
†e-mail: stasko@gatech.edu
‡e-mail: endert@gatech.edu

fect [20] (continued reliance on a piece of evidence, even after it has been discredited). As a result, some mitigation strategies can, to varying degrees, have an effect on multiple types of biases [1]. Hence, in this design space, we do not focus on any specific type of cognitive bias. Rather, we find it prudent to introduce design considerations for mitigating bias and improving analytic *processes*, agnostic to a specific type of bias.

Within this context, the contribution of this work is the derivation of 8 dimensions of vis design that designers should consider when developing systems to mitigate biased decision making, or retrofitting such capabilities in existing tools. These dimensions represent aspects of a vis system that can be manipulated in specific contexts to mitigate biased decision making. We concretize these dimensions through examples. In supplemental materials, we further include the design of a hypothetical VA system, `fetch.data`, to illustrate potential bias mitigation interventions. While prior work on bias mitigation in the context of vis and visual analytics is limited, we find it timely to scaffold design efforts going forward when building systems that can mitigate biased decision making.

2 Related Work

We organize related work as (1) existing strategies for mitigating cognitive bias (primarily outside of the VIS community), and (2) driving areas of visualization research that may inform bias mitigation in the context of the VIS community.

2.1 Existing Strategies for Mitigating Bias

Prior work categorized bias mitigation strategies in the context of intelligence analysis as either *training interventions* or *procedural interventions* [27]. We utilize the same terminology to refer more broadly to bias mitigation strategies in the context of data analysis and decision making, described next.

Training Interventions. Training interventions often come in the form of education (e.g., training courses, videos, and reading material [17, 20]) that may examine past errors to inform future decision making. Serious games (video games designed with educational purposes) have proven a more effective alternative to traditional means of bias training (e.g., [36]). These techniques educated analysts about cognitive biases, but nonetheless did little to mitigate negative effects when biases inevitably occurred in the analytic process. They reinforce that an analyst must be pro-active using feedback to adjust their behaviors to mitigate the negative effects of bias.

Procedural Interventions. Procedural interventions are integrated in the analytic process. While some work has theorized futuristic technologies that automatically correct user-injected bias in AI systems by modeling user trust [35], current practices are largely realized through non-technological means, includingstructured analytic techniques [21]. Perhaps the most known and accepted is Analysis of Competing Hypotheses (ACH) [20], a tactic that can be used during the analytic process to evaluate the likelihood of multiple hypotheses in an unbiased way. While an effective analytic tool, ACH is a time-consuming process not always used in practice. Other decision making strategies (e.g., "consider the opposite") have shown promise to reduce biases like overconfidence, hindsight, and anchoring [1]. However, these procedural strategies come at the cost of potential cognitive overload, which can ultimately amplify some biases [34]. Herein, we focus on machine-assisted strategies, specifically utilizing visualization, that can lighten the cognitive burden of bias mitigation in real-time.

2.2 Driving Areas in Visualization Research

While prior work on mitigating cognitive bias in the visualization domain is sparse [9, 14], we are motivated to define a design space in this emergent area. Hence, in the context of visualization research, we derive inspiration from sub-fields of visualization research that may be leveraged to mitigate biased decision making processes.

Guidance. According to Ceneda et al., guidance can be defined as "a computer-assisted process that aims to actively resolve a knowledge gap encountered by users during an interactive VA session" [6]. In other words, systems that guide users provide some form of assistance during interactive data analysis (e.g., Scented Widgets [44] in collaborative settings or VizAssist [4] for visualization creation).

Bias mitigation in VA can be loosely thought of as a form of guidance, where the goal is to impact the user's decision making in such a way as to promote a more balanced analytic *process* and / or a more reasonable *product* or final choice decision. Within Ceneda et al.'s [6] characterization of guidance,we focus on the *output means* in the context of bias mitigation.What can we show the user to facilitate an analytic process that is less prone to the potentially negative effects of cognitive bias?

Analytic Provenance. Analytic provenance is a description of the analytic process leading to a decision [31]. Many researchers have shown the impact of raising users' awareness of their process. Researchers have shown ways to measure or visualize the user's *coverage* of the data throughout analysis [2, 25], leading users to make more discoveries [44] and analyze the data more broadly [11, 28]. This body of research shows promise that provenance awareness can alter user behavior in the context of bias mitigation.

