

Beyond Visual Analytics: Human-Machine Teaming for AI-Driven Data Sensemaking

John Wenskovitch*
Corey Fallon†
Kate Miller‡

Pacific Northwest National Laboratory

Aritra Dasgupta§
New Jersey Institute of Technology

ABSTRACT

Detect the expected, discover the unexpected was the founding principle of the field of visual analytics. This mantra implies that human stakeholders, like a domain expert or data analyst, could leverage visual analytics techniques to seek answers to known unknowns and discover unknown unknowns in the course of the data sensemaking process. We argue that in the era of AI-driven automation, we need to recalibrate the roles of humans and machines (e.g., a machine learning model) as teammates. We posit that by realizing human-machine teams as a stakeholder unit, we can better achieve the best of both worlds: automation transparency and human reasoning efficacy. However, this also increases the burden on analysts and domain experts towards performing more cognitively demanding tasks than what they are used to. In this paper, we reflect on the complementary roles in a human-machine team through the lens of cognitive psychology and map them to existing and emerging research in the visual analytics community. We discuss open questions and challenges around the nature of human agency and analyze the shared responsibilities in human-machine teams.

1 INTRODUCTION

In this paper, we present our human-machine teaming vision as a communication paradigm for AI-driven data sensemaking and analyze the role of visual analytics [33] in realizing that vision. The founding principle of visual analytics, *detect the expected, discover the unexpected*, was anchored on meeting the needs for serendipitous data-driven discoveries by leveraging the best of both worlds of computing and the human vision system.

With the advent and growing adoption of AI and machine learning, there is a growing need to delineate human and machine responsibilities and rethink the role of visual analytic interfaces as a communication channel between the human and machine stakeholders. For the rest of this paper, we will use “machine” to denote an AI agent or a machine learning model. We argue that in the era of AI-driven automation, the roles of humans and machines need to be recalibrated as teammates throughout the foraging and sensemaking stages of the data analysis process [28]. Our vision for future human-machine interfaces aided by visual analytics is the following:

Visual analytic interfaces will act as a mutually intelligible communication channel between the human and the AI/ML model, so that human and machine teammates are in sync with their respective roles, responsibilities, and responsiveness towards each other’s actions.

This vision captures the need for bidirectional communication in human-machine systems. Bidirectional communication is im-

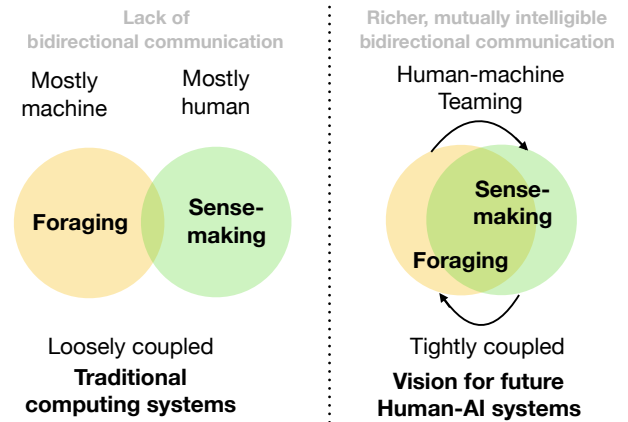


Figure 1: Future AI-driven data sensemaking [28] techniques will benefit from a richer human-machine communication where foraging tasks are more tightly coupled with an AI or machine learning model’s learning about human intent and reasoning.

portant for productive human-machine interaction that can benefit all team members and, consequentially, improve overall team performance. The human benefits from a machine that clearly communicates its capabilities, limitations, and actions both explicitly (e.g., transparent display design) [41] and implicitly (e.g., consistent predictable behavior) [7]. This communication is one of the vital system capabilities that allows a machine to function as a teammate by working to establish and maintain common ground with its human partner [13]. The machine benefits from the human’s corrective feedback by learning and improving its performance. The human’s increased understanding of machine capabilities and improved machine performance from corrective feedback leads to appropriate trust and reliance on the machine. Consequently, there is an overall improvement in human-machine system performance.

