

Four Perspectives on Human Bias in Visual Analytics

Emily Wall*
Georgia Tech

Leslie Blaha†
Pacific Northwest
National Laboratory

Celeste Lyn Paul‡
Department of
Defense

Kristin Cook§
Pacific Northwest
National Laboratory

Alex Endert¶
Georgia Tech

ABSTRACT

Visual analytic systems, especially mixed-initiative systems, can steer analytical models and adapt views by making inferences from users' behavioral patterns with the system. Because such systems rely on incorporating implicit and explicit user feedback, they are particularly susceptible to the injection and propagation of human biases. To ultimately guard against the potentially negative effects of systems biased by human users, we must first qualify what we mean by the term *bias*. Thus, in this paper we describe four different perspectives on human bias that are particularly relevant to visual analytics. We discuss the interplay of human and computer system biases, particularly their roles in mixed-initiative systems. Given that the term *bias* is used to describe several different concepts, our goal is to facilitate a common language in research and development efforts by encouraging researchers to mindfully choose the perspective(s) considered in their work.

Index Terms: H.5.0 [Information Systems]: Human-Computer Interaction—General

1 INTRODUCTION

Visual analytic applications foster exploratory data analysis by combining computational techniques with interactive visualizations. A critical aspect of visual analytics is understanding how to incorporate user feedback. Such human-in-the-loop approaches to analysis allow people to leverage their domain expertise and reasoning abilities to make sense of data and gain insight. Particularly relevant to *mixed-initiative systems* [20], the principles that frame our understanding of these systems include a balance of responsibility between systems and people (i.e., an understanding of who does which specific tasks). When successful, machines and people work together to engage in a dialog about the data.

In visual analytics, we observe a trend in how user interaction is incorporated. First, systems can take direct input from users to change views, direct analytic models, and perform other analytic tasks. Second, we observe a rise in systems that learn from people's interactions and behaviors, build user models, and adapt the system based on the system's interpretation of the user's interests, analytic process, etc. Generally, both approaches to incorporating user interaction result in people guiding the analytic process by adapting the ways data is computed, visualized, and otherwise transformed.

However, what if this human guidance is faulty? While people have immense sensemaking and reasoning abilities as well as valuable domain expertise, we also know that people are susceptible to innate biases. In current system designs and implementations, these biases can be incorporated and propagated throughout the system. For example, if someone exhibits confirmation or anchoring bias while analyzing data, the analytic models and views could amplify

the bias and lead to potentially biased or incorrect results. How do we build mixed-initiative visual analytic systems that are aware of this challenge and ideally guard against it? Recent work has begun to address how bias materializes in visual analytics [8, 16, 47].

An important first step toward understanding and leveraging bias is to review how we might define and formalize human bias in the scope of mixed-initiative visual analytics. Cognitive, behavioral, and social sciences have described many ways bias can occur in people's analytic processes [23, 35, 38], decision-making strategies [3, 9], and other behaviors. Motivated by the overloaded use of the term "bias" to describe different models and concepts, we thus present four perspectives on human bias, including (1) bias as a cognitive processing error, (2) bias as a filter for information, (3) bias as a preconception, and (4) bias as a model mechanism. These perspectives represent four commonly adopted takes on the term "bias." The four perspectives are not mutually exclusive; rather, they present different, potentially overlapping perspectives on bias in the context of visual analytics.

To more concretely discuss how bias can affect visual analytics, consider the following example. Suppose Susan is using a visual analytic tool to explore possibilities for purchasing a new home. She uses the tool to browse photos, explore different areas of the city, and refine her understanding of what features of a home are important to her. From her exploration, she intends to go view the homes in person and ultimately make a purchasing decision. Throughout the following sections, we will describe how each perspective on bias can impact Susan's process and visual analytics in general. For each perspective, we provide a brief description, present an example scenario, and discuss how it influences visual analytics.

2 BIAS AS A COGNITIVE PROCESSING ERROR

Description: From heuristics and bias research, bias is an error resulting from an unconscious deviation from rational behavior. Cognition is frequently conceptualized as a dual-process [7]. The two processes are often termed "intuition" and "reason" [22], the former being responsible for making quick, automatic decisions, and the latter being responsible for making deliberate, reflective decisions. It is one's quick judgments that are subject to errors.