Mixed-Initiative VA. Mixed-initiative [23] VA tools explore the balance between human and machine effort and responsibilities. Some systems leverage users' interaction sequences to infer about their goals and intentions in an analytic model (e.g., [5, 43]). These types of mixed-initiative tools inspire potential ways of mitigating cognitive bias as people use visualizations. In particular, the machine could operate as an unbiased collaborator that can act on behalf of the user, or *take initiative*, to mitigate biased analysis processes.

3 Design Space

In this section, we describe 8 dimensions (D1-D8) important to the design of bias mitigation strategies in VA tools. These 8 dimensions are not strictly orthogonal, nor are they exhaustive. Rather, they represent our view of the aspects of visualization systems that may be manipulated for the purposes of mitigating bias given current technologies. Due to limited prior work on bias mitigation in VA, the process for deriving this design space was largely ad-hoc, guided primarily by literature review in related areas of vis research (Section 2). Many of the dimensions are related (Figure 1).

To ground our design space, we describe applied examples using a common scenario. Suppose a hiring manager at a tech company uses a VA tool to analyze tabular data about job applicants. From potentially hundreds of applications on file, the hiring manager wants to select a handful of candidates to interview. Suppose the system is comprised of three interactive views: (A) a scatterplot view, (B) a filter panel, and (C) a ranking table view. Scatterplot axes can be configured, the table sorted, and filters used to adjust the subset of data viewed. For each design dimension below, we describe the concept and revisit this example to illustrate how a visualization could be retrofitted to mitigate biased decision making. [1]

D1: Visual Representation- *Concept.* There are many possibilities for representing information that may have a debiasing effect. For bias interventions intended not to impose significant disruption to the user's natural analytic process, designers may opt for *peripheral* or *ex-situ* visualizations. Peripheral visualizations would appear in a separate view of the interface, potentially available on demand, and hence may be less likely to call the user's attention away from the primaryvisualization. On the other hand, *in-situ* visualizations

[1] We have designed a hypothetical VA system, `fetch.data` to demonstrate some of these concepts, which can be seen in supplemental material.

would appear within existing views. For example, in-situ visualizations could encode debiasing information in previously unused visual channels (e.g., opacity, color, position, etc). The choice between in-situ v. peripheral display of debiasing information should be informed by (1) type of debiasing information, and (2) intended level of user attention to that information. Furthermore, the representation of information should follow conventions described in vis research. For example, chart types [33] and optimal visual encoding [8] should be considered based on the type of data presented.

Example. Supposethe hiring manager is subject to anchoring bias, or the tendency to rely too heavily on initial "anchoring" information [10] (the first few résumés received). If the first handful of candidates happened to be males, successful bias mitigation strategies could draw the user's attention away from potential gender bias. Some metrics of bias (e.g., [41,42]) compare the distribution of user interactions to the underlying distributions of the data. This could be shown in a *peripheral* view showing both distributions. Alternatively, a single metric quantifying the severity of the bias could be encoded as an ambient background display where color or opacity represents level of bias. In another example, history (provenance) could be shown *in-situ* by encoding size of scatterplot points as time spent examining each candidate, drawing attention to those (female candidates) who may have been unintentionally ignored.

D2: Interaction Design- *Concept.* Altering the interaction design may be another impactful way to mitigate bias. For example, a designer's choice between a rectangle or lasso selection may have implications about bias. Similarly, a system could disable interaction with data / views when biased behavior is detected. However, altering interaction design and affordances to mitigate bias can often come at the expense of perceived user control and system usability. Designers of bias mitigation interventions should weigh the tradeoffs of these choices so usability is not unduly compromised.

Example. Consider a filtering widget designed to mitigate bias. If the hiring manager applies a filter to exclude female candidates in the data, a typical system response would be to remove female candidates from the views in the visualization. The system could instead respond by presenting a split or duplicated scatterplot view: one in which the manager's intended data is shown (male candidates), and one in which the filtered data is shown (female candidates).