We are already witnessing real-world examples of such interaction in systems based on conversational AI [21]. We posit that the expressive power of visualization can facilitate such communication and allow humans and machines to take equal responsibility in the data sensemaking process. To realize this vision, we propose that the human in this partnership may assume one of several roles in the human-machine team depending on the sensemaking task. Each role is defined by a unique set of task goals and associated interactions with the machine. We also suggest that, regardless of the role the human assumes, this new human-machine team paradigm requires the human to forfeit some control to the machine. We suggest that this loss of control may pose a challenge to the human-machine team and offer some suggestions to overcome this challenge.

2 CONCEPTUALIZING HUMAN-MACHINE TEAMS

The traditional paradigm of computing-mediated data analysis focused primarily on a hypothesis-driven approach, in which information foraging was mainly a machine-level task (Fig. 1), whereas

*john.wenskovitch@pnnl.gov

†corey.fallon@pnnl.gov

‡kate.miller@pnnl.gov

§aritra.dasgupta@njit.edu

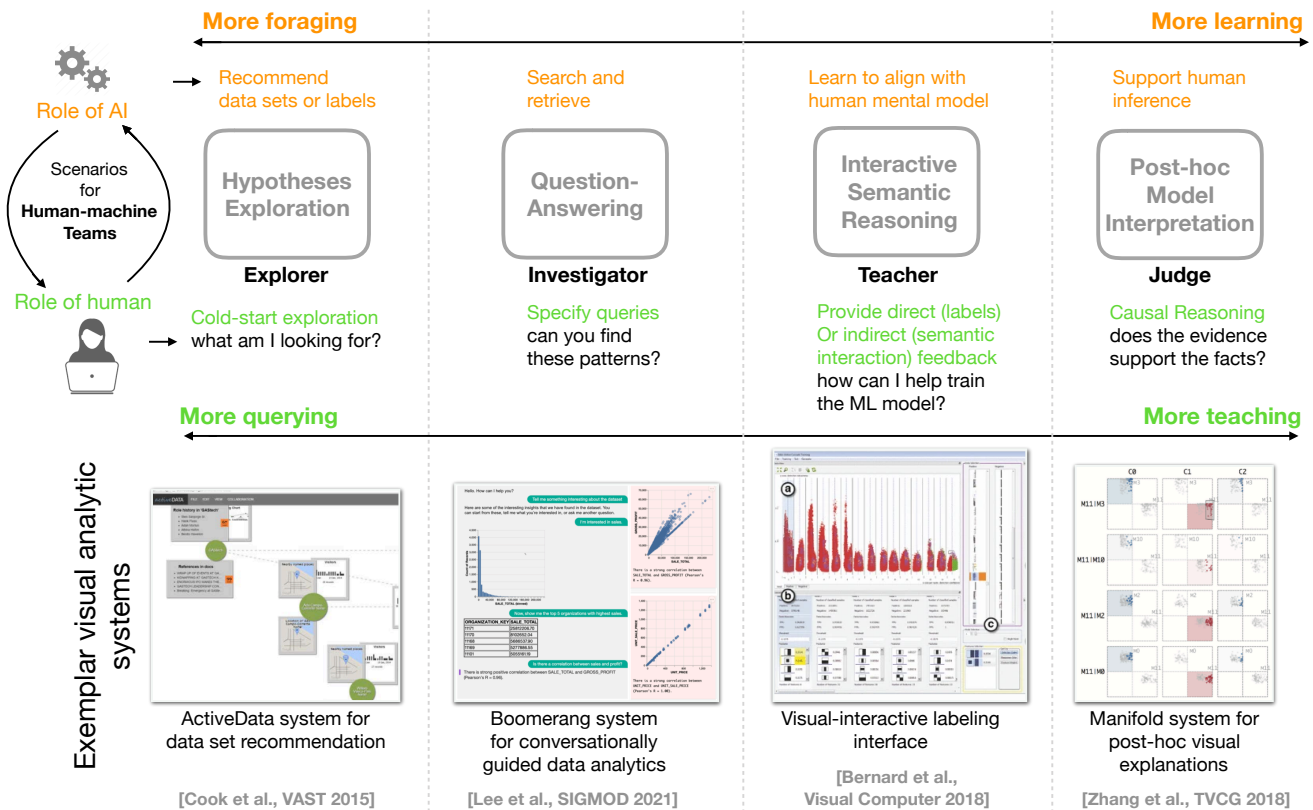


Figure 2: Analyzing the roles of humans and machines as teammates in diverse data analysis scenarios.

sensemaking was the responsibility of the human stakeholders. A domain expert or a data analyst had particular hypotheses and/or research questions and relied on analytics to explore them. The analysts needed to understand the strengths and limitations of the analytic and identify when its output might be incorrect. This understanding allowed the analyst to build appropriate trust in the tool and appropriate use. It was also important for the analyst to maintain their skills so that they could identify problems with the analytic and find workarounds if the tool was not providing accurate output.

