Balancing Trust and Over-Reliance in Visual Analytics: The AI-in-the-Loop Dilemma

Alvitta Ottley

Visualization Interaction and Behavior Exploration

http://visualdata.wustl.edu/

Increasing integration of AI in VA

Extreme Behaviors: Underutilization vs Overreliance

Appropriate reliance

Underutilization

Overreliance

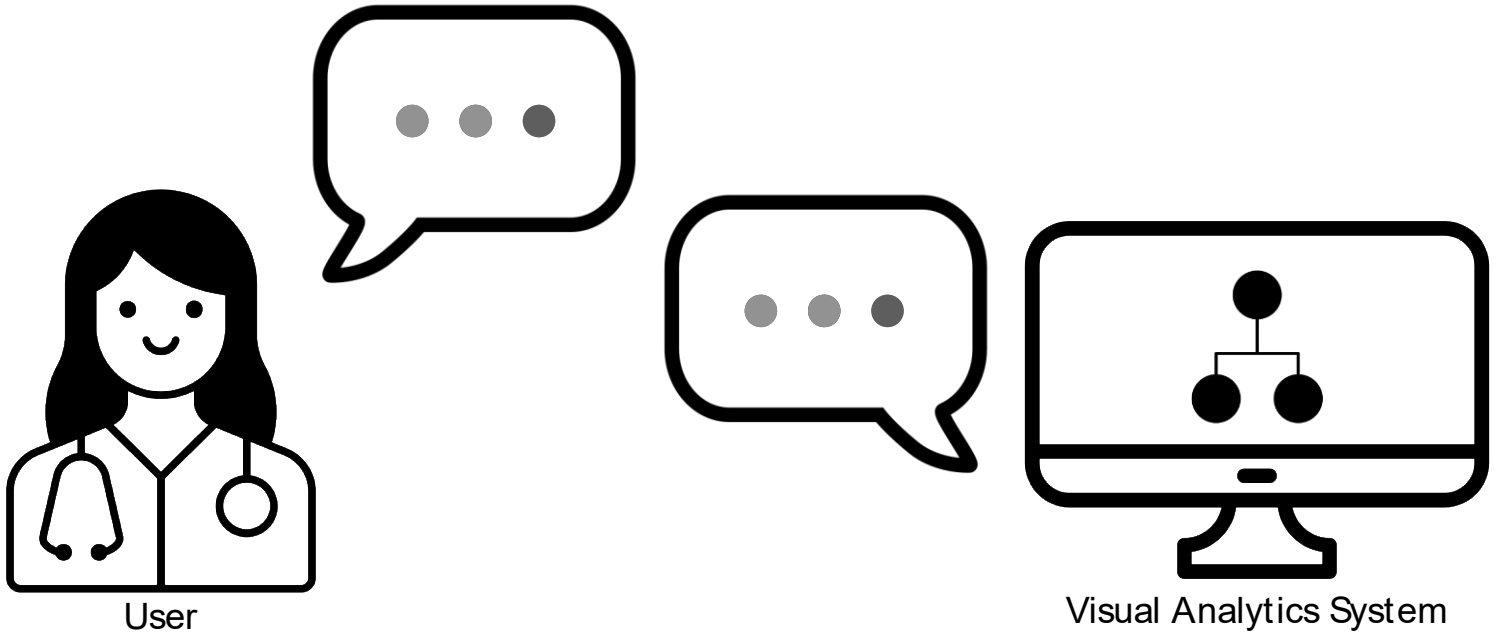

Can AI help solve a VAST Challenge?

Guided Data Discovery in Interactive Visualizations via Active Search

Shayan Monadjemi*
Washington University in St. Louis

Sunwoo Ha†
Washington University in St. Louis

Quan Nguyen
Washington University in St. Louis

Henry Chai
Carnegie Mellon University

Roman Garnett
Washington University in St. Louis

Alvitta Ottley‡
Washington University in St. Louis

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

Intelligence Analysis Scenario: Vastopolis

- Published by the VAST (Visual Analytics Science and Technology) community in 2011
- Involves a fictitious terrorist attack where a truck accident over a major river contaminates the water and air with harmful chemicals
- Contains 1,023,077 social media posts from various parts of town during a 21-day period (04/30/2011 - 05/20/2011)

Machine Learning for Time Series

- Samples arrive sequentially
- Sample size is unknown and varies
- Data are not available during training
- Waiting for until time T to accumulate a batch may not be feasible
  - Eliminates recurrent Neural Network (rNN), i.e., long short-term memory network
    - Notoriously difficult to train
    - Require temporal relationships of the past and future to be similar

Ottley, A., Garnett, R., & Wan, R. (2019, June). Follow the clicks: Learning and Anticipating Mouse Interactions During Exploratory Data Analysis. In Computer Graphics Forum (Vol. 38, No. 3, pp. 41-52).

Providing Guidance with *k*-NN and Active Search

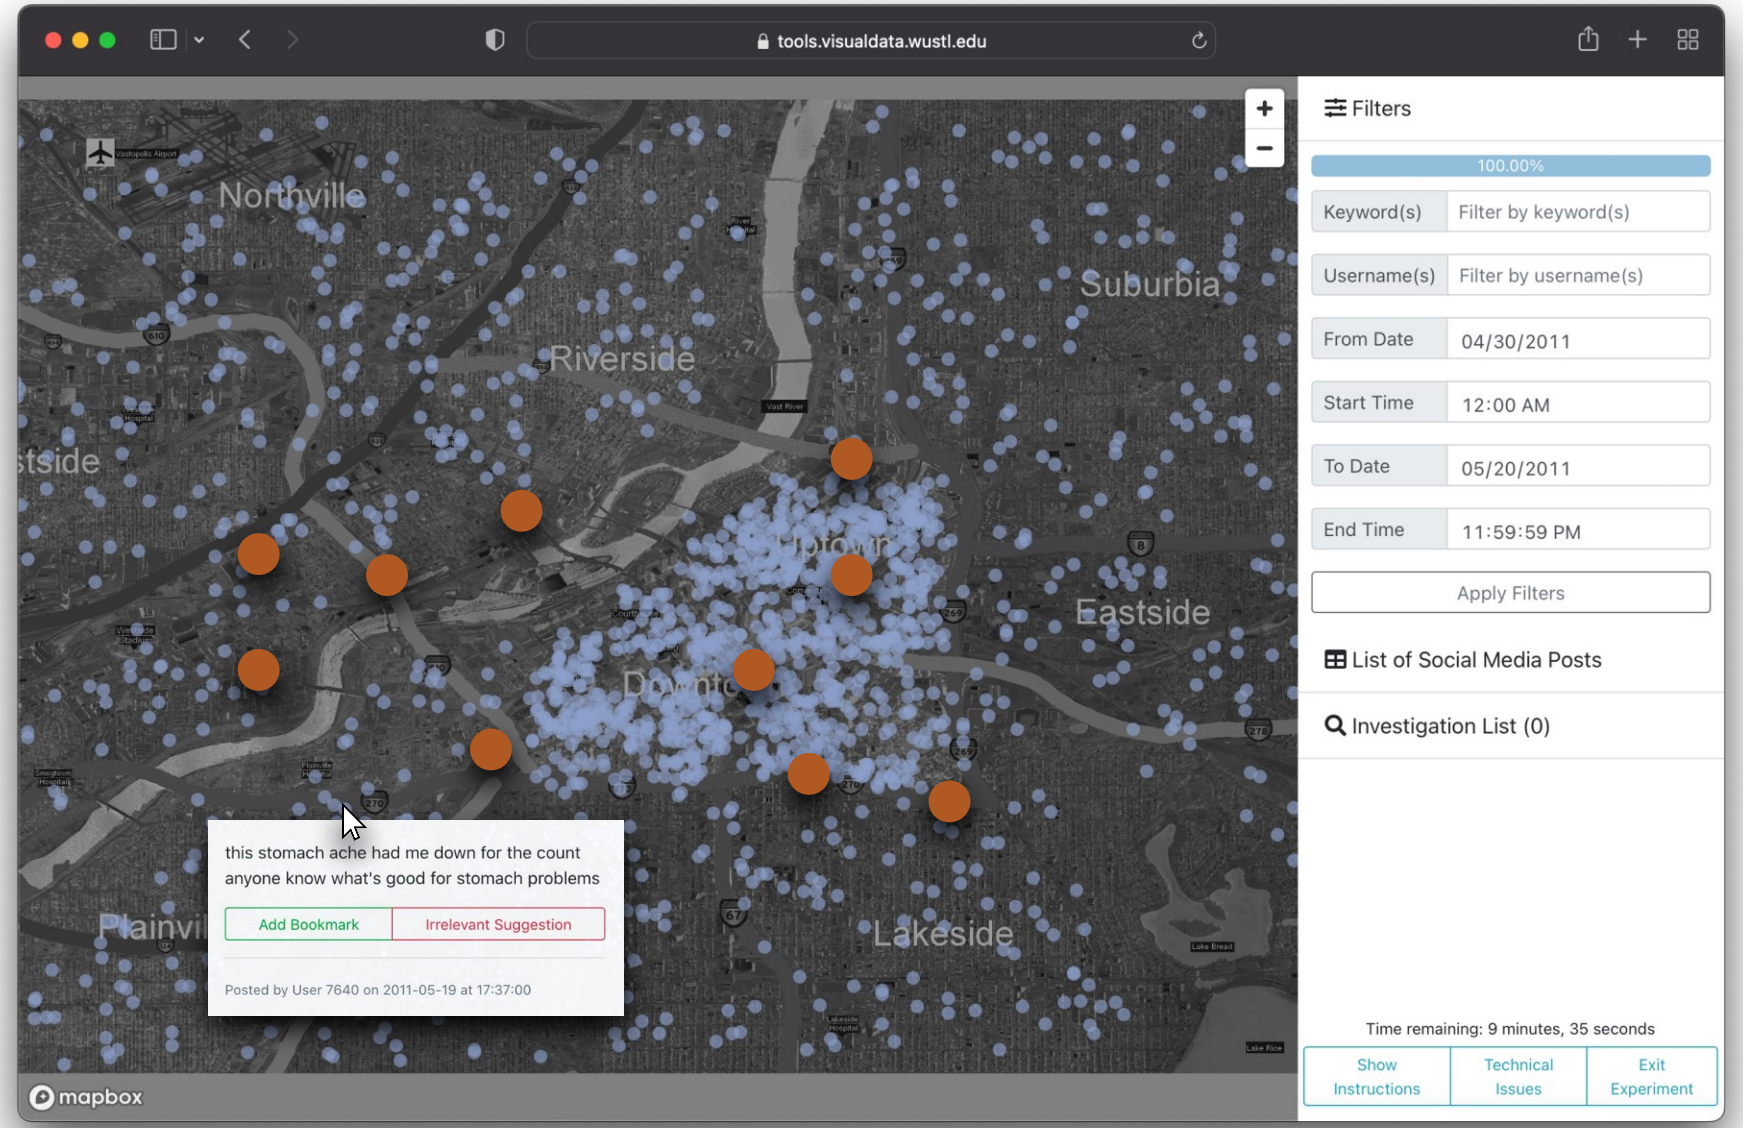

k-NN

Predicts relevant data points in light of past interactions

Active Search

Decides which points to suggest given the *k*-NN model's belief

Same interface and task

You are assisting contact tracers and public health officials in their efforts to identify the source of a new outbreak.

“Tag as many posts mentioning residents having illness-related symptoms as possible within 10 minutes.“