D3: Supporting User Feedback- *Concept.* While the primary objective of bias mitigation interventions is to communicate information from the system to the user, supporting user feedback is likewise important. In real-world systems that may be able to characterize user bias with limited accuracy, it can enable the user to communicate information unknown to the underlying model of bias (e.g., that a presumed bias is not due to unconscious error, but rather an external task constraint). When user feedback is supported, users may be given an increased sense of mutual understanding or common ground with the system. Further, models of user bias might be improved as a result.

Example. In the hypothetical hiring scenario, suppose the system detects a strong (gender) bias in that the hiring manager has primarily interacted with male candidates. One system response could be to recommend female candidates. However, the hiring manager's focus could be the result of a constraint on the task unknown to the system (e.g., a division of labor between two managers). If the manager dismisses the recommendation of female candidates, the system can elicit feedback (e.g., via a pop-up dialog) to clarify information potentially outside the system's purview. Reasons may include things like a repetitive recommendation, an irrelevant recommendation, or an external task constraint. According to the hiring manager's selection, the system may alter the underlying model of bias to account for these preferences or constraints.

D4: Type of Debiasing Information- *Concept.* A primary consideration in designing bias mitigation strategies is the type of debiasing information that the system will capture and communicate to the user. Types of debiasing information that could promote user awareness includes things like analytic provenance, summative metrics that quantify the analytic process [12, 24], and so on. We could further conceive of future systems that are able to identify specific types of bias the user may be subject to by name (e.g., confirmation bias [30], anchoring bias [10], etc). Systems should ideally communicate information about potential biases in a way that guides users to counteract them (i.e., they should be *informative* and *actionable*).

Example. Suppose the hiring manager is exhibiting signs of availability bias [39], or a heavy reliance on information that is most easily remembered or most recent (i.e., the most recent application received). When bias is detected (i.e., the hiring manager is exhibiting signs of availability bias), the system could show provenance information to the hiring manager by adding an additional view to the interface that shows a snapshot of various stages of history of the manager's analytic process (e.g., like the history shown in [19]). Alternatively, the system could show the results of summative interaction metrics, similar to the metric visualization in [41]. This could enable the hiring manager to reflect on their process and adjust.

D5: Degree of Guidance- *Concept.* Degree of guidance is analogous to Ceneda et al.'s *guidance degree* in VA guidance [6]. It can be thought of as a spectrum that refers to how much the system "helps" the user. On one end of the spectrum, the system provides little intervention, while on the other end, the system more aggressively steers the user. Ceneda et al. describe three scenarios for degrees of guidance: *orienting*, *directing*, and *prescribing*, examples of which are described below. The degree of guidance adopted must be considered alongside tradeoffs of user experience. Systems that deny user control may come at the expense of perceived usability issues.

Example. An *orienting* bias mitigation strategy would promote user awareness of their biases. For example, the system could size candidates in the scatterplot according to the hiring manager's focus (where larger points represent neglected candidates). A *directing* bias mitigation strategy could suggest candidates to the hiring manager to consider from the pool of candidates who have not been analyzed. A *prescribing* bias mitigation strategy would involve the system assuming initiative or otherwise taking control from the user. An example of this might be disabling filters or interactions with specific candidates.

D6: Intrusiveness- *Concept.* Intrusiveness refers to how much the system interrupts or otherwise intrudes on the user's analysis process. On the low end of the spectrum, bias information may be presented peripherally or even on demand (i.e., user attention optional). Highly intrusive mitigation strategies may present information front and center requiring the user's attention until the perceived bias is addressed. The level of intrusiveness of the intervention should not outweigh the intended benefit, however. In lower-cost decisions (e.g., analyzing a dataset of food to construct a weekly menu), a highly intrusive bias mitigation strategy would likely be unwelcome to the user. On the other hand, the intrusion may be acceptable for decisions that carry greater importance (e.g., criminal intelligence analysis). This is distinct from D5 (DEGREE OF GUIDANCE). Consider the following analogy: suppose a person asks her friend for directions from point A to point B. The friend may draw a map, suggest GPS, or walk her friend there herself (i.e., DEGREE OF GUIDANCE). If she walks with her friend, she may exhibit a spectrum of INTRUSIVENESS (e.g., how closely does she stand to her friend).