Stanovich and West referred to the two cognitive processes as system 1 (intuition) and system 2 (reason) [40]. In this analogy, system 1 is largely subconscious and prone to making errors (bias), while system 2 is responsible for recognizing and correcting errors through intentional deliberation. These types of errors result from shortcuts in cognition, broadly referred to as heuristics [22]. Bias then is described as the method or mechanism by which the error occurs. However, the process of heuristic decision making does not always lead to errors; it usually facilitates fast decision making.

Example: From this perspective, there are dozens of types of bias. One such example is anchoring bias [45], which refers to the tendency to be heavily reliant on an initial value or anchor. It is analogous to a center of mass: people are unlikely to strongly deviate from center. In Susan's home-buying scenario, she will likely be subject to anchoring bias during the price negotiation of her purchase. That is, the home's initial list price forms an anchor point and will thus subconsciously impact the amount she is willing to offer. Susan's offer for the home might have been very different had she made an offer given a different initial list price. She might

*e-mail: emilywall@gatech.edu

†e-mail: leslie.blaha@pnnl.gov

‡e-mail: clpaul@tycho.nsc.mil

§e-mail: kris.cook@pnnl.gov

¶e-mail: endert@gatech.edu

even pay more money for the same home due to the tendency not to strongly deviate from the anchor point. Thus, systems apprised of probable cognitive errors like anchoring bias have the potential to help users make better decisions by guarding against such errors.

Relevance to Visual Analytics: Common heuristic errors include confirmation bias [31] which describes the way people tend to accept confirmatory evidence of a pre-existing hypothesis and dismiss contrary information. Another common error is availability bias [44], where people tend to rely more heavily on information that is easily remembered (e.g., most recent). Similarly, the attraction effect [21] describes the tendency for a decision to be influenced by an inferior alternative. Collectively, these errors shape the way people search for and interpret information. Recently, Dimara et al. [8] showed that the attraction effect is present in users of information visualizations. Hence, bias impacts users outside of laboratory decision making studies and can lead to incorrect decisions and inefficiencies in visualization-supported analytic processes.

3 BIAS AS A FILTER FOR INFORMATION

Description: Bias acts as a filter through which we manage and perceive information. The challenge of information overload [28] motivates this analogy. Information overload, now commonly leveraged in consumer research to influence purchasing behavior, refers to a point beyond people's cognitive and perceptual limits where performance and decision making suffer [27]. One's filter or bias thus determines how sensory information is distinguished and interpreted [15].

The literature on goal-directed attention and resource allocation posits that all perception is guided by top-down influences, such as the allocation of endogenous attention [10, 36, 41]. Top-down perception governs which sensory information is identified in a scene based on goals. Bias does not make for a purely objective filter for information, however. Heuer refers to perception as an "active" process (compared to passive) that "constructs" reality (rather than records it) [19]. Similarly, obvious or important information is sometimes filtered out. For example, in one classic selective perception task, participants were asked to count how many times basketball players on a team passed the ball [39]. Most participants count the appropriate number of passes but about half fail to perceive a glaringly misfit player walk across the court. In contrast to top-down perception, bottom-up perception refers to the way external factors influence attention [37]. When there is a loud noise or someone says your name across the room, you notice despite top-down attentional and perceptual focus.

Example: In our home-buying scenario, Susan may experience information overload [28] as she explores homes on the market in a visual analytic tool. She might see hundreds of homes available in the area, each with dozens of attributes. Thus, her filter or bias will govern which information she perceives and which she dismisses. By leveraging knowledge about people's perceptual strengths and limitations, systems can present information in ways that are easy for users to understand.

Relevance to Visual Analytics: A great deal of research in perception has been leveraged by researchers in information visualization and visual analytics to present information in ways that are most perceptually accessible [12]. Preattentive processing theory [42], for example, describes the nature and limits of visual information processing. In creating visual representations of data, this is often used by designers as a guide to prevent overwhelming users' perceptual limitations. Similarly, Gestalt principles [24] refer to the relationships inferred by the visual system based on proximity, groupings, symmetry, etc. between visual elements. Thus, understanding how people's filters work can inform things like which visual widgets or elements to place in close proximity to one another or which graph layout algorithm is most appropriate.