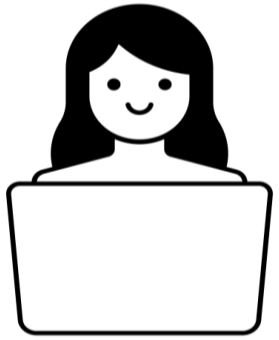

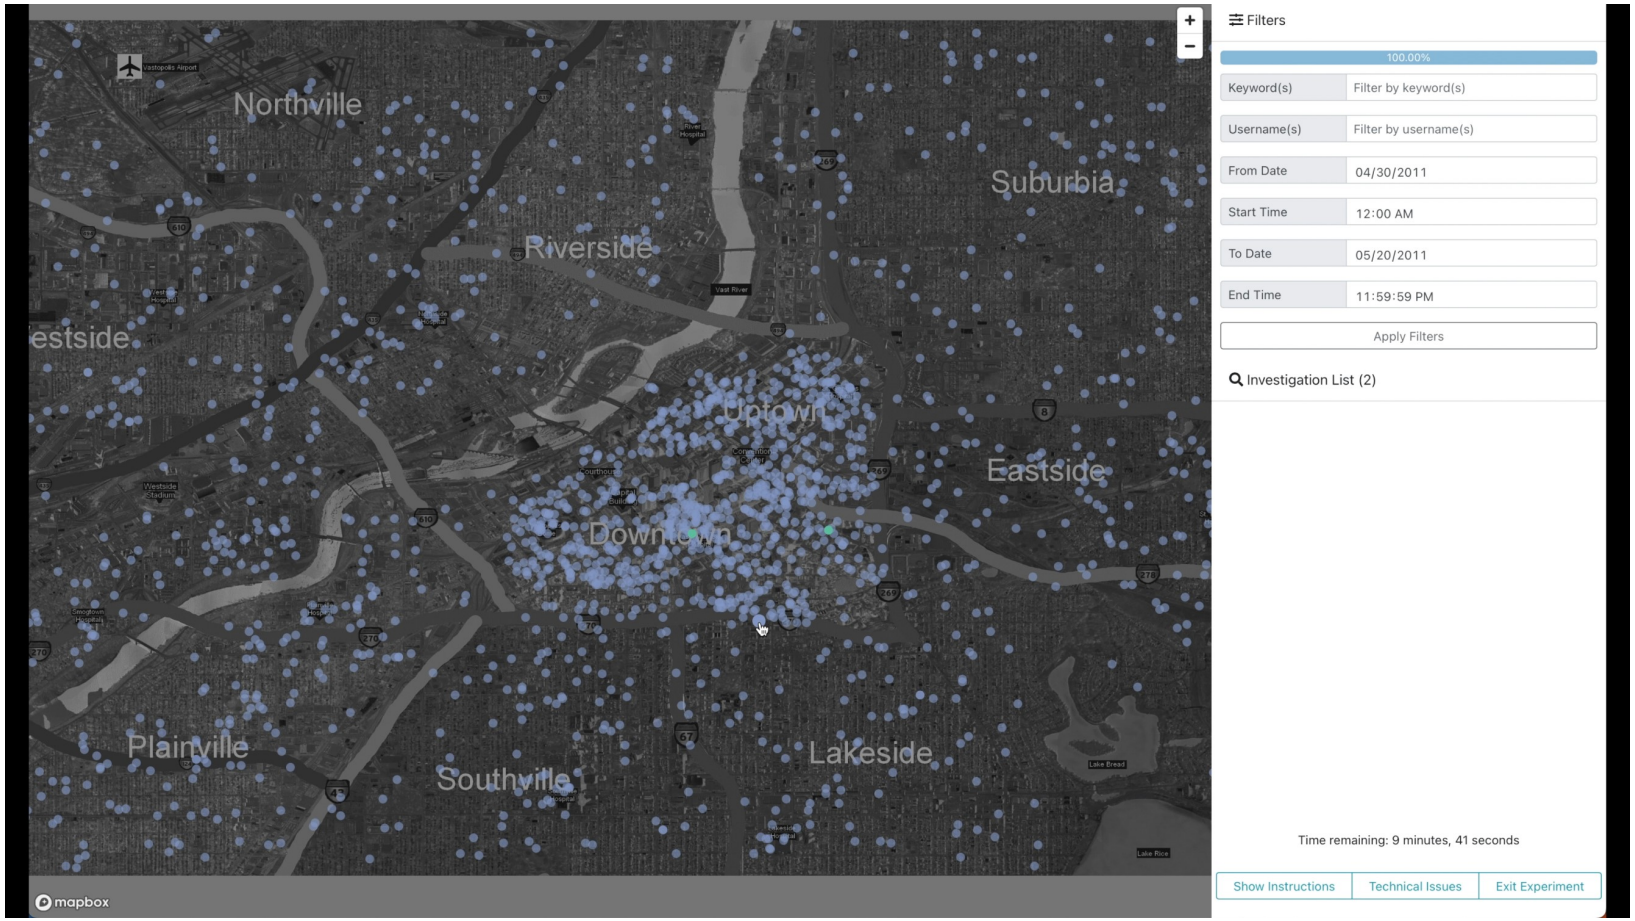

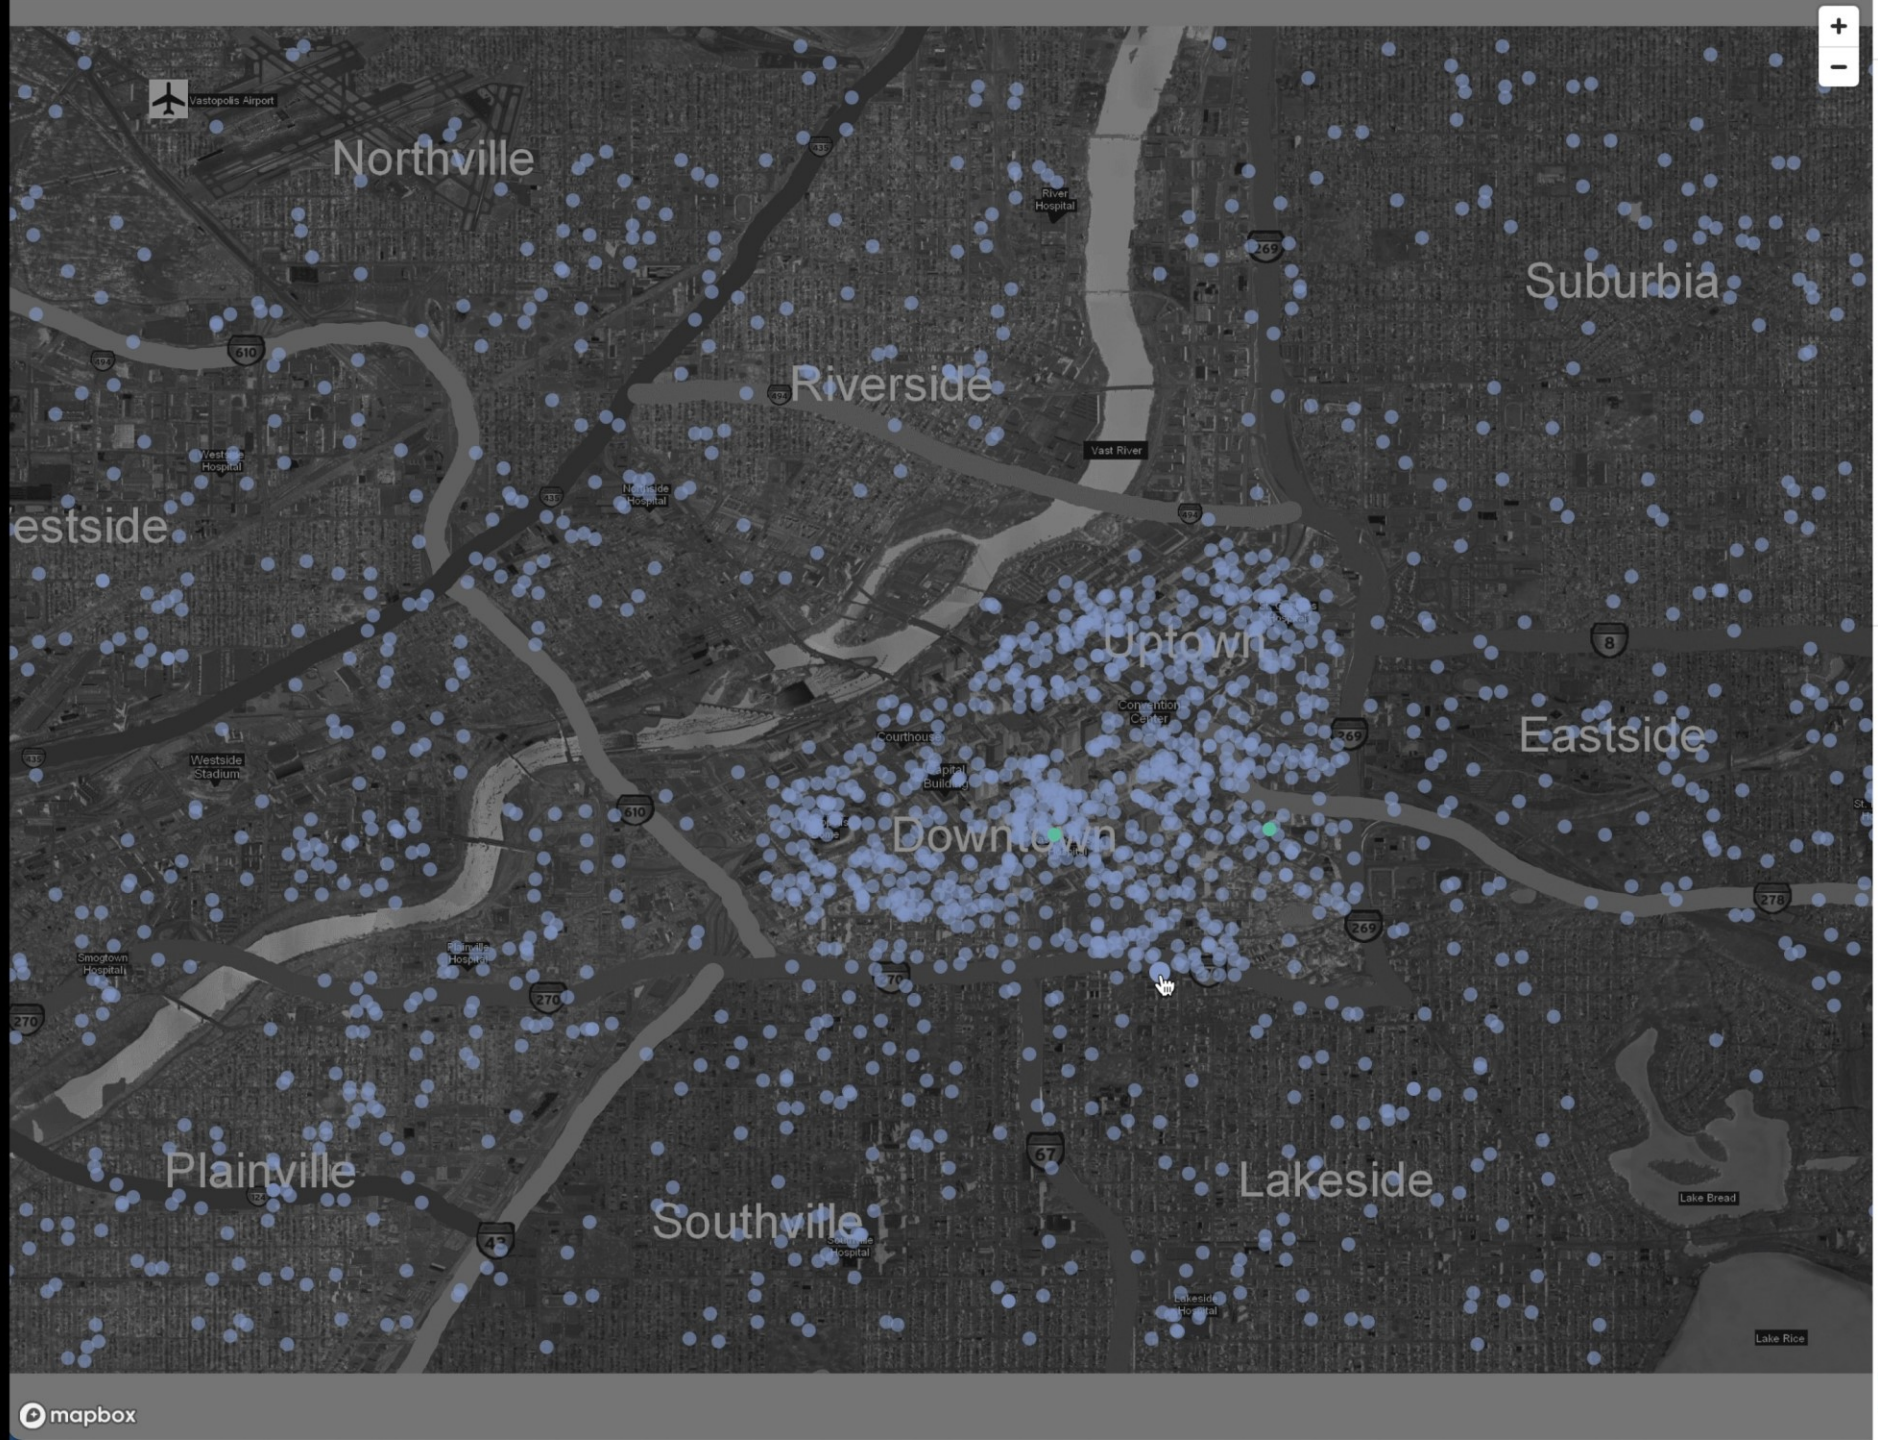

Filters

100.00%

Keyword(s)	Filter by keyword(s)
Username(s)	Filter by username(s)
From Date	04/30/2011
Start Time	12:00 AM
To Date	05/20/2011
End Time	11:59:59 PM

Apply Filters

Investigation List (2)