Example. In our hypothetical scenario, a minimally intrusive mitigation strategy may present bias information to the hiring manager only *on-demand*. For example, there may be a tab in the interface that reveals information about the model of user bias when clicked on. A more intrusive bias mitigation strategy could be a pop-up notification repeatedly alerts the hiring manager until a less biased analysis state is reached.

D7: Task Presentation and Framing- *Concept.* Changes to the presentation of information can have an impact on the analytic process and outcome. Framing has been found to strongly shape decision-making [38], including richness of language used and positive v. negative terminology to describe logically equivalent information [40]. In one study, researchers showed that people chose one treatment (surgery) over another (radiation therapy) when it was described as having a 90% short-term survival rate v. a 10% immediate mortality rate [29]. In addition to language, visual framing or anchoring can also shape decision making [7]. In situations where designers of bias mitigation interventions have control of the task, thoughtful consideration should be given to the often subtle-seeming aspects of task presentation.

Example. This contextual consideration is primarily limited to situations in which the designer has control over the presentation of the task (e.g., in a user study). In our hypothetical scenario, the job description can impact the analysis process. For example, the framing of criminal background criteria may alter the hiring manager's threshold for minimally viable candidates (e.g., the negative framing "does not have a criminal record" may lead to a lower decision threshold than the positive framing "has a clean record"). Visual framing of information can also impact decision making (i.e., the relative size and spatial arrangement of multiple views, the order in which the hiring manager is trained to use them, etc).

D8: Collaboration- *Concept.* Collaborative contexts have potential to mitigate bias by allowing others to check an analyst's work. By leveraging "wisdom of crowds", collaboration mitigates that no sub-optimal individual decision prevails [22, 32]. Analysts teaming on a project may be alerted to biased behaviors, to ensure they cross-validate each other's work. In this case, prior work on fostering awareness in collaborative settings can be informative [2, 3, 16]. Collaboration is contextually relevant, as it may be infeasible in many scenarios due to the nature of the decision (e.g., a personal healthcare decision) or other constraints (e.g., division of labor).

Example. To leverage collaboration to mitigate biased decision making, designers of the vis tool could show traces of other hiring managers' exploration behaviors. For example, this could entail coloring points in a scatterplot based on which have been previously examined by other hiring managers (e.g., [2]).

4 Characterizing Existing Systems

Two recent works have designed interventions within visualization systems to mitigate cognitive bias [9, 14]. For each, we describe the context of the problem, the bias intervention, and how it fits within the aforementioned design space.

Mitigating Selection Bias. In analyzing high dimensional data sets, many dimensions may exhibit correlations. Hence, when attempting to select a sample from a larger dataset, the analyst may unintentionally filter out a representative part of a population (i.e., selection bias) [14]. To mitigate selection bias, Gotz et al. modified a visualization tool, DecisionFlow. Specifically, they modified an existing view in the visualization (D1, in-situ) by adding a color-coded bar after each subsequent data selection to depict the similarity of the subset to the original dataset. The color-coding of the bar was based on a computed value (D4, bias metric) that quantified the differences in variable distributions between the two datasets. They also added a secondary view (D1, ex-situ) that provided details about how variables of the data were constrained either via direct or unintentional filtering via correlation. These modifications represent an *orienting* degree of guidance (D5) that is relatively unintrusive (D6). They did not modify the interaction design (D2), task presentation (D7), or collaborative nature (D8) of the system, and did not enable user feedback (D3).

Mitigating the Attraction Effect. Dimara et al. designed an experiment to test two different strategies for mitigating the attraction effect (the phenomenon where a person's decision between two alternatives is altered by the introduction of an irrelevant third option) in scatterplots [9]. In one strategy, they highlighted optimal choices with a brightly colored stroke (D1, in-situ) before users clicked to select their choice point. This constitutes an *orienting* degree of guidance (D5). In another design, they altered the task framing (D7) and interaction design (D2) from "select a point" to "eliminate points until only one remains". While this was more effective than the first strategy, it could have usability implications as it represents a more intrusive (D6) design. For both strategies, they do not support user feedback (D3) or collaboration (D8). By virtue of these mitigation strategies taking place within an experiment, the debiasing information (D4) was a precondition to the study.