4 BIAS AS A PRECONCEPTION

Description: Analysts approach mixed-initiative systems bringing all their experiences and internal influences that unconsciously shape their approaches to the analysis process. This, in turn, influences the ways they interact with systems. The consequence is that the user model within the system, the analytic products, and provenance may be shaped by each individual's unconscious biases. These types of bias may seem to have little to do directly with the task at hand. Yet, because they shape the person, there is a high likelihood they can influence mixed-initiative sensemaking.

Unconscious biases arise in a number of ways. They derive from a person's cultural beliefs and traditions, which include their implicit assumptions and expectations regarding stereotypes. Unconscious biases result from general self-confidence or self-esteem, as well as comfort or familiarity level with the capabilities of the machine analytics and interface functions. Related personality traits render some people more risk seeking or risk averse, shaping how they push boundaries exploring a space of hypotheses or push the capabilities of the computational system. These characteristics are thus seen as a source of individual variability between people.

Example: Susan is avoiding listings for houses downtown in the city. Having lived in the suburbs for many years, Susan assumes that neighborhoods near downtown have higher crime rates and lower economic stability. She believes she should not make a housing investment there. The availability of recent census results and police reports within the real estate analytic tools enable Susan to explore her assumptions and refine her thinking. A mixed-initiative system may detect her avoidance of downtown properties and could prompt her to challenge her assumptions with the related data.

Relevance to Visual Analytics: Unconscious biases shape analysts' assumptions and stereotypes about analytical tools and mixed-initiative aids, and they shape assumptions and stereotypes about the data / analytical subjects (e.g., presumed reliability or trustworthiness of certain sources). Implicit attitudes shape the formulation of hypotheses and the questions about the assumptions and the consequences of those hypotheses. Klein and colleagues posited that the entire sensemaking process begins with a practitioner framing the problem, and the selected framework, however minimal, then shapes what an analyst thinks about and what structure they think with [23]. Frames reflect a perspective an analyst takes to make sense of data or to solve a problem. As implicit attitudes shape an analyst's perspective, they shape the analyst's frames, thereby shaping the sensemaking process.

Expertise, derived from general experience as well as explicit training, further shapes the analytical process and is shaped by implicit biases. Expertise can impact expectations and perceptions of a mixed-initiative system and the interpretations of the information visualizations under consideration. Expertise in forensic analytics, for example, may make analysts more conservative in their judgments, shaped in part by their expert understanding of the consequences of their decisions. Expertise often also provides the user with a better understanding of the limitations of the analytical tools or data collection practices, which can shape more nuanced interpretations during the analysis process.

Because they are built to record a number of different types of user behaviors throughout the analysis process, mixed-initiative systems may be particularly well-positioned to aid in the assessment of unconscious biases of analysts. We argue that it is possible for a mixed-initiative system to capture and integrate unconscious, preconception biases into analytics through the user model.

5 BIAS AS A MODEL MECHANISM

Description: Bias is the term often used in cognitive modeling to describe a decision boundary or a tendency toward one response option over another. Cognitive models are mathematical and computational approaches to formally describe mechanisms supporting

perception, memory, decision making, and other cognitive functions [5]. A number of these models include a mechanism explicitly called bias, or they use a combinations of mechanisms to capture the ways the aforementioned types of bias manifest in measurable behaviors, like response choice and speed. Models with explicit bias mechanisms often contain a bias parameter or measure bias as a relationship between parameters. Here, we will review two major perspectives on bias as a model mechanism, one which formalizes bias within models of mental organization and another which formalizes bias in models of decision making dynamics. Both types of behavior are necessary in visual analytics, as analysts work through their sensemaking processes of organizing information and weighing evidence against potential hypotheses and interpretations. As interactive visual analytic systems aid in the externalization of analysts' mental models, model mechanisms can help us interpret how bias is reflected in the patterns and dynamics of their interactions.

One approach to modeling bias addresses the question: where do people mentally “draw the line” between one response option and another when performing an analytic task? Many models of perceptual choice or organization describe information representation with two mechanisms. One mechanism is spatial organization that groups pieces of information by similarity/proximity; like objects are close in space or clustered together. The second mechanism is at least one boundary that divides the space into response regions; object labels or choices are made according to the response regions defined by the boundary. Examples of these models include the theory of signal detection for finding signals in noise [18, 26] or categorization models [33] for clustering and labeling groups of objects. Bias in these models is described by a weighting of boundary regions (representing bias toward certain responses) or a feature weighting (representing how much the respondent emphasized certain features over others).