Time remaining: 9 minutes, 41 seconds

Show Instructions | Technical Issues | Exit Experiment

Crowd-sourced User Study

Control Group

Active Search Group

10 Minutes

Crowd-sourced User Study

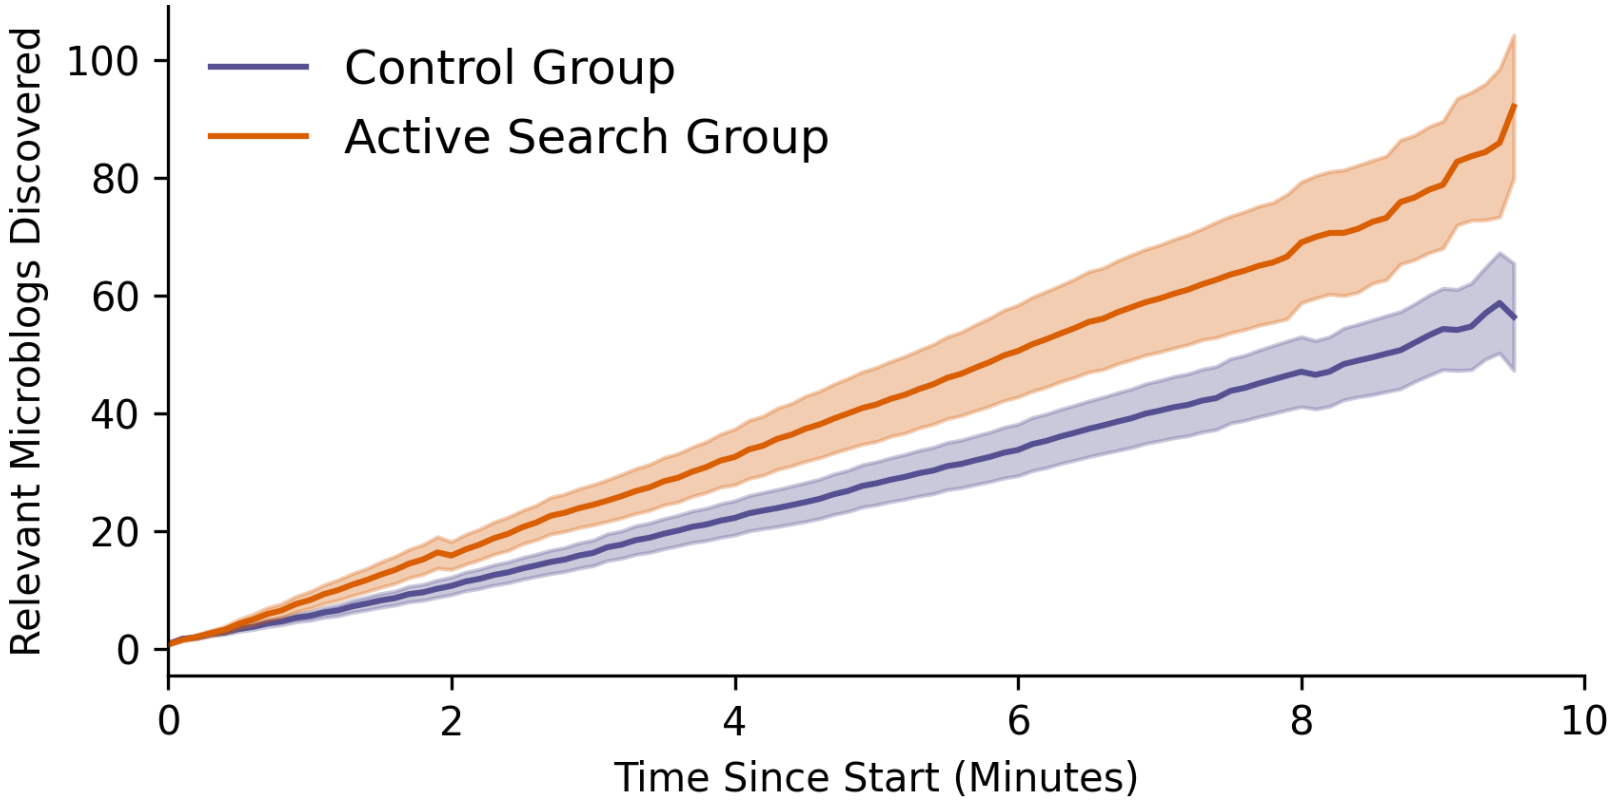

Crowd-sourced User Study

	Control Group	Active Search Group	
Hovers per Minute	16.7 ± 1.19	14.3 ± 1.23	$p = 0.0112$
Relevant Hovers per Minute	6.7 ± 0.68	9.2 ± 1.12	$p = 0.0001$
Hover Purity	0.39 ± 0.02	0.63 ± 0.05	$p < 0.0001$

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

Crowd-sourced User Study

	Control Group	Active Search Group	
Hovers per Minute	16.7 ± 1.19	14.3 ± 1.23	$p = 0.0112$
Relevant Hovers per Minute	6.7 ± 0.68	9.2 ± 1.12	$p = 0.0001$
Hover Purity	0.39 ± 0.02	0.63 ± 0.05	$p < 0.0001$

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

Crowd-sourced User Study

	Control Group	Active Search Group	
Hovers per Minute	16.7 ± 1.19	14.3 ± 1.23	$p = 0.0112$
Relevant Hovers per Minute	6.7 ± 0.68	9.2 ± 1.12	$p = 0.0001$
Hover Purity	0.39 ± 0.02	0.63 ± 0.05	$p < 0.0001$

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

Crowd-sourced User Study

	Control Group	Active Search Group	
Hovers per Minute	16.7 ± 1.19	14.3 ± 1.23	$p = 0.0112$
Relevant Hovers per Minute	6.7 ± 0.68	9.2 ± 1.12	$p = 0.0001$
Hover Purity	0.39 ± 0.02	0.63 ± 0.05	$p < 0.0001$

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

A large percentage of people ignored the recommendations

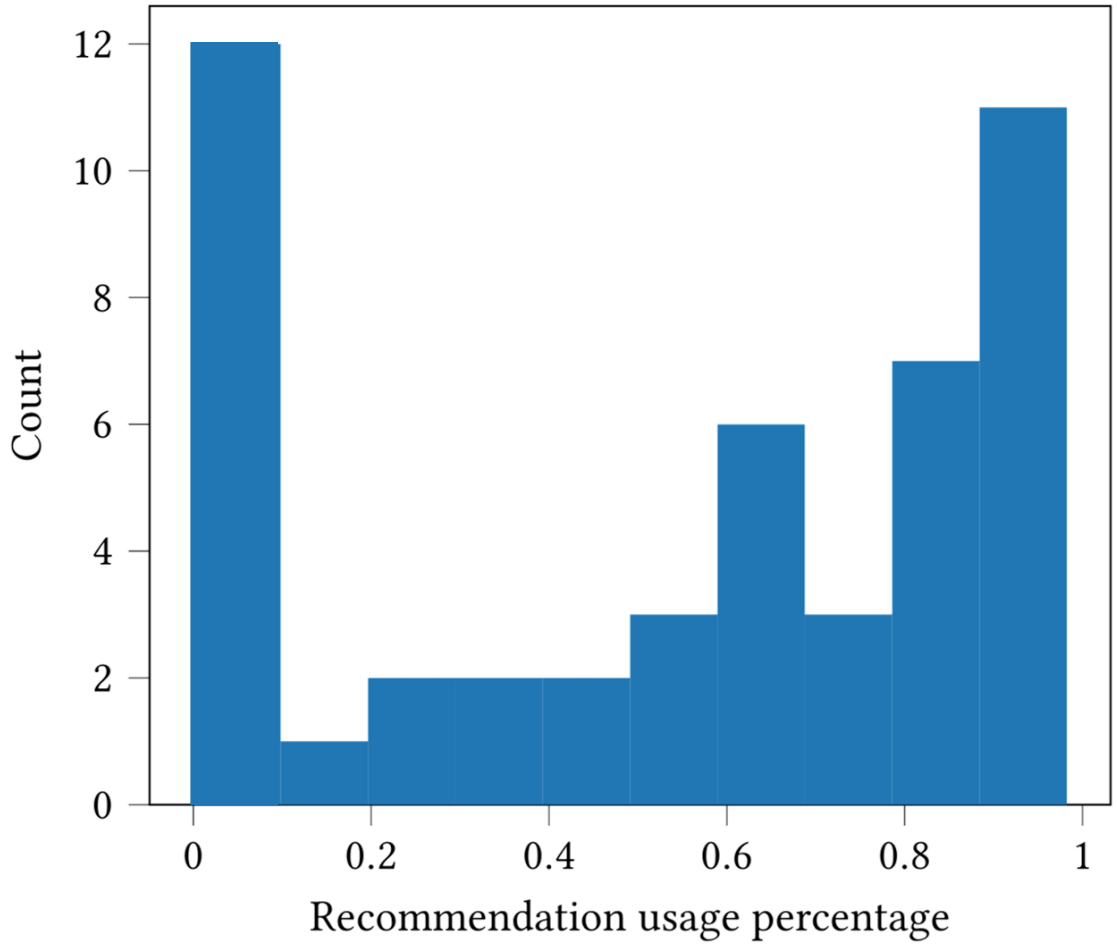

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

A large percentage of people ignored the recommendations

Blind trust can be equally problematic

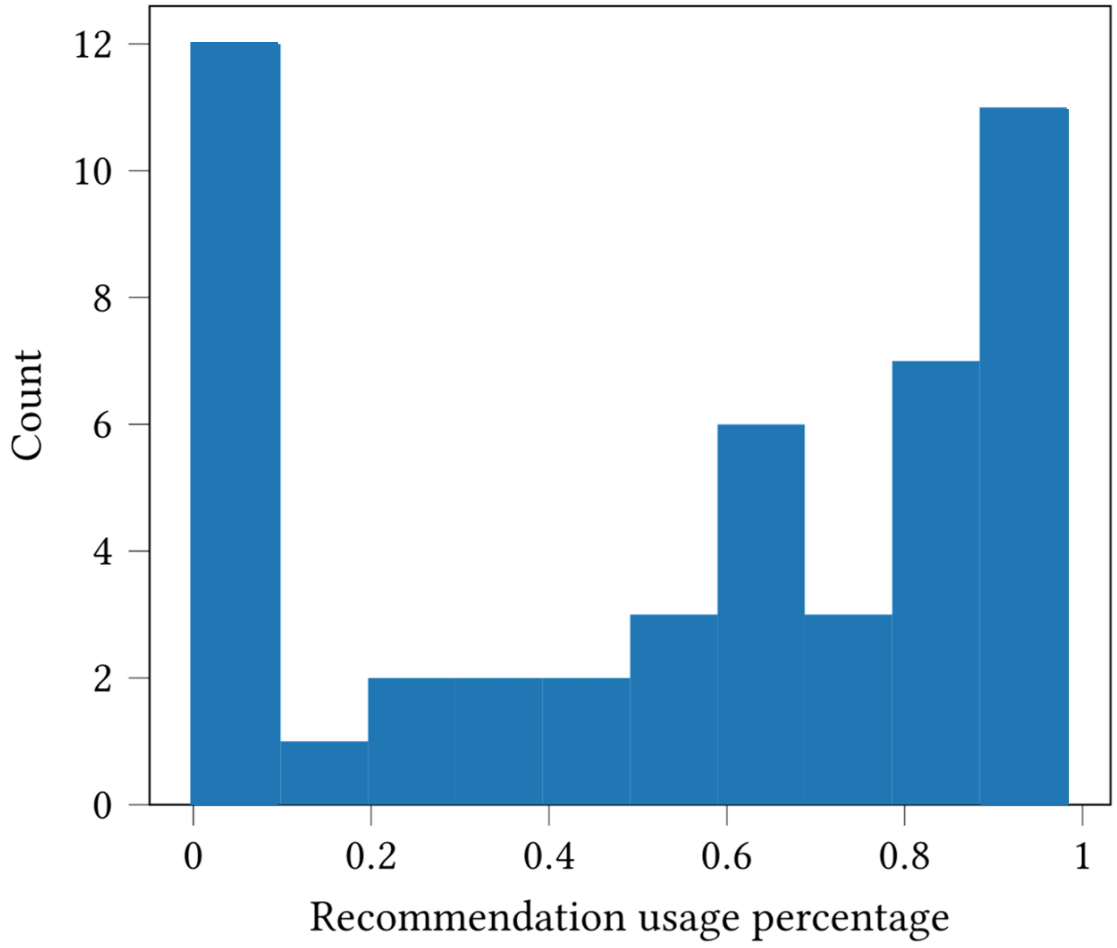

Monadjemi, S., Ha, S., Nguyen, Q., Chai, H., Garnett, R., & Ottley, A. (2022, October). Guided data discovery in interactive visualizations via active search. In 2022 IEEE Visualization and Visual Analytics (VIS) (pp. 70-74). IEEE.

Extreme Behaviors: Underutilization vs Overreliance

Appropriate reliance

Underutilization

Overreliance

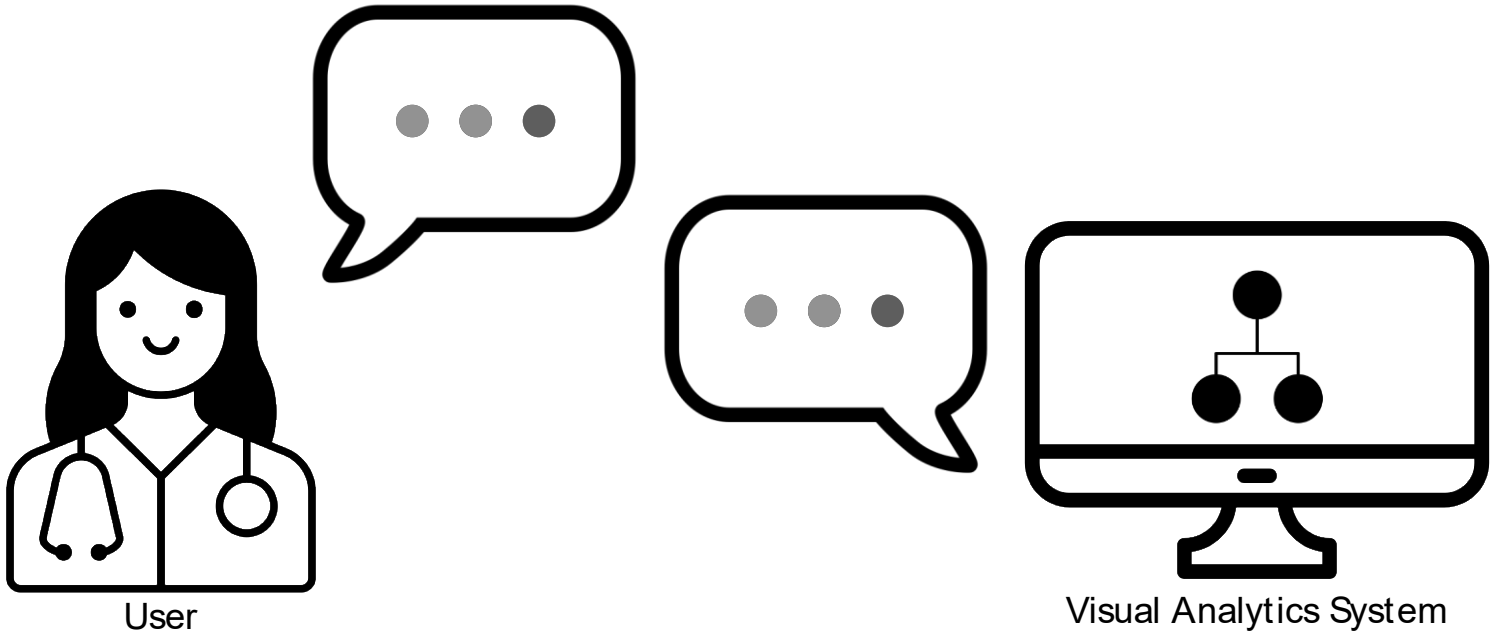

Research Questions

1. How do different levels of transparency provided by VA systems impact users' trust and acceptance of suggestions?

2. How does the decision to utilize or not utilize suggestions affect users' data exploration patterns and decisions?

3. How might task difficulty affect observed behaviors?

Same interface and task

You are assisting contact tracers and public health officials in their efforts to identify the source of a new outbreak.