5 Discussion

This design space does not exhaustively include all possible contextual design considerations when building visualization systems that can mitigate biased decision making. For example, the device type may drive design choices that are compatible with varying input modalities or screen sizes (e.g., haptic feedback or other non-visual channels when screen real-estate is limited). Mitigation strategies may also be adaptive to the type of user of the system (casual user, domain expert, data analyst, etc). Furthermore, while we have focused on improving decision making *processes*, agnostic to a specific type of bias, there may be more targeted mitigation strategies that address a specific type of bias. Some of these limitations could be overcome by future systematic literature review (e.g., revisiting ad-hoc dimensions).

Choices within this design space must be balanced with potentially conflicting design considerations. For example, higher levels of INTRUSIVENESS may mitigate bias, but at the expense of user frustration in using the system. In addition, we have assumed that the TYPE OF DEBIASING INFORMATION is given a priori. However, the collection of this information within a system may necessitate its own design considerations. Systems that compute bias metrics based on user interaction sequences (e.g., [41, 42]) will have constraints on VISUAL REPRESENTATION and INTERACTION to ensure that the user's interactions adequately capture their cognitive process. Hence, this may conflict with bias mitigation strategies that involve altering that design.

6 Conclusion

We presented 8 dimensions of visualization design that can be manipulated to promote a less biased decision making process. The design space was developed based on related literature in cognitive science and visualization research. Because bias mitigation in VA is an emergent topic, few of these strategies have been evaluated. Hence, future work should begin to implement and assess the various strategies within this design space to understand in which contexts each strategy might excel or fail. Furthermore, it is important to understand the relationship and implications of combining multiple mitigation strategies.

This design space can inform the design of future visualization systems that can better support human decision making processes. It is our hope that this work may serve as a call to action within the visualization community to address problems around the design and evaluation of systems that can guard against people's inherent cognitive limitations.

Acknowledgments

This work was supported in part by the National Science Foundation grant IIS-1813281, and the Siemens FutureMaker Fellowship.

References

[1] H. R. Arkes. Costs and benefits of judgment errors: Implications for debiasing. *Psychological bulletin*, 110(3):486, 1991.

[2] K. Badam, Z. Zeng, E. Wall, A. Endert, and N. Elmqvist. Supporting team-first visual analytics through group activity representations. *Graphics Interface*, 2017.

[3] A. D. Balakrishnan, S. R. Fussell, and S. Kiesler. Do visualizations improve synchronous remote collaboration? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1227–1236. ACM, 2008.

[4] F. Bouali, A. Guettala, and G. Venturini. Vizassist: an interactive user assistant for visual data mining. *The Visual Computer*, 32(11):1447–1463, 2016.

[5] E. T. Brown, J. Liu, C. E. Brodley, and R. Chang. Dis-Function: Learning Distance Functions Interactively. *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 83–92, 2012.

[6] 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, 2017.

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

[8] W. S. Cleveland and R. McGill. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association*, 79(387):531–554, 1984.

[9] E. Dimara, G. Bailly, A. Bezerianos, and S. Franconeri. Mitigating the attraction effect with visualizations. *IEEE transactions on visualization and computer graphics*, 25(1):850–860, 2019.

[10] M. Englich and T. Mussweiler. Anchoring effect. *Cognitive Illusions: Intriguing Phenomena in Judgement, Thinking, and Memory*, p. 223, 2016.

[11] M. Feng, C. Deng, E. M. Peck, and L. Harrison. Hindsight: Encouraging exploration through direct encoding of personal interaction history. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):351–360, 2017.

[12] 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, 2019.

[13] S. Few. *Now you see it: simple visualization techniques for quantitative analysis*. Analytics Press, 2009.

[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. ACM, 2016.

[15] T. Green, W. Ribarsky, and B. Fisher. Visual analytics for complex concepts using a human cognition model. *2008 IEEE Symposium on Visual Analytics Science and Technology*, pp. 91–98, 2008. doi: 10.1109/VAST.2008.4677361

[16] C. Gutwin, M. Roseman, and S. Greenberg. A usability study of awareness widgets in a shared workspace groupware system. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work*, pp. 258–267. ACM, 1996.

[17] D. F. Halpern. Teaching critical thinking for transfer across domains: Disposition, skills, structure training, and metacognitive monitoring. *American psychologist*, 53(4):449, 1998.