Another major use of bias parameters is found in models of information processing dynamics behind the time to make a decision. These dynamic decision models characterize the choice between two options as a stochastic process whereby information about the options is incrementally sampled and accumulated, often in a random walk fashion, until some threshold is reached for one of the response options [6]. The evidence accumulation process governs a person's response speed and is influenced by the salience and complexity of the choice options. Bias in these models is captured by the relationship between the starting value of the evidence accumulators and the response thresholds. If the accumulator starts at zero, then the process is not biased; all responses are equally likely. If the bias parameter is non-zero, then the process is biased toward the response threshold closer to the bias value. This bias mechanism captures behaviors wherein some responses, correct or erroneous, are selected more frequently or more quickly than others.

Example: Homes for sale are comprised of a large number of attributes drawn from real estate descriptions. Susan is likely to have certain features along which she is organizing the options available on the market, such as number of bedrooms, number of bathrooms, basement square footage, and proximity to schools. This forms a four-dimensional representation space into which the houses can be organized. If she is weighing numbers of bedrooms and bathrooms equally, we can describe her decision bias as equidistant from the category centroids or close to zero. However, Susan has strong opinions about basement square footage and proximity to schools. Based on how she organizes houses into desirable and undesirable categories, we might use models to infer that she is biased toward liking houses that are within a 10 minute walk to schools but have small basements less than 400 square feet.

Relevance to Visual Analytics: Visual analytic systems designed to support data exploration capture an externalization of the analyst's mental organization in the form of interaction. By leveraging analytic provenance [32], researchers can better understand users'

strategies [9], processes that led to insights [17], and ultimately better support the sensemaking process [48]. Different spatial layouts and data encodings (colors, shapes, etc.) reflect mental organization patterns, including perception of similarity between data points. Characterizing the biases in this mental organization process provides a quantifiable way to describe the information representation space and decision boundaries. For example, we can use the perceptual organization models to infer if the analyst is biased toward some data attributes or certain clusters/labels. We could use the sequential sampling model to identify biases in how analysts are weighing the relative utility or value of a piece of evidence.

From the perspective that bias is a model mechanism, we can also formally characterize bias from the other three perspectives described in Sections 2–4. Although these models are implemented in a way that is rather agnostic to errors in reasoning, the bias parameters enable inferences about how errors from decision heuristics occur. For example, anchoring bias would be captured as a bias toward one of the response thresholds close to the anchor value in a decision dynamics model. Bias as a filter can be formalized as a bias node or parameter in a neural network or hierarchical model of vision [43]. This would reflect the way information might be differently sampled by an analyst based on the goal-related task s/he is performing. Preconception bias can be included in models as latent factors or correlates of measurable behaviors. As latent factors, biases such as gender or race stereotypes can modulate other mechanisms in the mental models, such as the organization of similar objects or response preferences [46].

6 DISCUSSION

These four perspectives of bias illustrate the diversity in how people process information and form a model of the world. Each are valid perspectives that greatly shape how bias is framed in visual analytics research. However, the multiplicity in definitions sometimes leads to challenges in sharing and collaboration due to a lack of common ground. One goal of this paper is to present these definitions, so that we as a community have a starting point for discussing how these perspectives fit within the visual analytic research agenda. Additionally, when considering all of these perspectives, the space in which to study bias in visual analytics increases dramatically. This leads to several open challenges and opportunities for the visual analytics community.

6.1 Does bias endanger mixed-initiative visual analytics?

Visual analytic applications continue to model users and adapt interfaces, visualizations, and analytic models based on their interactions. However, how do such systems differentiate between valuable subject matter expertise (which should be incorporated), and biased input? Without such techniques for identifying and guarding against biased input, applications run the risk of showing users biased views of their data that correspond to what they want to use, rather than truthful representations of the information.

For example, in model steering situations, user input guides analytic models to focus on salient aspects of the domain being studied [11]. Without guarding against potentially biased user input, the system may overfit the model to the biased input. The result may be a system that shows users the views they want to see, but is essentially an “echo chamber” for their own biases.