“Tag as many posts mentioning residents having illness-related symptoms as possible within 10 minutes.“

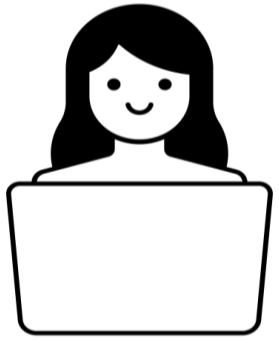

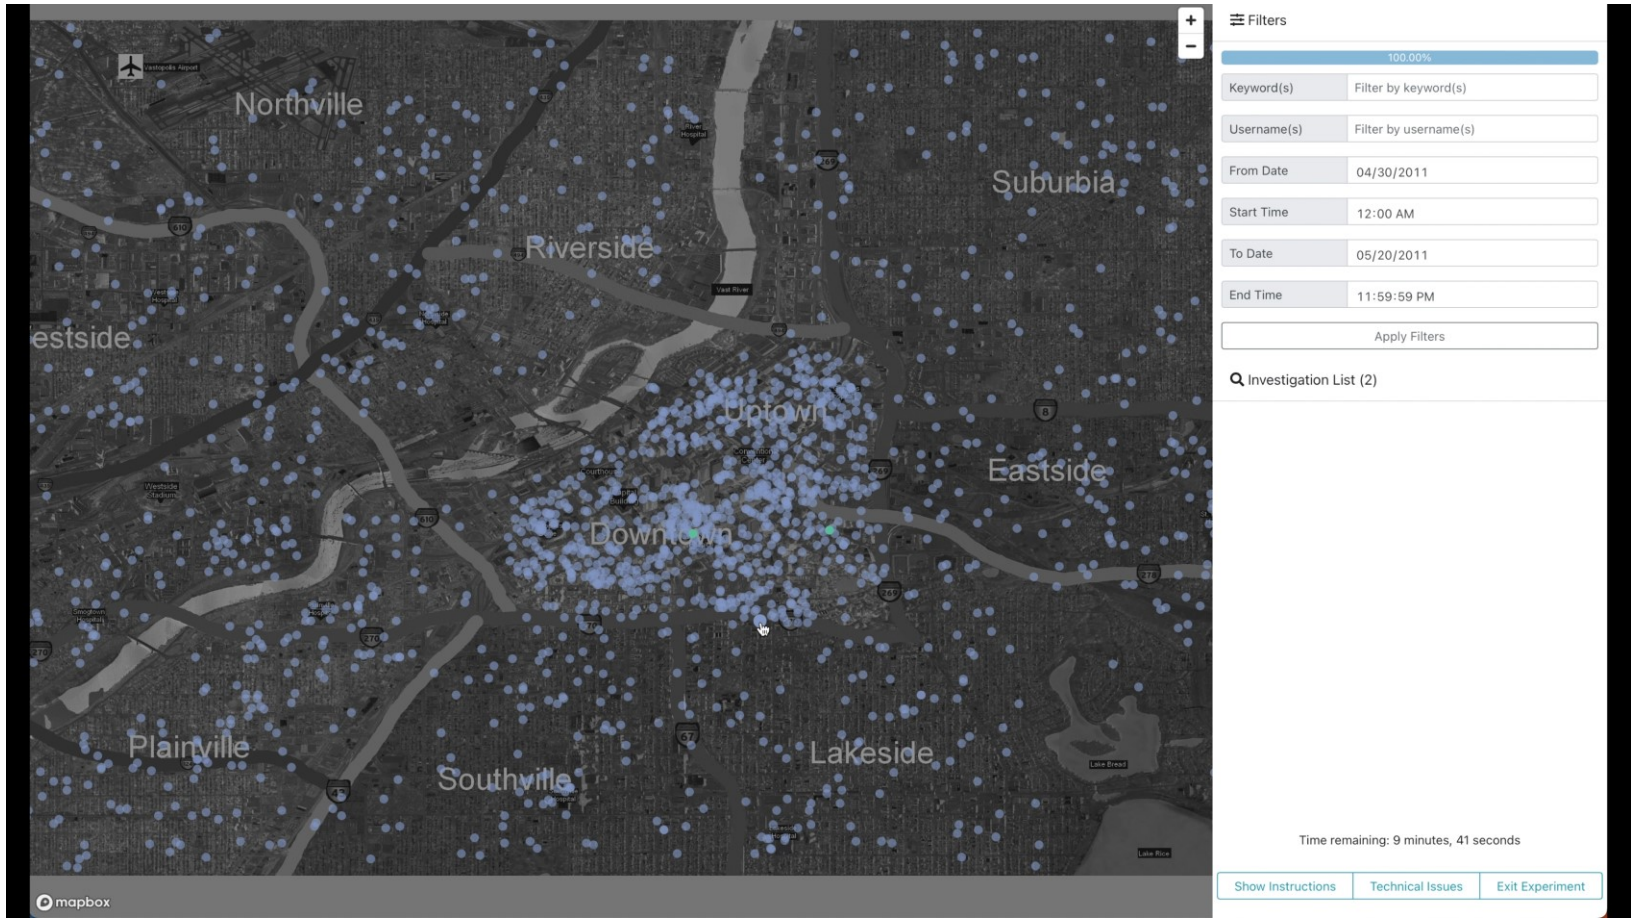

Experiment Design: Conditions

- 2 (task difficulty) x 4 (transparency level) between-subject design
  - Task difficulty
    - Both datasets were a subset of 2000 points

Easy (36%)

Hard (9%)

Experiment Design: Transparency Variants

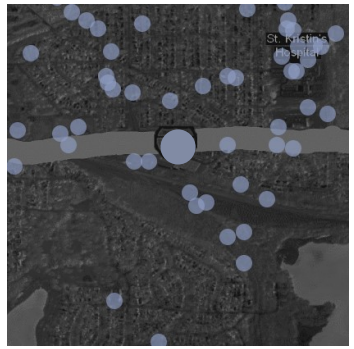
Control

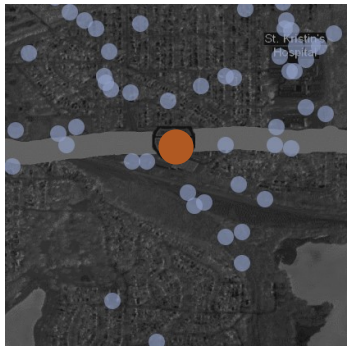
No Explanations

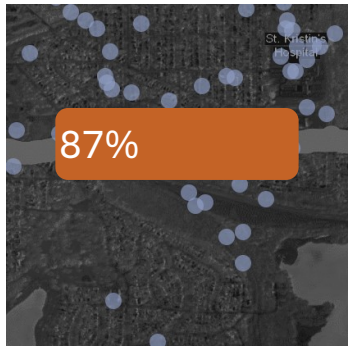

Confidence

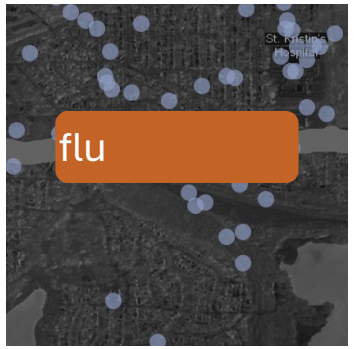

Keyword

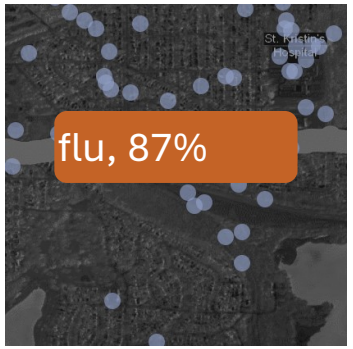

Keyword + Confidence

Ha, S., Monadjemi, S., & Ottley, A. (2024, June). Guided By AI: Navigating Trust, Bias, and Data Exploration in AI-Guided Visual Analytics. In Computer Graphics Forum (Vol. 43, No. 3, p. e15108).

Is it more likely to utilize suggestions when completing a more difficult task?

Is it more likely to utilize suggestions when completing a more difficult task?

(U = 10178.5, p = .0394, η^2 = .014)

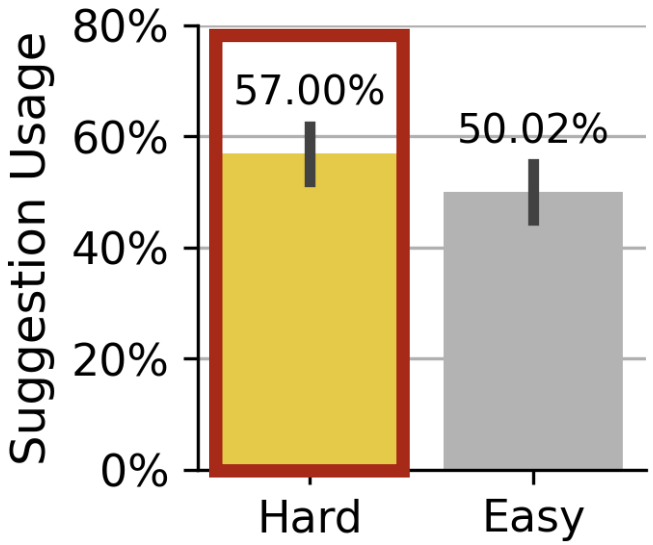

Participants in the hard task utilized suggestions more

(U = 16790.5, p < .001, η^2 = .135)

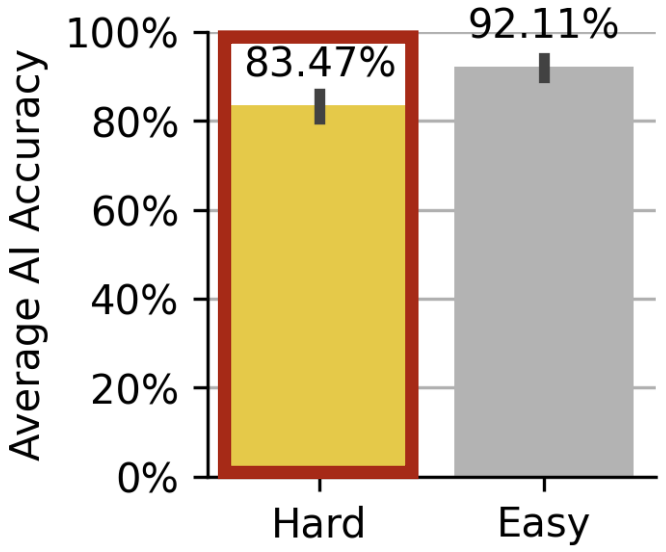

The AI provided less accurate suggestions than for participants in the hard task

Is it more probable to use suggestions when there is more transparency?

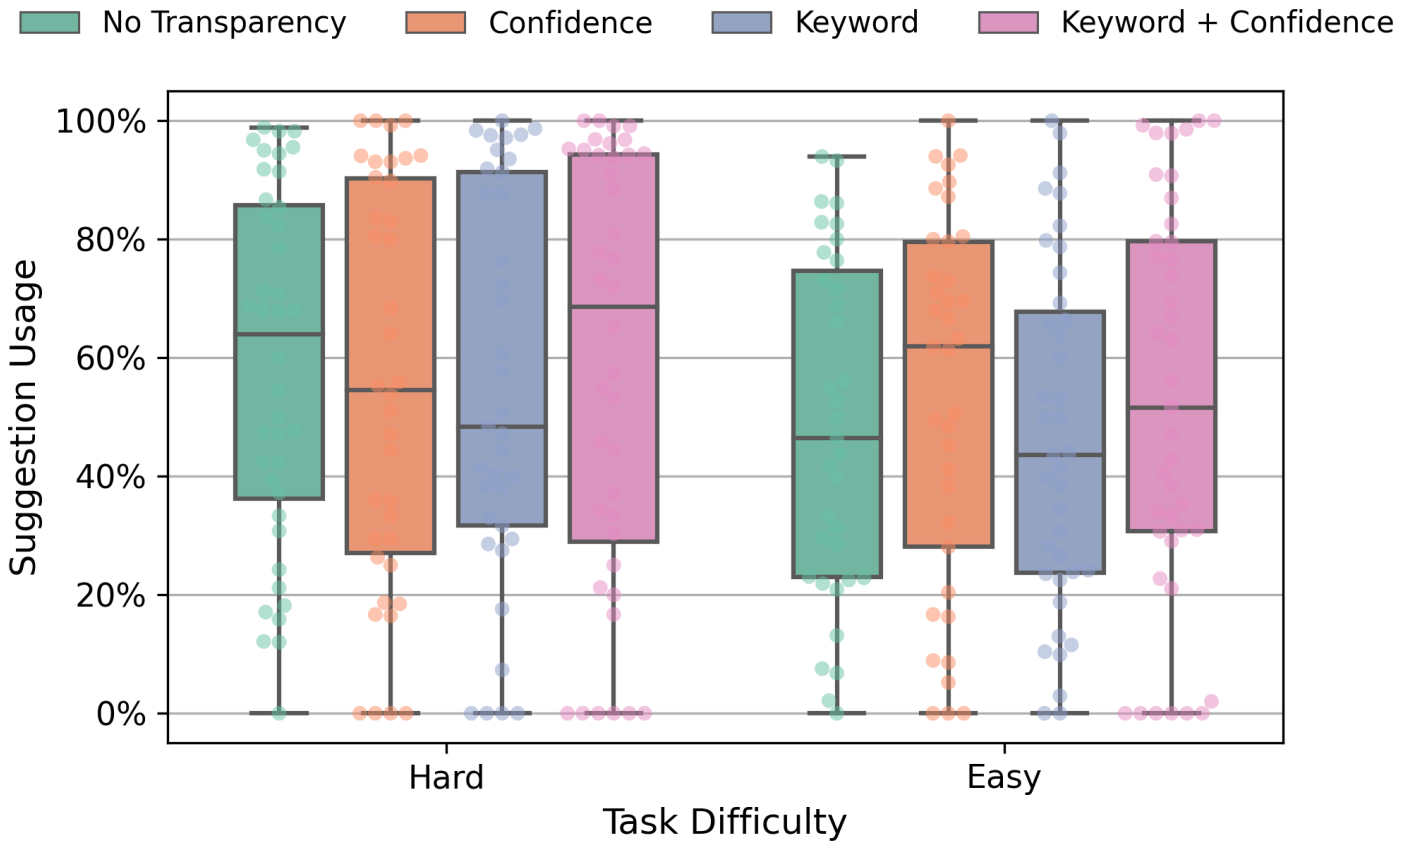

No difference in suggestion usage based on amount of transparency relayed to the participant

Is it more probable to use suggestions when there is more transparency? (cont.)