We are transitioning into a new paradigm of human-machine teaming. In this new paradigm, the technology has the ability to learn and identify new patterns in the data and allow the analyst to generate new research questions. In its ability to learn from the analyst, the tool behaves more like a teammate than a tool and requires new roles and responsibilities for the human. We suggest that the human may assume one of several roles in the human-machine teaming paradigm depending on the sensemaking task (Fig. 2). In this section, we describe each role, machine as teammates, and discuss how visual analytic systems can facilitate team communication throughout the data analysis process.

2.1 Roles as Teammates

Explorer: Within a human-machine team, each entity has a role to play in the foraging process, with a human focus on creativity to understand what relevant information they are looking for and a machine focus on computation and enrichment to identify and provide relevant information. The creativity of the human is particularly important at the beginning of exploration, as there is little or no existing information to direct a search: a cold-start scenario. This role still continues throughout the exploration process, as the human works orient to the data landscape and begin to connect individual pieces of information into a coherent structure. In the old paradigm, the human was then responsible for the data searches that follow this

creative exercise. This responsibility could take the form of manually inspecting documents, structuring database queries, or filtering a returned dataset. In the human-machine teaming paradigm, the machine can play a role in initiating parts of these foraging actions. We see examples of mixed-initiative visual analytic systems that instantiate such foraging through recommendations of datasets [8,9], labels [14], or visualizations [44], or by performing preprocessing on a dataset to ensure that the most relevant values are prepared for further analysis. Though we are far from the ability to instill an adequate degree of “creativity” into a machine assistant to operate independently, the machine can still work semi-autonomously to support hypothesis exploration.

Investigator: Following a process of foraging for data, a human-machine team can begin to search through the data to identify connections, detect anomalies, and synthesize findings. This often takes the form of human-directed investigation, though from a Question-Answering perspective [20,21] rather than as an information retrieval challenge. The difference from the pure exploration scenario is in the degree of specificity that a human analyst has in asking questions: they are primarily trying to seek answers for known unknowns. This mode can be construed as a hybrid of the traditionally disparate modes of exploratory and confirmatory data analysis process, and require deep, inferential reasoning from analysts for validating the answers provided or patterns detected by a model [17]. In the old paradigm, the machine was responsible for searching through data and retrieving answers to queries, while the human guided this process through a sequence of instructions, directing the machine towards finding the necessary data. In a human-machine team however, the machine can take on some further autonomy to identify the interests of the human, automatically searching for and retrieving information that has potential relevance to the exploration. The machine can further assist by placing this new supporting evidence into context in the display for the human to evaluate and provide feedback.

When questions are known a priori, the expressive power of natural languages can be used to aid the communication process during investigative scenarios. In the spirit of question-answering systems like Boomerang [21], we can observe the recent interest in the visual analytics community towards integrating natural language interfaces for guiding the design and refinement of visualizations [31,42].

Teacher: As part of a human-machine team, the analyst has additional communication and monitoring responsibilities. Although these teamwork activities are vital to the health of the human-machine team, there is a cost associated with this coordination in terms of time and effort. In the human-machine teaming paradigm the analyst must communicate with the tool in a new way to provide the technology with corrective feedback. By adding a feedback loop to the interaction, interactive ML [4, 12, 39] requires the analyst to teach the tool, implicitly or explicitly. This new responsibility may detract from the time the analyst could spend developing their expertise. For example, providing basic corrective feedback to teach an image classifier the difference between vehicles and trees may contribute little to the analyst's expertise. The analyst may also need to spend time deciphering how feedback should be provided (e.g. formatted, curated) to ensure that the ML is learning in an expected manner. The work required to communicate corrective feedback to the machine is one of the coordination costs needed in this new human-machine teaming paradigm. In the old paradigm, monitoring tool performance for accuracy was an important aspect of an analyst's activities. In the human-machine teaming paradigm, a new level of monitoring is now necessary when working with ML tools. Analysts must *monitor* the technology's progress as it learns from feedback. The analyst must assess not only the accuracy of the machine's output, but also how its output has changed in response to the analysts' feedback. These observations in turn guide future corrective feedback. Both communication and monitoring teamwork activities can be challenging in human-human teams. These activities are particularly challenging when interacting with a machine that processes information and formulates patterns in ways that may be unpredictable and alien to humans without effective explanatory features. Existing visual analytic approaches in this scenario involve those where human analysts provide explicit (e.g., using active learning) [5] or implicit labels (e.g., using semantic interaction) [39] to a machine learning model and the latter attempts to proactively learn about the mental model of the analysts. Analysts have agency in training the model, leveraging synchronous communication that is common in techniques that exploit interactive machine learning.