A recent example that showcases the potential consequences of human bias in systems is the AI chatbot, Tay [1, 25]. The artificial intelligence was intended to be a friendly chatbot that appealed to young adults. The underlying model was continually trained by incoming tweets, causing Tay to tweet increasingly racist and misogynistic messages shortly after going live. While a vulnerability in Tay was exploited, the chatbot nonetheless conveys what can happen when human bias is introduced unchecked into a system.

An awareness of these potential risks will help us develop better systems, and ultimately foster better data-driven decisions.

One approach for making the distinction between valuable domain expertise and biased input might be to consider the consistency or inconsistency of a user's interaction sequences. More sophisticated approaches could be derived by studying the differences in interaction sequences of domain experts and novices who are biased.

6.2 How to keep the machine “above the bias”?

Designing mixed-initiative visual analytic systems to reduce negative effects of biased user input is an interesting and important line of research leveraging our bias classifications. As noted by Friedman and colleagues, there are three types of bias that can influence computer systems: pre-existing, technical, and emergent biases [13, 14]. Pre-existing bias arises from the attitudes or societal norms/practices that the software designers might impart into system designs. This is akin to our bias as a preconception perspective. Concerted efforts can be made to address pre-existing bias throughout the visual analytics design process, such as using the recent GenderMag method to address gender biases in interface designs [4].

Technical biases are a consequence of technical considerations, such as choice of hardware or algorithm. Computational technical biases are unique from the various definitions of human bias we summarized herein. But because they will contribute to biases in mixed-initiative system performance, careful technical choices should be made and appropriate details should be made available to the user to facilitate informed interpretation of system behaviors.

Emergent biases arise from the use of a system, resulting from changing context or knowledge in which a system is being used. Friedman argues that these are more difficult to know in advance or even identify in practice [14]. Emergent biases are highly likely to occur in mixed-initiative systems, particularly as the interface or algorithms are shaped by any of the aforementioned biases that are influencing the user's interactions. Theoretically, the role of the machine is to be unbiased and to present a rational result based on clear rules. However, there are limitations to this approach, namely the lack of tacit knowledge and analytic context that cannot be easily modeled. This has led to the rise of user-driven machine learning that goes beyond a “supervisory” role in training [2]. Yet, as soon as the human is re-introduced into the system, the rationality of the machine is affected. How can we judge when this human-machine teaming is succeeding or failing?

We propose that mixed-initiative systems are uniquely suited to aid in the identification and mitigation of emergent biases, exactly because mixed-initiative systems reflect the user's analytic process. To do this then, we must be able to correctly interpret the user's biases as they are captured by the computational system. The four perspectives we have outlined will help the bias interpretation process. Each provides a way to identify how that source of bias plays out in the analytic process. To the degree that formal models are available for each bias perspective, those can be integrated into the system for more automated interpretations.

6.3 Could the mixed-initiative system impart bias to the user?

Yes. A less-emphasized aspect of emergent bias is that the structure of the user interface may influence and bias the interactions of the user. Reliance on machine automation and automated decision aids can result in automation bias. This is the heuristic use of automation instead of more vigilant information seeking and decision making [29, 30, 34]. The errors resulting from automation bias are concerning for mixed-initiative systems, wherein those errors might be integrated into the analytic results/visualizations or even the analytic processes. Of particular concern in this domain are automation commission errors. These errors are inappropriate actions resulting from over-attending to automated aids without attention to the

context or other critical environmental information sources. Commission errors occur when a user accepts the recommendation of some machine analytics even when there is contrary evidence from other information sources, either internal or external to the analytics system.¹ The design of an interactive analytic interface may lend itself to overemphasizing some analytic results or mixed-initiative recommendations, such as highlighting recommendations or even in the size or color that might make some recommendations stand out over others. Automation bias to accepting the most strongly emphasized recommendations could lead the analyst down a biased analysis path. Does the system or the user bear the responsibility for mitigating automation bias? We argue that if mixed-initiative systems can cultivate emergent biases in both the machines and the users, then mixed-initiative systems also offer new opportunities for humans and machines to team up to mitigate negative bias effects.

6.4 Is bias good or bad?

The term bias tends to carry a negative connotation. It is perceived as something that we should strive to eradicate. However, bias is not always bad. Each of the four perspectives on bias differs in how it impacts cognitive and perceptual processes.