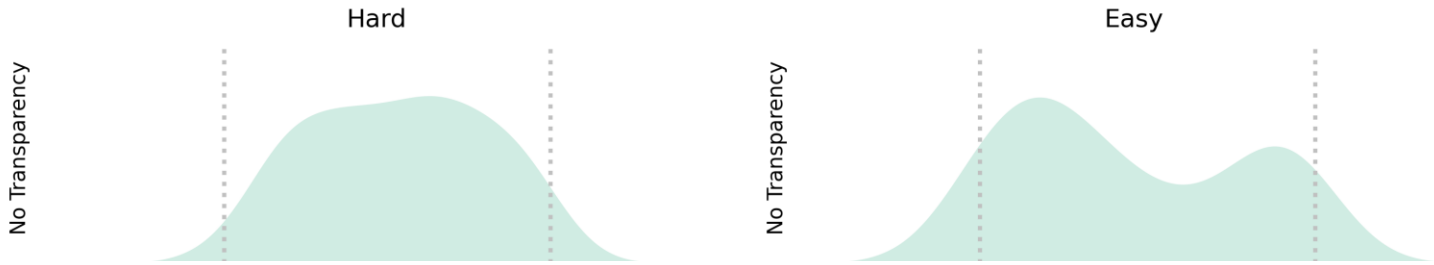

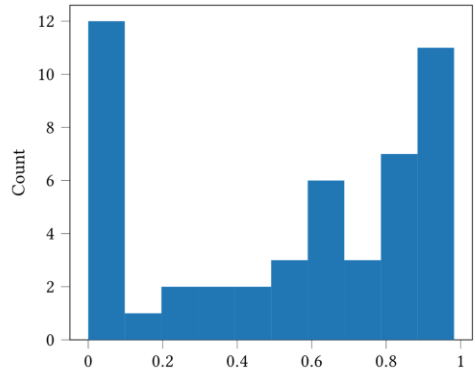

Is it more probable to use suggestions when there is more transparency? (cont.)

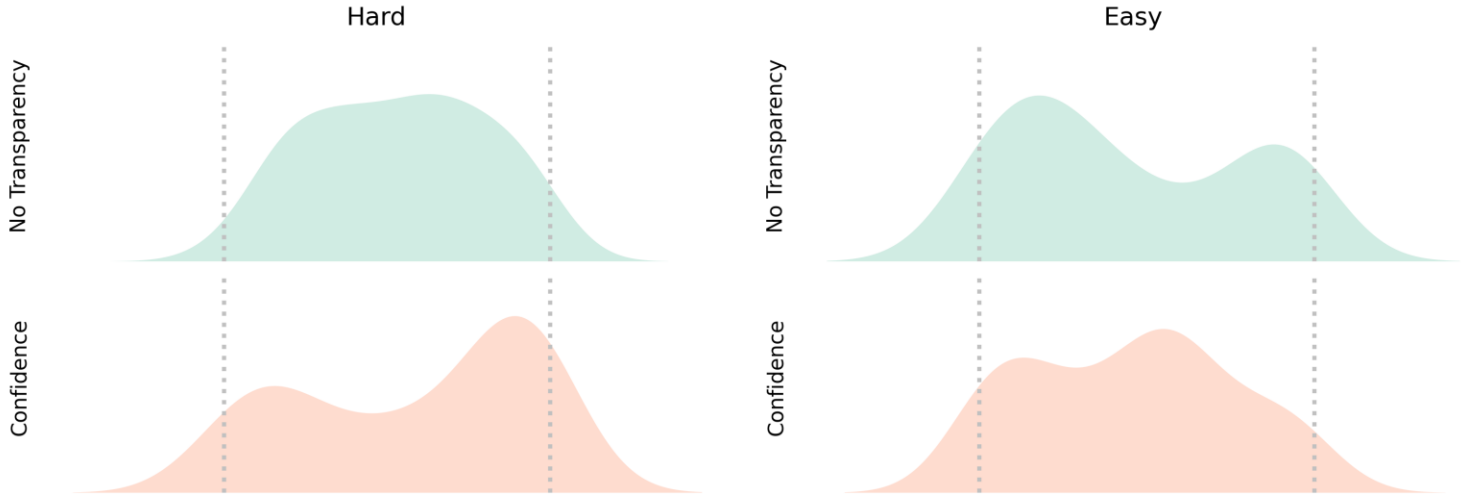

Is it more probable to use suggestions when there is more transparency? (cont.)

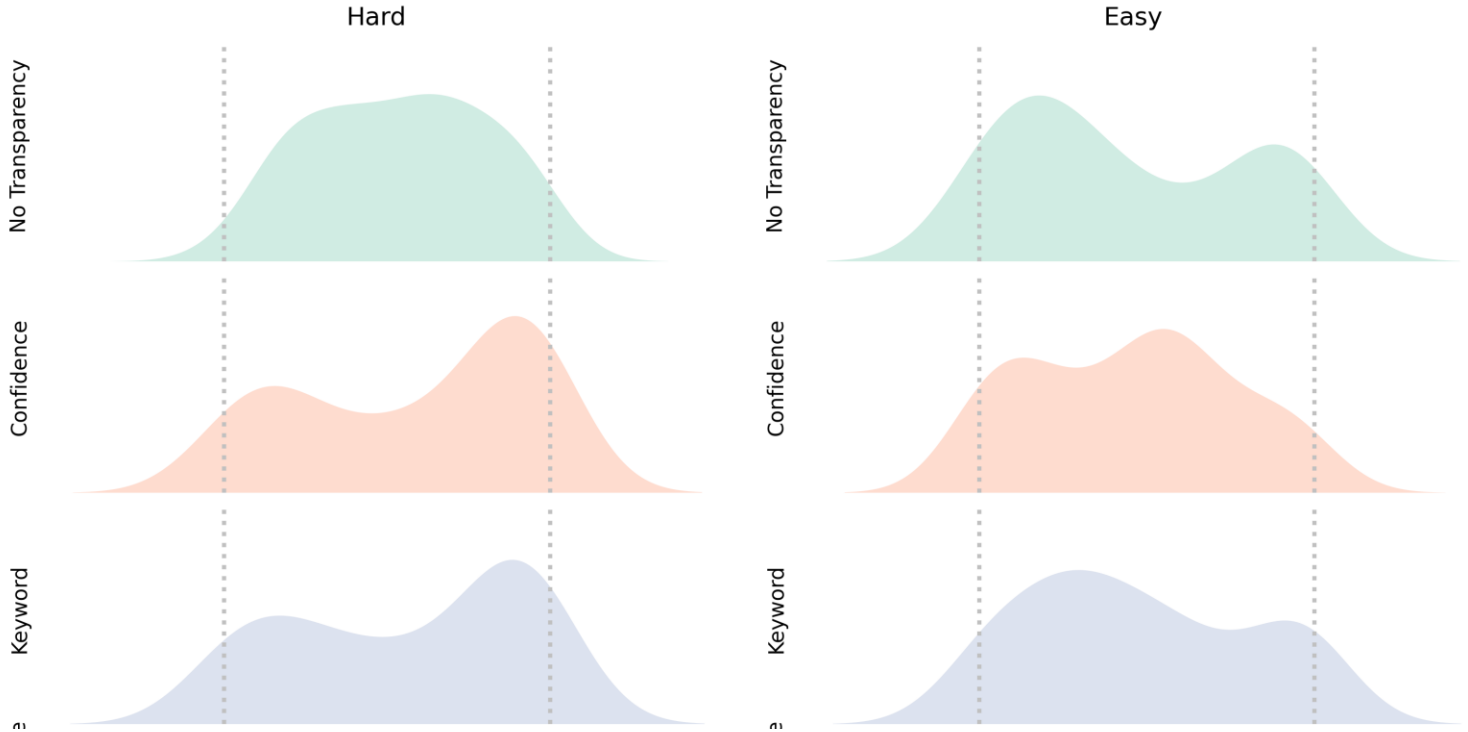

Is it more probable to use suggestions when there is more transparency? (cont.)

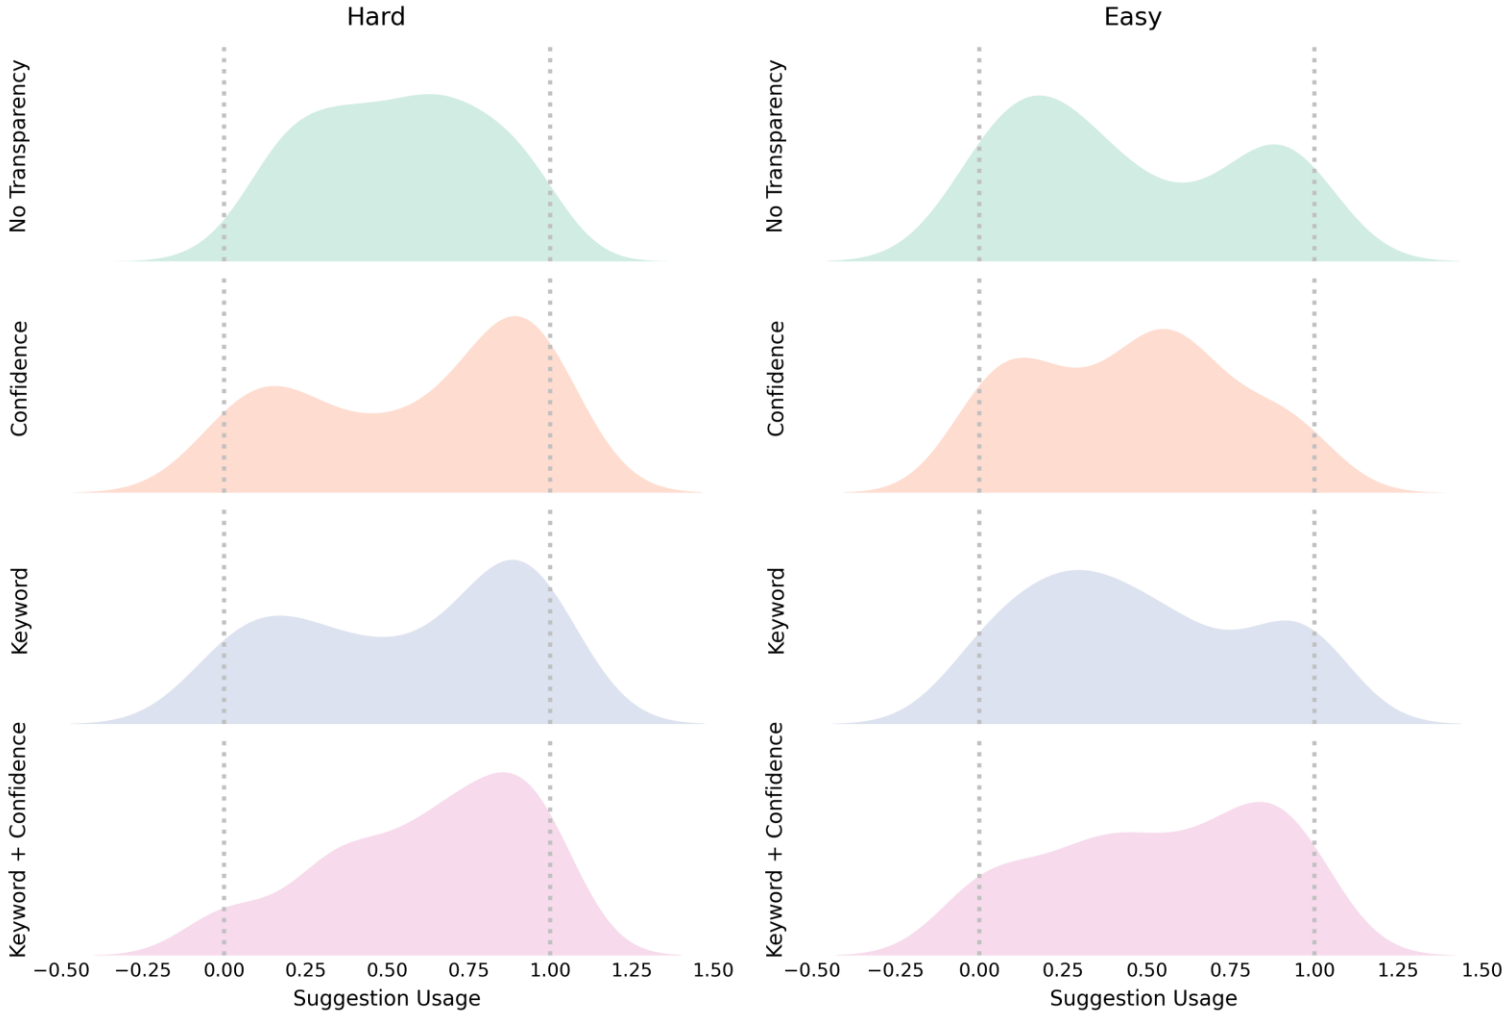

Is it more probable to use suggestions when there is more transparency? (cont.)