Judge: Finally, a human-machine team can function collaboratively to evaluate the information that was identified and connected as a part of the exploration and conclusions. Implicit in this evaluation is an understanding of the quality of the model itself, working to judge the degree to which additional information can be identified and further injected into the findings to confirm the conclusions that have already been reached. In the old paradigm, the human had sole responsibility for judging whether the evidence and conclusion was supported by the facts, but a very-developed machine learning model can further support the human inference process by highlighting potential bias, conflicting information, or overlooked avenues of exploration. Existing visual analytic approaches in this scenario involve those where human analysts engage in post-hoc evaluation [10] and interpretation of model decisions using visual explanations [18,43]. In these scenarios, analysts do not have agency in the model training process. They are mainly trying to recover the semantic reasoning of the model and judge if, based on the available explanations, the model predictions can be trusted and acted upon.

2.2 Communication Considerations

As described by Wenskovitch and North [40], communication within an interacting human-machine pair should be considered at four

points: (1) the information that the human externalizes, (2) the machine interpretation of that externalization, (3) the machine response to that interpretation, and (4) the human understanding of that response. Though this discussion is described for communication between closed-box actors, the idea generalizes to human and machine roles in a team. Other research argues that effective communication between humans and machine intelligence goes beyond the algorithms themselves, requiring evaluation of datasets and performance metrics [34], as well as working to clearly define interpretability and its necessity in human-machine collaboration [11].

At the Hypotheses Exploration end of the spectrum from Fig. 2, the human takes on the role of an Explorer while the machine takes on the role of a Recommender. In this feedback loop, the human must communicate to the machine the type, quantity, and constraints of the information that it is seeking. Such information could be presented to the machine directly either via a straightforward search box or a more complex dynamic querying interface, or it could be inferred from user actions such as annotations or highlights within documents [38]. The machine response to this information depends on the directness by which it is presented, permitting the machine to either construct a database query based on clearly-presented search constraints or to rank the recommendations that it returns using a learned metric. In the opposing communication direction, the machine is returning its recommended data sets, analytics, and/or labels to the user. Here, the machine should provide a rationale that accompanies the recommendation, communicating the means by which the output was generated. The human understanding of that rationale will directly inform the level of trust that the human has in the information provided by the machine.

A similar communication relationship can be found in the Post-hoc Model Interpretation [18,43] case at the other end of the spectrum. Here, the human acts as a Judge while the machine works as a Supporter of human inference. The support that the machine is able to provide is directly tied to the machine understanding of those human inferences. The human must then clearly externalize their conclusions and evidence, so that the machine is able to provide adequate support. In the other direction, the level of trust that the human has in the provided supporting evidence will be informed by the quality of the machine output and how it is subsequently internalized.

3 CHARACTERIZING AGENCY IN HUMAN-MACHINE TEAMS

Regardless of their role in the human-machine team, the human analysts must relinquish more control of the sensemaking process to the machine in this new paradigm. This loss of control can be a challenging adjustment for the human and lead to performance degradation. The need that arises within a human to maintain supervisory control over a machine teammate is thought to be rooted in feeling a lack of agency, or having some decisions made by the system out of their direct control. Indeed, active research seeks to understand the degree to which users want to automate analytical processes and interact with machine learning [32], in some cases finding no need for complete automation and subsequently a continued role for the human in the loop [35].