From the perspective that bias is an error, we should work to minimize it; however, it should not be confused with the heuristic decision making processes that lead to such biases. We emphasize that heuristic decision making is not inherently bad; it usually results in more efficient decision making. Thus, it is imperative that in attempting to mitigate bias as an error, we do not unduly limit heuristic decision making processes in general.

From the perspective that bias is a model mechanism, it is neither good nor bad. In this case, it is an objective characterization of the decision making process. While the decision making process itself may be suboptimal or erroneous (as is the case of bias as an error), here bias just describes the boundary between response options.

From the perspective that bias is a filter and the perspective that bias is a preconception, it can be both beneficial and detrimental depending on circumstances. Perceptual filters prevent us from experiencing information overload. However, they can also cause us to inadvertently filter out information relevant to a given decision. Unconscious biases like innate risk-aversion tendencies can help us to make deliberate, mindful decisions, but on the other side of the spectrum can lead to impulsive high-risk decisions. Thus, because different perspectives on bias vary widely in their potential benefits or risks, it is imperative to thoughtfully define the perspective and scope considered for bias detection or mitigation efforts.

7 CONCLUSION

Bias is a particularly important consideration in the design of mixed-initiative visual analytic systems, where biased human input can shape analytical models. Thus, in this paper we have described four perspectives on bias particularly relevant to such human-machine collaborative systems. We hope that by discussing and differentiating these perspectives on the overloaded term “bias,” researchers and developers can thoughtfully define which perspective they take in their work on bias.

ACKNOWLEDGMENTS

The research described in this document was sponsored by the U.S. Department of Defense. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

¹Commission errors are contrasted with automation omission errors, which occur if the human-machine team fails to respond to system irregularities or the system fails to provide an indicator of a problematic state. In visual analytics, an omission error could occur if a system “knows” an algorithm might be mis-matched to a data type but does not alert the analyst.