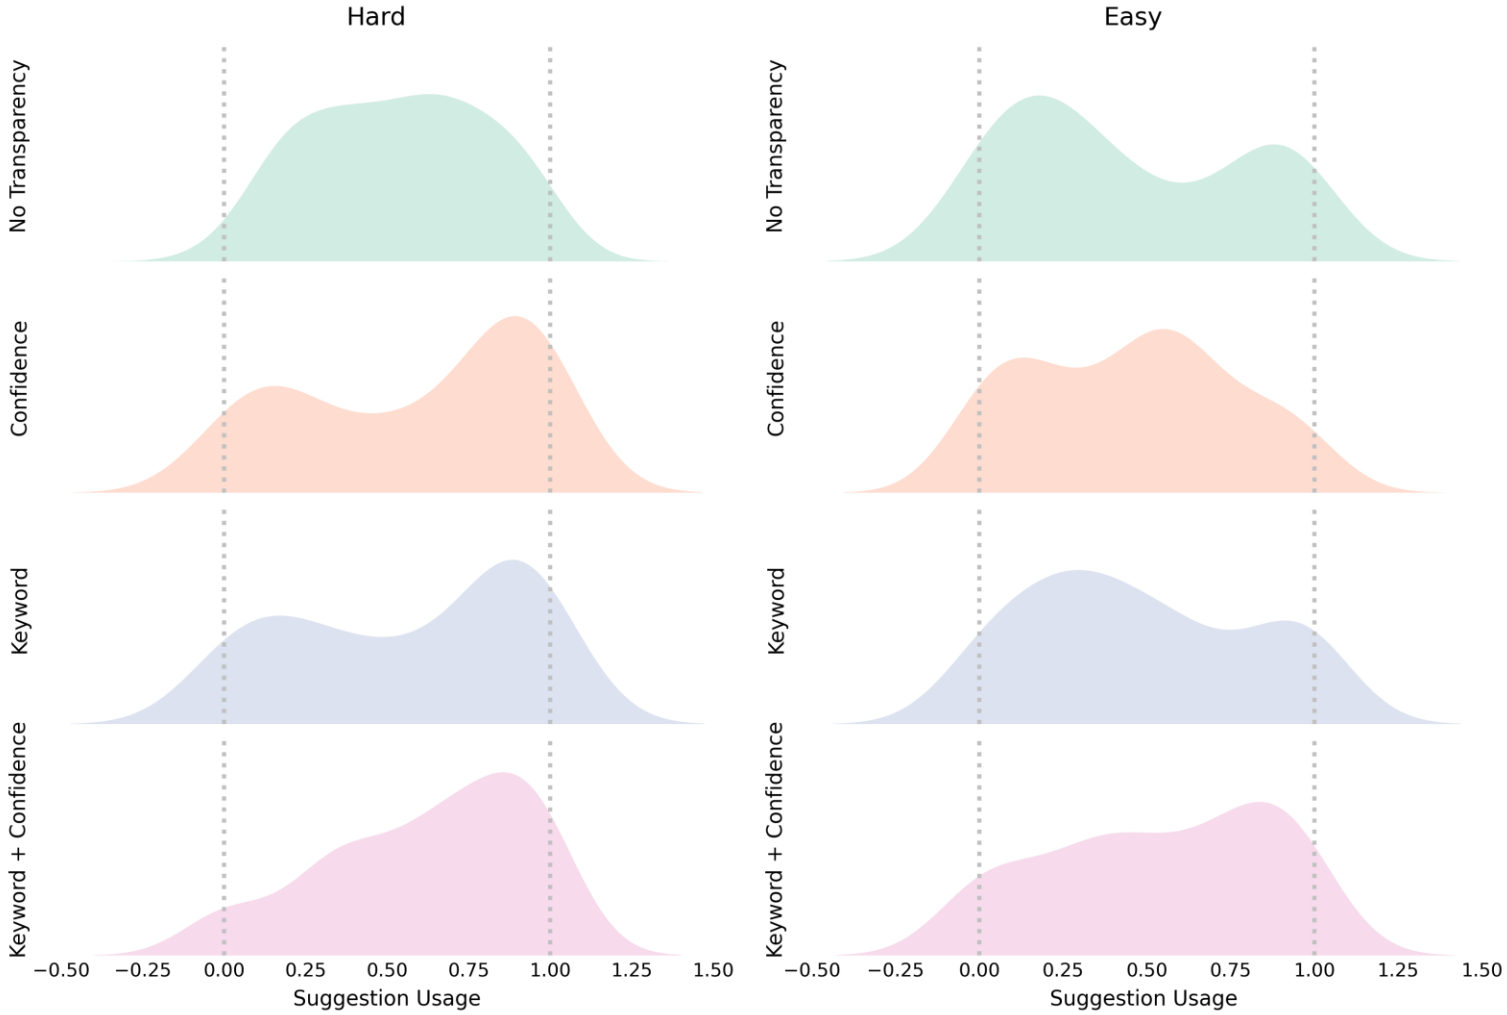

Is there evidence of overreliance?

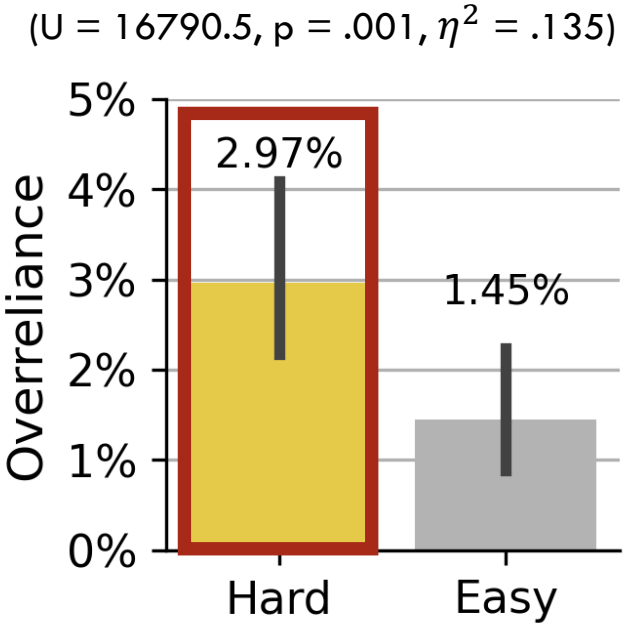

Participants in the hard task overrelied on suggestions more

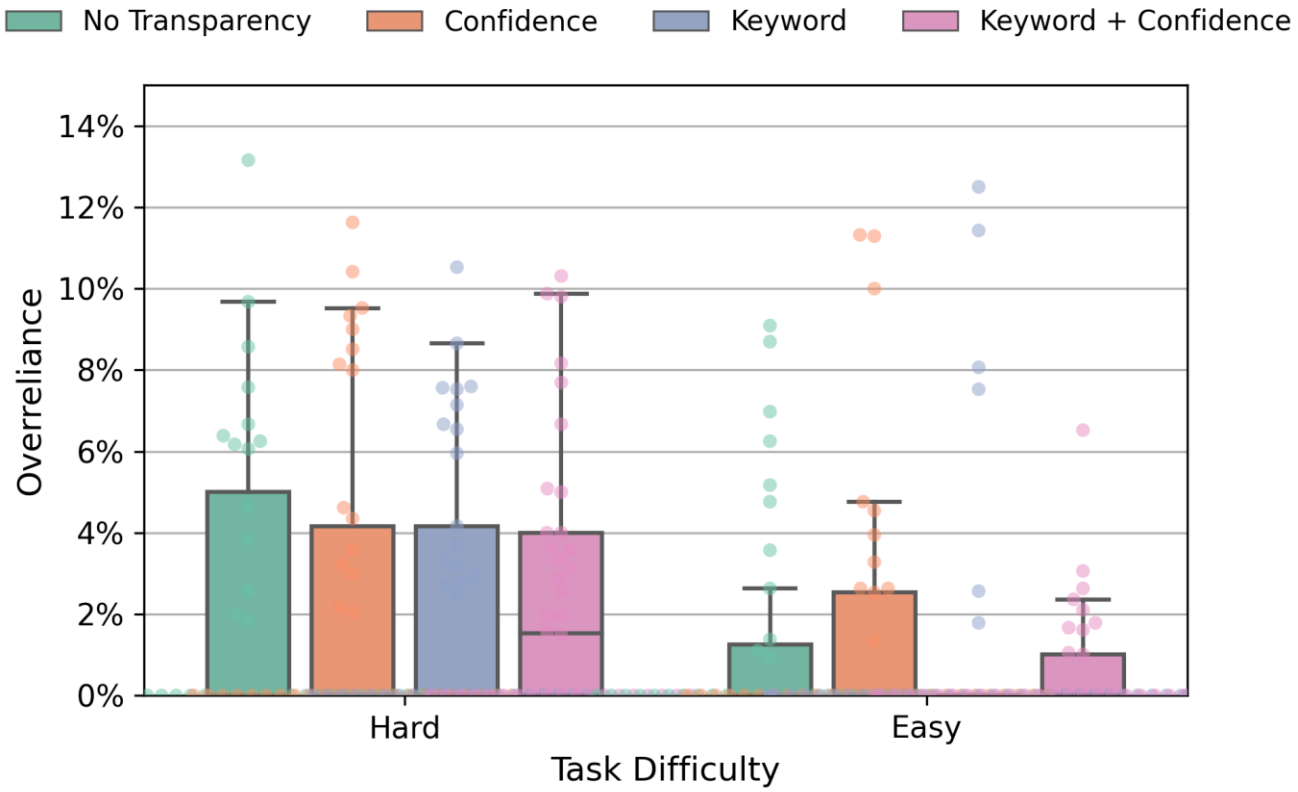

No difference in overreliance based on amount of transparency relayed to the participant

Does transparency affect subjective trust?

High levels of subjective trust

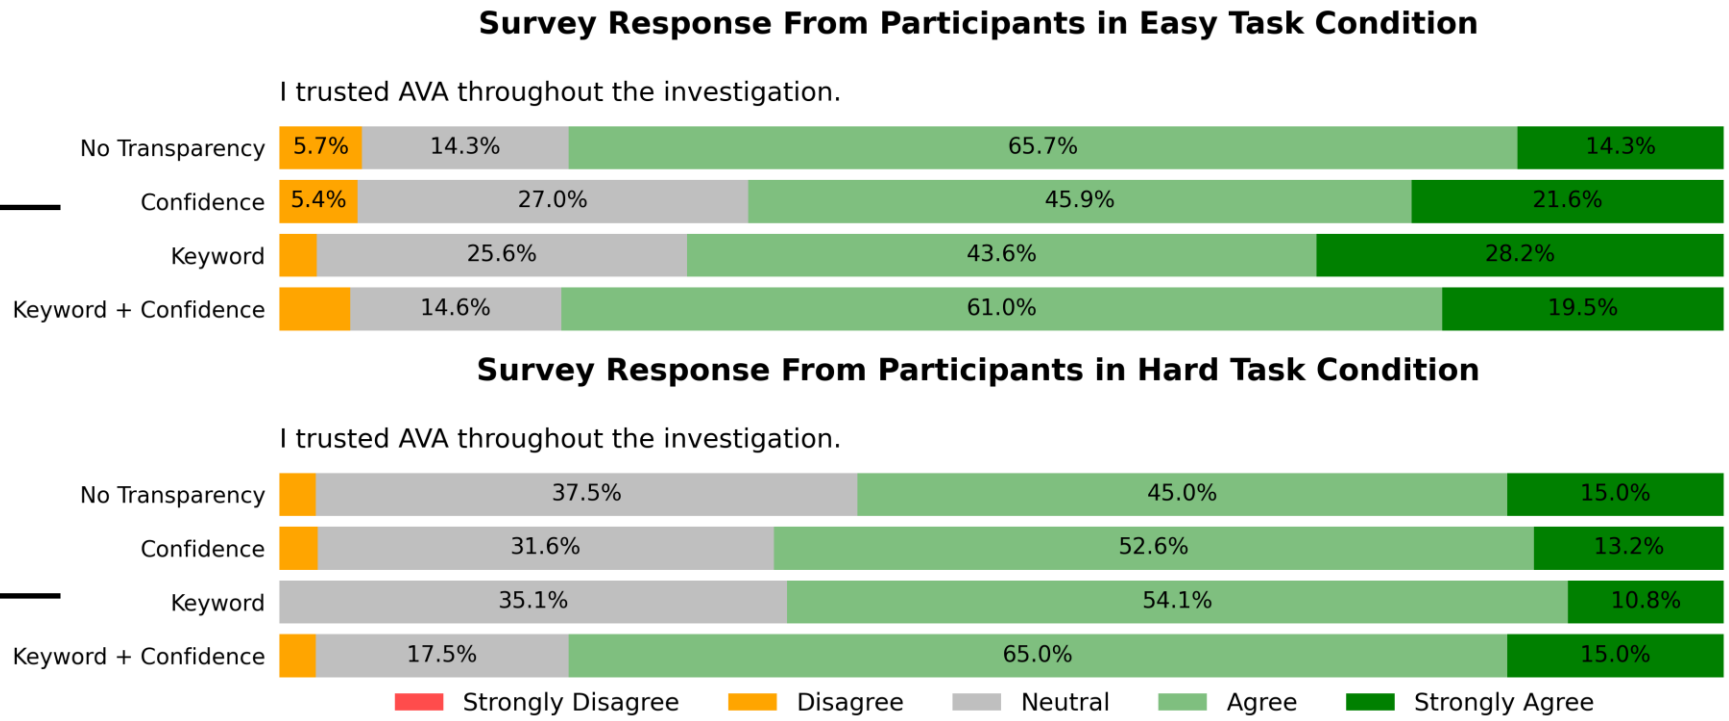

71% of recruited participants either agreed or strongly agreed that they trusted the system's guidance

Ha, S., Monadjemi, S., & Ottley, A. (2024, June). Guided By AI: Navigating Trust, Bias, and Data Exploration in AI-Guided Visual Analytics. In Computer Graphics Forum (Vol. 43, No. 3, p. e15108).