Haslam defines agency as one's ability to act independently to make decisions [16]. Berberian found that changes in agency led to impacts in operator performance, as well as in system acceptability and adoption [3]. For example, Baron [1] reported that "the major human concern of pilots in regard to introduction of automation is that, in some circumstances, operations with such aids may leave the critical question, who is in control now, the human or the machine?" Shneiderman & Plaisant agree, noting that "users strongly desire the sense that they are in charge of the system and that the system responds to their actions" [30].

A feeling of a loss of agency can lead to other effects as well, such as a degradation of the quality of monitoring and action. These

effects have particular influence on the Teacher role, who must monitor and provide quality feedback to the machine in order to optimize the contributions provided by the machine, as well as to the Judge working to confirm conclusions reached in collaboration with the machine. Both theory [26] and evidence [25] in the human factors literature suggests the removal of human agency in favor of automation can lead to a negative impact on human performance when the automation fails. Humans struggle to gain situation awareness and provide prompt corrective action in the wake of these failures, and this degradation in performance intensifies as the level of automation increases [25]. Kuhn et al. [19] present an electrophysiological justification for this, showing that feelings of agency are linked to changes in the processing of sensory events. Likewise, Bednark & Franz identified a connection between action and outcome that is dependent upon the attribution of agency in the presence of others [2]. In more extreme circumstances, a feeling of being coerced into taking a step more closely resembles passive movement rather than intentional action [6].

Loss of agency can also interfere with the attribution of responsibility. Haggard [15] argues that the sense of agency holds a great importance in human civilization as a whole because it underpins the concept of responsibility in human societies. This argument has also been validated experimentally. Li et al. [22] present an experiment that studied the subjective sense of responsibility, finding that participants reported a higher sense of responsibility when they were told that a game was uncontrollable. Moretto et al. [24] support this finding, identifying a tight association between action and effect that occurs even in scenarios when actions had unpredictable consequences. Caspar et al. [6] found that induced disbelief in free will impact the sense of agency over the consequences of one's actions.

Thankfully, solutions exist to mitigate this perceived loss of agency. Priming the user, or providing some predictive information of what will happen next, has shown benefits. Wegner et al. [36] conducted several sensory experiments, finding that priming a movement with instructions led to a greater feeling of agency than when the instructions followed the movement. Moore et al. found similar benefits when studying voluntary vs. involuntary movements [23]. Finally, Sato [29] studied whether the sense of agency depended primarily on a connection between preview information and sensory feedback or between prediction and sensory feedback. The answer is "both." These priming actions and preview information are particularly relevant to recommendations systems that support users with the introduction of new data, and particularly influence the Explorer and Investigator roles during the data foraging and synthesis processes.

Further, Peters et al. [27] proposed that examining the interactions between human and machine to identify appropriate interruption timings can be beneficial. Interrupting the human without disrupting the teaming interactions with an intelligent interruption system that monitors the actions of the human teammate enhances the human's feelings of control and agency. Wegner & Wheatley [37] argue simply for the identification of whether or not an action is perceived as "willed," listing these three principles: i) The thought precedes the action at a proper interval (the *Priority Principle*), ii) the thought is compatible with the reaction (the *Consistency Principle*), and iii) the thought is the only apparent cause of the action (the *Exclusivity Principle*).

Fallon et al. emphasize the importance of having a machine teammate with the ability to establish and maintain common ground in human-machine team [13]. Although autonomy is one characteristic necessary for a machine to function as a teammate, autonomy alone is not sufficient. Without common ground the human may lose trust and reduce their reliance on the machine. These researchers propose that a machine that strives to both inform and stay informed may help mitigate some of the challenges associated with the machine assuming greater agency in the human-machine partnership.

4 CONCLUSION AND FUTURE WORK

In this paper, we have outlined a vision for human-machine teaming grounded in the existing literature in cognitive psychology and emerging technological advances in the machine learning, AI, and visual analytics communities. We analyzed why human agency is important, discussed four different scenarios for teaming and described the communication processes in such teaming scenarios. We believe that this futuristic vision will help guide researchers interested in human-AI interaction to develop solutions that are not only address data-driven challenges but also the challenges of intelligible communication among humans and machines.