REFERENCES

- [1] F. Alaieri and A. Vellino. Ethical decision making in robots: Autonomy, trust and responsibility. In A. Agah, J.-J. Cabibihan, A. M. Howard, M. A. Salichs, and H. He, eds., *Social Robotics: 8th International Conference*, pp. 159–168. Springer International Publishing, Kansas City, MO, 2016.
- [2] S. Amershi, M. Cakmak, W. B. Knox, and T. Kulesza. Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 35(4):105–120, 2014.
- [3] 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. doi: 10.1109/TVCG.2014.2346575
- [4] M. Burnett, S. Stumpf, J. Macbeth, S. Makri, L. Beckwith, I. Kwan, A. Peters, and W. Jernigan. GenderMag: A method for evaluating software’s gender inclusiveness. *Interacting with Computers*, 28(6):760–787, 2016.
- [5] J. R. Busemeyer and A. Diederich. *Cognitive Modeling*. Sage, Los Angeles, CA, 2010.
- [6] J. R. Busemeyer and J. T. Townsend. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432–459, 1993.
- [7] S. Chaiken and Y. Trope. *Dual-Process Theories in Social Psychology*. Guilford Press, New York, 1999.
- [8] E. Dimara, A. Bezerianos, and P. Dragicevic. The attraction effect in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):471–480, 2017.
- [9] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. R. Lipford, and R. Chang. Recovering reasoning process from user interactions. *IEEE Computer Graphics & Applications*, May/June(March):52–61, 2009.
- [10] H. E. Egeth and S. Yantis. Visual attention: Control, representation, and time course. *Annual Review of Psychology*, 48(1):269–297, 1997.
- [11] A. Endert, W. Ribarsky, C. Turkay, B. Wong, I. Nabney, I. D. Blanco, and F. Rossi. The state of the art in integrating machine learning into visual analytics. In *Computer Graphics Forum*. Wiley Online Library, 2017.
- [12] J.-D. Fekete, J. Van Wijk, J. Stasko, and C. North. The value of information visualization. *Information Visualization*, pp. 1–18, 2008.
- [13] B. Friedman. Value-sensitive design. *Interactions*, 3(6):16–23, 1996.
- [14] B. Friedman and H. Nissenbaum. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)*, 14(3):330–347, 1996.
- [15] J. P. Frisby and J. V. Stone. *Seeing: The Computational Approach to Biological Vision*. The MIT Press, Cambridge, MA, 2010.
- [16] D. Gotz, S. Sun, and N. Cao. Adaptive contextualization: Combating bias during high-dimensional visualization and data selection. *Proceedings of the 21st International Conference on Intelligent User Interfaces - IUI '16*, pp. 85–95, 2016. doi: 10.1145/2856767.2856779
- [17] D. Gotz and M. X. Zhou. Characterizing users’ visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009.
- [18] D. M. Green, T. G. Birdsall, and W. P. Tanner Jr. Signal detection as a function of signal intensity and duration. *The Journal of the Acoustical Society of America*, 29(4):523–531, 1957.
- [19] R. J. Heuer Jr. *Psychology of Intelligence Analysis*. Washington, DC, 1999.
- [20] E. Horvitz. Principles of mixed-initiative user interfaces. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 159–166, 1999.
- [21] J. Huber, J. W. Payne, and C. Puto. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9(1):90–98, 1982.
- [22] D. Kahneman and S. Frederick. A model of heuristic judgment. *The Cambridge Handbook of Thinking and Reasoning*, pp. 267–294, 2005.
- [23] G. Klein, B. Moon, and R. R. Hoffman. Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems*, 21(5):88–92, 2006.
- [24] K. Koffka. *Principles of Gestalt Psychology*, vol. 44. Routledge, 2013.
- [25] P. Lee. Learning from Tay’s introduction. <https://blogs.microsoft.com/blog/2016/03/25/learning-tays-introduction/>, 2016.
- [26] N. A. Macmillan and C. D. Creelman. *Detection Theory: A User’s Guide*. Psychology Press, 2004.
- [27] N. K. Malhotra. Information load and consumer decision making. *Journal of Consumer Research*, 8(4):419–430, 1982.
- [28] J. T. Milord and R. P. Perry. A methodological study of overloadx. *The Journal of General Psychology*, 97(1):131–137, 1977.
- [29] K. L. Mosier and L. J. Skitka. Human decision makers and automated decision aids: Made for each other. In R. Parasuraman and M. Mouloua, eds., *Automation and Human Performance: Theory and Applications*, pp. 201–220. Lawrence Erlbaum Associates, Mahwah, NJ, 1996.
- [30] K. L. Mosier and L. J. Skitka. Automation use and automation bias. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 43, pp. 344–348. Sage Publications, 1999.
- [31] R. S. Nickerson. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220, 1998.
- [32] 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
- [33] R. M. Nosofsky. Stimulus bias, asymmetric similarity, and classification. *Cognitive Psychology*, 23(1):94–140, 1991.
- [34] R. Parasuraman and D. H. Manzey. Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52:381–410, 2010.
- [35] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. *Proceedings of International Conference on Intelligence Analysis*, 2005:2–4, 2005.
- [36] M. I. Posner. Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32(1):3–25, 1980.
- [37] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2(11):1019–1025, 1999.
- [38] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim. Knowledge generation model for visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1604–1613, 2014.
- [39] D. J. Simons and C. F. Chabris. Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9):1059–1074, 1999.
- [40] K. E. Stanovich and R. F. West. Advancing the rationality debate. *Behavioral and Brain Sciences*, 23(5):701–717, 2000.
- [41] A. Torralba, A. Oliva, M. S. Castelhano, and J. M. Henderson. Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological Review*, 113(4):766–786, 2006.
- [42] A. Treisman. Preattentive processing in vision. *Computer Vision, Graphics, and Image Processing*, 31(2):156–177, 1985.
- [43] J. K. Tsotsos. *A Computational Perspective on Visual Attention*. MIT Press, Cambridge, MA, 2011.
- [44] A. Tversky and D. Kahneman. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232, 1973.
- [45] A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185:1124–1131, 1974.
- [46] J. Vandekerckhove. A cognitive latent variable model for the simultaneous analysis of behavioral and personality data. *Journal of Mathematical Psychology*, 60:58–71, 2014.
- [47] 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. To appear.
- [48] K. Xu, S. Attfield, T. Jankun-Kelly, A. Wheat, P. H. Nguyen, and N. Selvaraj. Analytic provenance for sensemaking: A research agenda. *IEEE Computer Graphics and Applications*, 35(3):56–64, 2015.