Does task difficulty affect subjective trust?

Participants who completed the hard task with guidance elicited beliefs of lower trust than those in who completed the easy task with guidance.

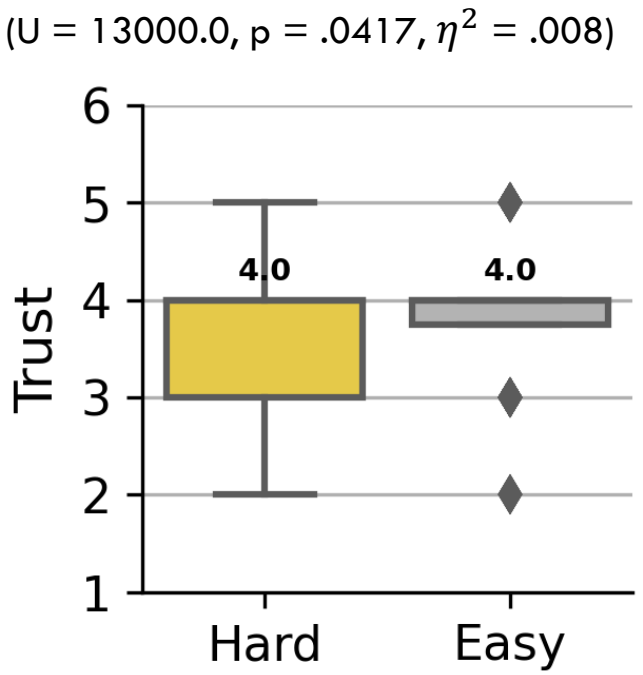

Does AI guidance encourage/discourage symptom diversity?

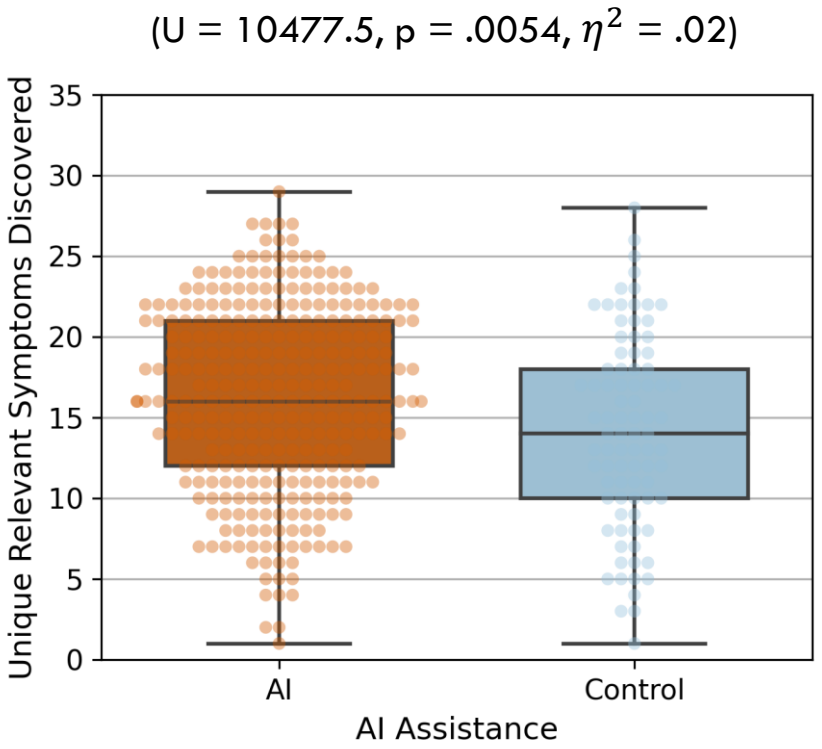

Participants who received AI guidance found more unique symptoms related to the epidemic than the control group.

Did AI Assistance increase the likelihood of uncovering the transmission sources?

Did AI Assistance increase the likelihood of uncovering the transmission sources?

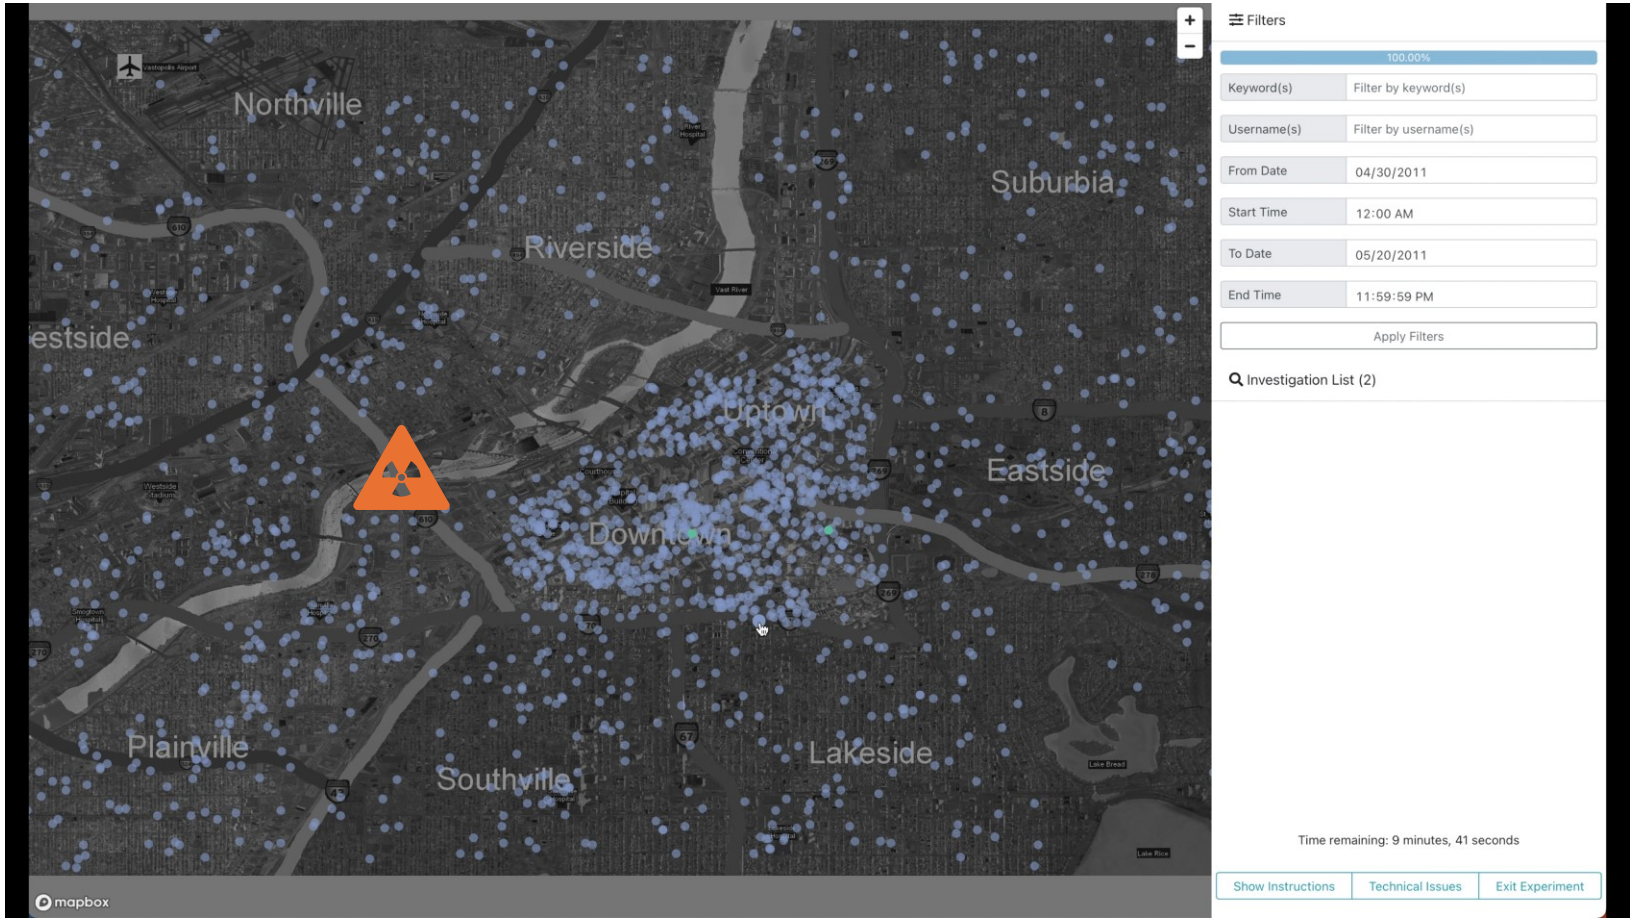

Did AI Assistance increase the likelihood of uncovering the transmission sources?

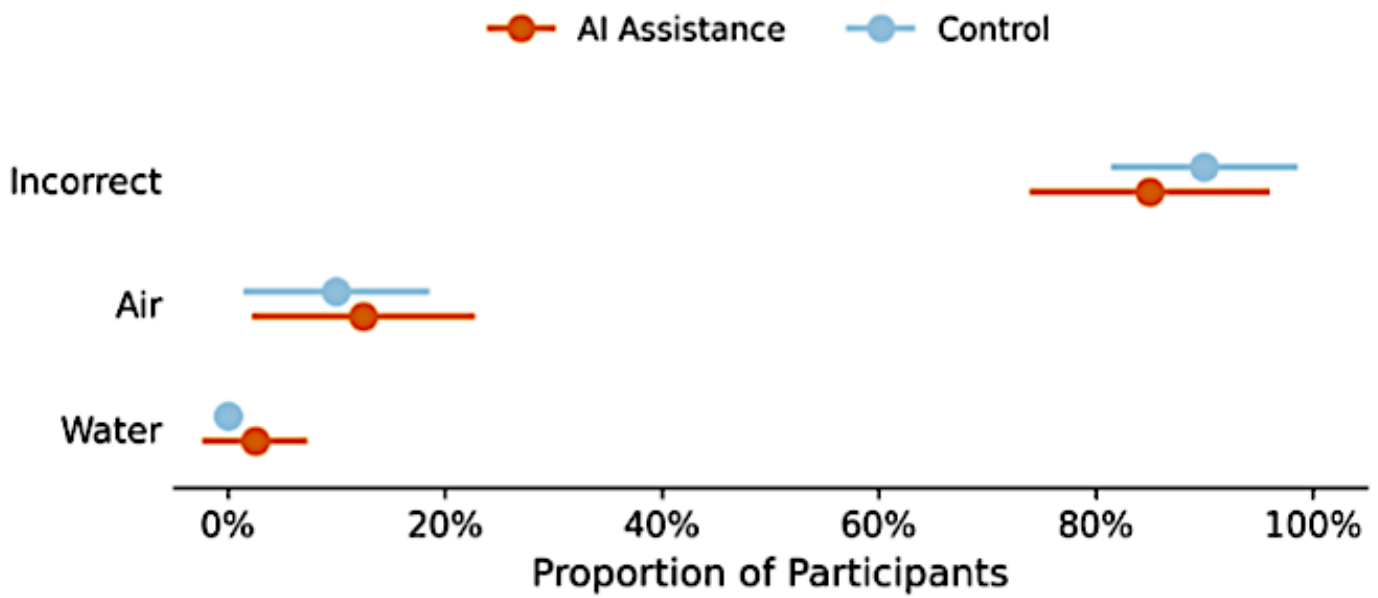

Successful identification of the epidemic's sources of transmission did not depend on AI assistance.