REFERENCES

- [1] S. Baron. Pilot control. In *Human Factors in Aviation*, pp. 347–385. Elsevier, 1988.
- [2] J. G. Bednark and E. A. Franz. Agency attribution: Event-related potentials and outcome monitoring. *Experimental Brain Research*, 232(4):1117–1126, 2014.
- [3] B. Berberian. Man-machine teaming: A problem of agency. *IFAC-PapersOnLine*, 51(34):118–123, 2019.
- [4] J. Bernard, M. Zeppelzauer, M. Lehmann, M. Müller, and M. Sedlmair. Towards user-centered active learning algorithms. *Computer Graphics Forum*, 37(3):121–132, 2018.
- [5] J. Bernard, M. Zeppelzauer, M. Sedlmair, and W. Aigner. Vial: a unified process for visual interactive labeling. *The Visual Computer*, 34(9):1189–1207, 2018.
- [6] E. A. Caspar, J. F. Christensen, A. Cleeremans, and P. Haggard. Coercion changes the sense of agency in the human brain. *Current biology*, 26(5):585–592, 2016.
- [7] A. Chavaillaz, D. Wastell, and J. Sauer. System reliability, performance and trust in adaptable automation. *Applied Ergonomics*, 52:333–342, Jan. 2016. doi: 10.1016/j.apergo.2015.07.012
- [8] K. Cook, N. Cramer, D. Israel, M. Wolverson, J. Bruce, R. Burtner, and A. Endert. Mixed-initiative visual analytics using task-driven recommendations. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 9–16. IEEE, 2015.
- [9] A. Dasgupta, J.-Y. Lee, R. Wilson, R. A. LaFrance, N. Cramer, K. Cook, and S. Payne. Familiarity vs trust: A comparative study of domain scientists' trust in visual analytics and conventional analysis methods. *IEEE transactions on visualization and computer graphics*, 23(1):271–280, 2016.
- [10] A. Dasgupta, H. Wang, N. O'Brien, and S. Burrows. Separating the wheat from the chaff: Comparative visual cues for transparent diagnostics of competing models. *IEEE transactions on visualization and computer graphics*, 26(1):1043–1053, 2020.
- [11] F. Doshi-Velez and B. Kim. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv e-prints*, p. arXiv:1702.08608, Feb. 2017.
- [12] A. Endert, P. Fiaux, and C. North. Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 473–482, 2012.
- [13] C. K. Fallon, L. M. Blaha, K. Cook, and T. Billow. Common ground and autonomy: Two critical dimensions of a machine teammate. In *International Conference on Applied Human Factors and Ergonomics*, pp. 14–26. Springer, 2019.
- [14] C. Felix, A. Dasgupta, and E. Bertini. The exploratory labeling assistant: Mixed-initiative label curation with large document collections. In *31st Annual ACM Symposium on User Interface Software and Technology, UIST 2018*, pp. 153–164. Association for Computing Machinery, Inc, 2018.
- [15] P. Haggard. Sense of agency in the human brain. *Nature Reviews Neuroscience*, 18(4):196, 2017.
- [16] N. Haslam. Dehumanization: An integrative review. *Personality and Social Psychology Review*, 10(3):252–264, 2006.
- [17] J. Hullman and A. Gelman. Designing for interactive exploratory data analysis requires theories of graphical inference. 2021.
- [18] J. Krause, A. Dasgupta, J. Swartz, Y. Aphinyanaphongs, and E. Bertini. A workflow for visual diagnostics of binary classifiers using instance-level explanations. *stat*, 1050:21, 2017.