[18] J. Heer and M. Bostock. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 203–212. ACM, 2010.

[19] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE transactions on visualization and computer graphics*, 14(6), 2008.

[20] R. J. Heuer Jr. *Psychology of Intelligence Analysis*. Washington, DC, 1999.

[21] R. J. Heuer Jr, R. J. Heuer, and R. H. Pherson. *Structured analytic techniques for intelligence analysis*. Cq Press, 2010.

[22] R. M. Hogarth. A note on aggregating opinions. *Organizational behavior and human performance*, 21(1):40–46, 1978.

[23] E. Horvitz. Principles of Mixed-Initiative User Interfaces. *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, (May):159–166, 1999. doi: 10.1145/302979.303030

[24] T. Jankun-Kelly. Using visualization process graphs to improve visualization exploration. In *International Provenance and Annotation Workshop*, pp. 78–91. Springer, 2008.

[25] T. Jankun-Kelly and K.-L. Ma. A spreadsheet interface for visualization exploration. In *Proceedings of the Conference on Visualization'00*, pp. 69–76. IEEE Computer Society Press, 2000.

[26] D. Kahneman and S. Frederick. A model of heuristic judgment. *The Cambridge Handbook of Thinking and Reasoning*, pp. 267–294, 2005.

[27] D. R. Kretz. Experimentally evaluating bias-reducing visual analytics techniques in intelligence analysis. In *Cognitive Biases in Visualizations*, pp. 111–135. Springer, 2018.

[28] P.-M. Law and R. C. Basole. Designing breadth-oriented data exploration for mitigating cognitive biases. In *Cognitive Biases in Visualizations*, pp. 149–159. Springer, 2018.

[29] B. J. McNeil, S. G. Pauker, H. C. Sox Jr, and A. Tversky. On the elicitation of preferences for alternative therapies. *New England Journal of Medicine*, 306(21):1259–1262, 1982.

[30] R. S. Nickerson. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220, 1998.

[31] C. North, R. May, R. Chang, B. Pike, A. Endert, G. A. Fink, and W. Dou. Analytic Provenance: Process + Interaction + Insight. *29th Annual CHI Conference on Human Factors in Computing Systems, CHI 2011*, pp. 33–36, 2011. doi: 10.1145/1979742.1979570

[32] D. Prelec, H. S. Seung, and J. McCoy. A solution to the single-question crowd wisdom problem. *Nature*, 541(7638):532, 2017.

[33] B. Saket, A. Endert, and C. Demiralp. Task-based effectiveness of basic visualizations. *IEEE transactions on visualization and computer graphics*, 2018.

[34] L. J. Sanna, N. Schwarz, and S. L. Stocker. When debiasing backfires: Accessible content and accessibility experiences in debiasing hindsight. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3):497, 2002.

[35] F. Sperrle, U. Schlegel, M. El-Assady, and D. Keim. Human trust modeling for bias mitigation in artificial intelligence.

[36] C. Symborski, M. Barton, M. Quinn, K. S. Kassam, C. Symborski, M. Barton, and M. Quinn. Missing: A serious game for the mitigation of cognitive biases. *Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2014*, pp. 1–13, 2014.

[37] D. A. Szafir. Modeling color difference for visualization design. *IEEE transactions on visualization and computer graphics*, 24(1):392–401, 2018.

[38] A. K. Thomas and P. R. Millar. Reducing the framing effect in older and younger adults by encouraging analytic processing. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 2(139), 2011.

[39] A. Tversky and D. Kahneman. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232, 1973.

[40] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.

[41] 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. *IEEE Conference on Visual Analytics Science and Technology (VAST)*, 2017.

[42] E. Wall, L. M. Blaha, C. Paul, and A. Endert. A formative study of interactive bias metrics in visual analytics using anchoring bias. *Proceedings of the 17th IFIP TC 13 International Conference on Human-Computer Interaction (INTERACT'19)*, 2019.

[43] E. Wall, S. Das, R. Chawla, B. Kalidindi, E. T. Brown, and A. Endert. Podium: Ranking data using mixed-initiative visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 2017.

[44] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1129–1136, 2007.