- [19] S. Kühn, I. Nenchev, P. Haggard, M. Brass, J. Gallinat, and M. Voss. Whodunnit? electrophysiological correlates of agency judgements. *PloS One*, 6(12):e28657, 2011.
- [20] D. J.-L. Lee, J. Lee, T. Siddiqui, J. Kim, K. Karahalios, and A. Parameswaran. You can't always sketch what you want: Understanding sensemaking in visual query systems. *IEEE transactions on visualization and computer graphics*, 26(1):1267–1277, 2019.
- [21] D. J.-L. Lee, A. Quamar, E. Kandogan, and F. Özcan. Boomerang: Proactive insight-based recommendations for guiding conversational data analysis. In *Proceedings of the 2021 International Conference on Management of Data*, pp. 2750–2754, 2021.
- [22] P. Li, C. Han, Y. Lei, C. B. Holroyd, and H. Li. Responsibility modulates neural mechanisms of outcome processing: An erp study. *Psychophysiology*, 48(8):1129–1133, 2011.
- [23] J. W. Moore, D. M. Wegner, and P. Haggard. Modulating the sense of agency with external cues. *Consciousness and Cognition*, 18(4):1056–1064, 2009.
- [24] G. Moretto, E. Walsh, and P. Haggard. Experience of agency and sense of responsibility. *Consciousness and Cognition*, 20(4):1847–1854, 2011.
- [25] L. Onnasch, C. D. Wickens, H. Li, and D. Manzey. Human performance consequences of stages and levels of automation: An integrated meta-analysis. *Human Factors*, 56(3):476–488, 2014.
- [26] R. Parasuraman, T. B. Sheridan, and C. D. Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 30(3):286–297, 2000.
- [27] N. S. Peters, G. C. Bradley, and T. Marshall-Bradley. Task boundary inference via topic modeling to predict interruption timings for human-machine teaming. In *International Conference on Intelligent Human Systems Integration*, pp. 783–788. Springer, 2019.
- [28] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4. McLean, VA, USA, 2005.
- [29] A. Sato. Both motor prediction and conceptual congruency between preview and action-effect contribute to explicit judgment of agency. *Cognition*, 110(1):74–83, 2009.
- [30] B. Shneiderman, C. Plaisant, M. S. Cohen, S. Jacobs, N. Elmqvist, and N. Diakopoulos. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Pearson, 2016.
- [31] A. Srinivasan, N. Nyapathy, B. Lee, S. M. Drucker, and J. Stasko. Collecting and characterizing natural language utterances for specifying data visualizations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–10, 2021.
- [32] H. Suresh, S. R. Gomez, K. K. Nam, and A. Satyanarayan. Beyond expertise and roles: A framework to characterize the stakeholders of interpretable machine learning and their needs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2021.
- [33] J. J. Thomas and K. A. Cook. A visual analytics agenda. *IEEE computer graphics and applications*, 26(1):10–13, 2006.
- [34] J. W. Vaughan and H. Wallach. A human-centered agenda for intelligible machine learning. *Machines We Trust: Getting Along with Artificial Intelligence*, 2020.
- [35] D. Wang, Q. V. Liao, Y. Zhang, U. Khurana, H. Samulowitz, S. Park, M. Muller, and L. Amini. How Much Automation Does a Data Scientist Want? *arXiv e-prints*, p. arXiv:2101.03970, Jan. 2021.
- [36] D. M. Wegner, B. Sparrow, and L. Winerman. Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology*, 86(6):838, 2004.
- [37] D. M. Wegner and T. Wheatley. Apparent mental causation: Sources of the experience of will. *American Psychologist*, 54(7):480, 1999.
- [38] J. Wenskovitch, L. Bradel, M. Dowling, L. House, and C. North. The effect of semantic interaction on foraging in text analysis. In *2018 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 13–24, 10 2018. doi: 10.1109/VAST.2018.8802424
- [39] J. Wenskovitch, I. Crandell, N. Ramakrishnan, L. House, S. Leman, and C. North. Towards a systematic combination of dimension reduction and clustering in visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):131–141, 1 2018. doi: 10.1109/TVCG.2017.2745258
- [40] J. Wenskovitch and C. North. Interactive artificial intelligence: Designing for the ‘two black boxes’ problem. *Computer*, 53(8):29–39, 2020. doi: 10.1109/MC.2020.2996416
- [41] C. D. Wickens, B. A. Clegg, A. Z. Vieane, and A. L. Sebok. Complacency and Automation Bias in the Use of Imperfect Automation. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 57(5):728–739, Aug. 2015. doi: 10.1177/0018720815581940
- [42] B. Yu and C. T. Silva. Flowsense: A natural language interface for visual data exploration within a dataflow system. *IEEE transactions on visualization and computer graphics*, 26(1):1–11, 2019.
- [43] J. Zhang, Y. Wang, P. Molino, L. Li, and D. S. Ebert. Manifold: A model-agnostic framework for interpretation and diagnosis of machine learning models. *IEEE transactions on visualization and computer graphics*, 25(1):364–373, 2018.
- [44] J. Zhao, M. Fan, and M. Feng. Chartseer: Interactive steering exploratory visual analysis with machine intelligence. *IEEE Transactions on Visualization and Computer Graphics*, 2020.