Takeaways

- Users in the hard task were more likely to use suggestions, but had lower levels of trust (behavioral vs cognitive trust)
- Trust remained unaffected by transparency
- Balance communication of transparency and information overload

Let’s talk about trust

trust 1 of 2 noun

ˈtrəst

1 **a :** firm belief in the character, ability, strength, or truth of someone or something

b : a person or thing in which confidence is placed

trust 1 of 2 noun

ˈtrəst

1 **a :** firm belief in the character, ability, strength, or truth of someone or something

b : a person or thing in which confidence is placed

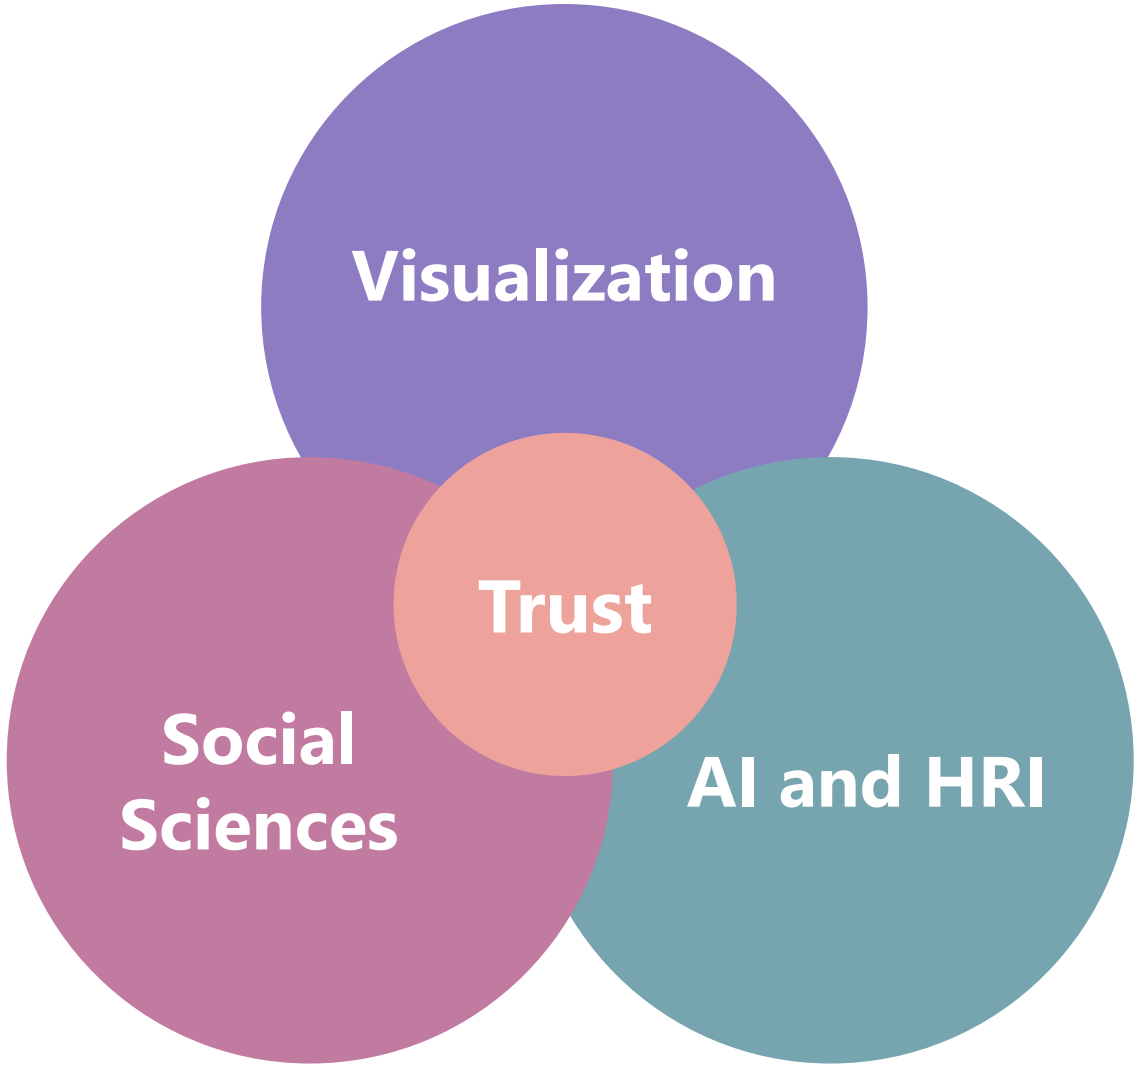

AI: Artificial Intelligence
HRI: Human-Robot Interaction

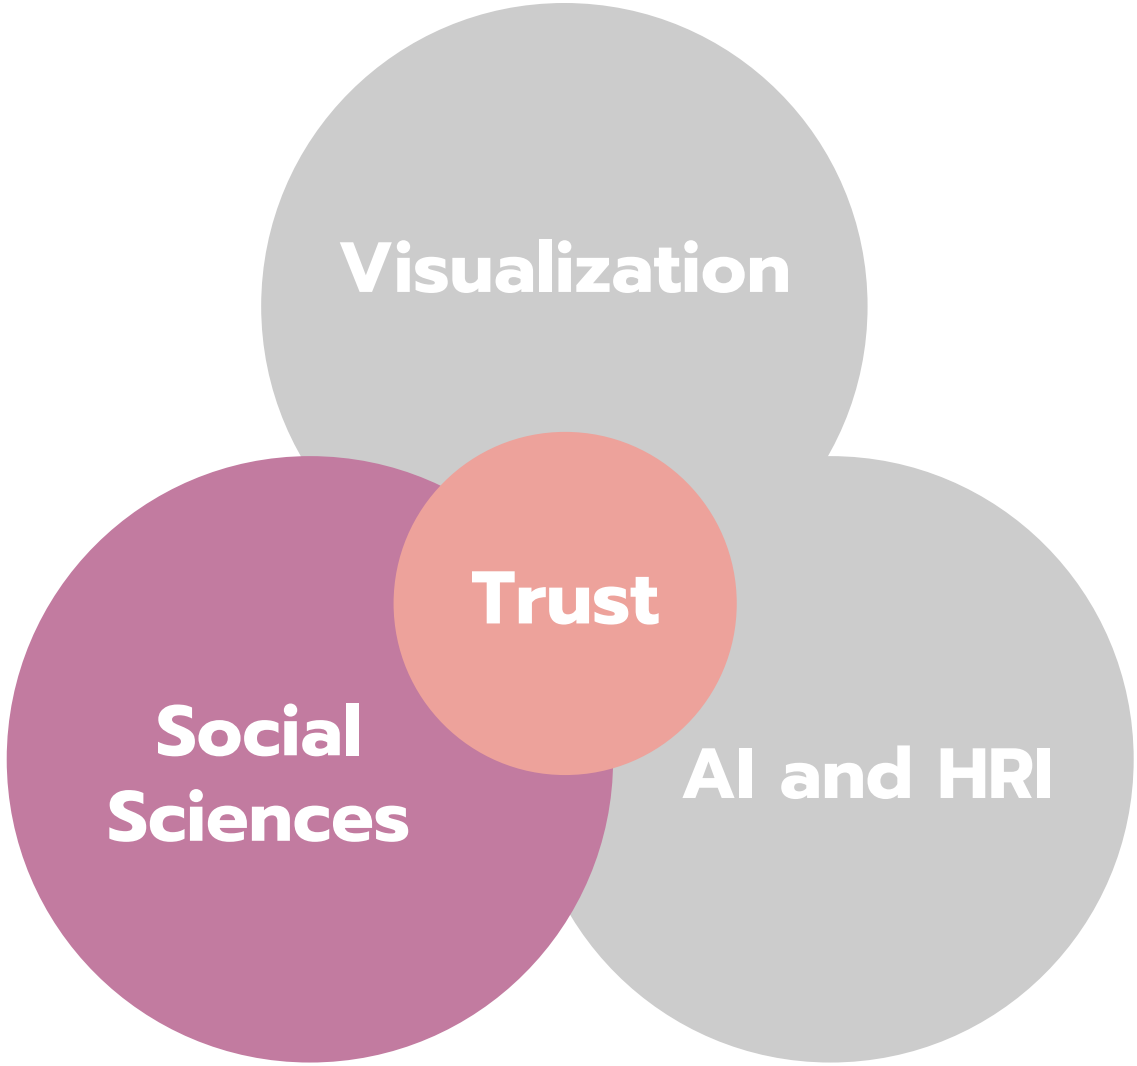

AI: Artificial Intelligence
HRI: Human-Robot Interaction

"Generalized expectancy that the oral or written statements of other people can be **relied** on."

Rotter, Julian B. "Generalized expectancies for interpersonal trust." *American psychologist* 26, no. 5 (1971): 443.

AI: Artificial Intelligence
HRI: Human-Robot Interaction

The mutual **confidence** that no party to an exchange will exploit another's vulnerabilities.

Sabel, Charles F. "Studied trust: Building new forms of cooperation in a volatile economy." *Human relations* 46, no. 9 (1993): 1133-1170.

AI: Artificial Intelligence
HRI: Human-Robot Interaction

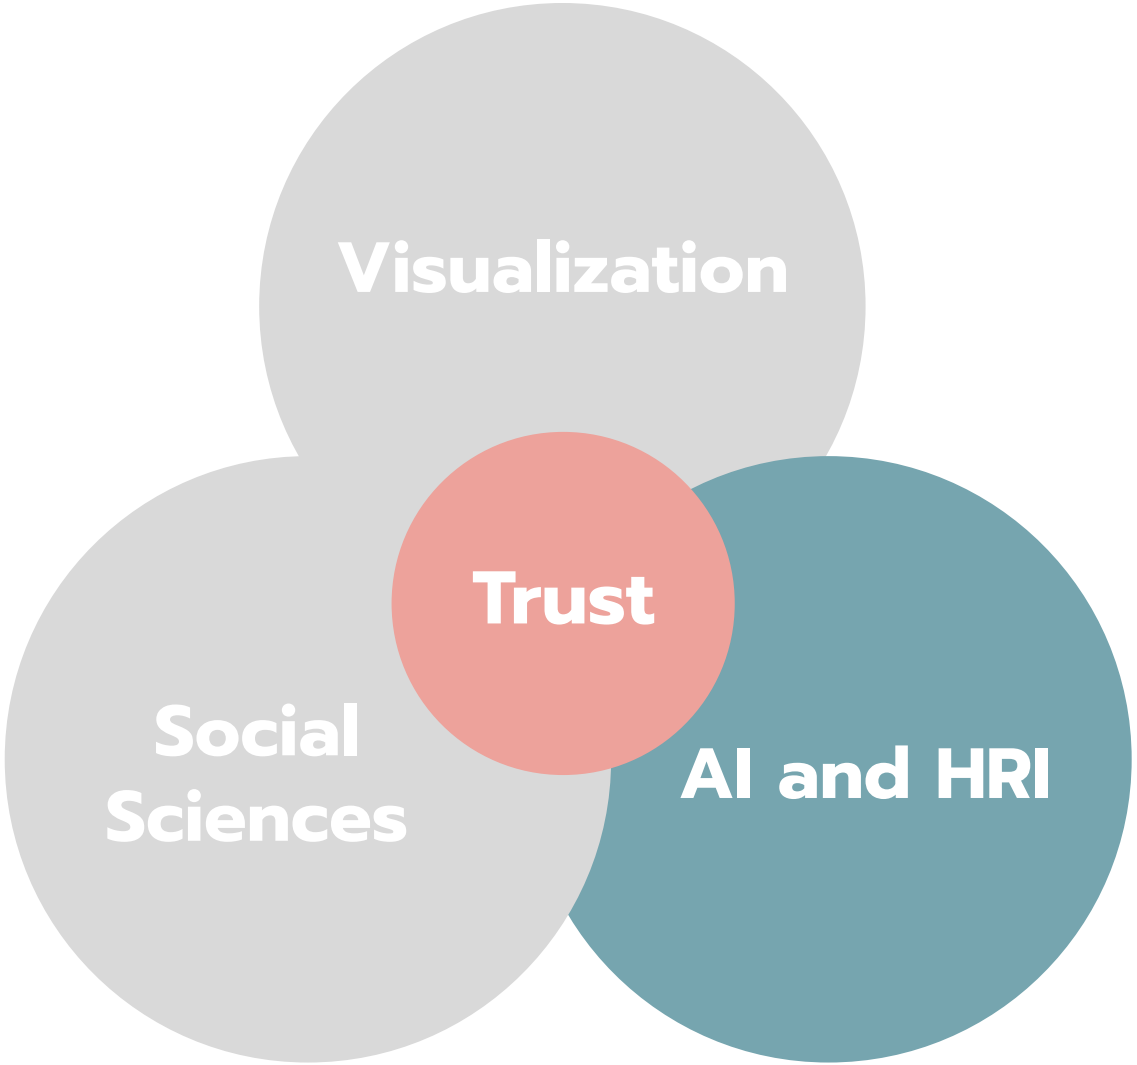

AI: Artificial Intelligence
HRI: Human-Robot Interaction

Trust = a dyadic relation in which one person accepts vulnerability because they expect that the other person's future action will have certain characteristics; these characteristics include some mix of performance (ability, reliability) and/or morality (honesty, integrity, and benevolence).

Ullman, Daniel, and Bertram F. Malle. "What does it mean to trust a robot? Steps toward a multidimensional measure of trust." In *Companion of the 2018 acm/ieee international conference on human-robot interaction*, pp. 263-264. 2018.

Reliance Intentions Scale

Items 1–4 were modified from Mayer et al. (1995).

1. If I had my way, I would NOT let the system have any influence over issues that are important to the task (surveillance). (reverse-coded)
2. I would be comfortable giving the system complete responsibility for the surveillance task.*
3. I really wish I had a good way to monitor the route decisions of the system. (reverse-coded)*
4. I would be comfortable allowing the system to implement its route decision, even if I could not monitor it.

Items 5–10 were added to gauge reliance intentions on the automation.

5. I would rely on the system without hesitation.*
6. I think using the system will lead to positive outcomes.*
7. I would feel comfortable relying on the system in the future.*
8. When the task was hard, I felt like I could depend on the system.*
9. If I were facing a very hard task in the future, I would want to have this system with me.*
10. I would be comfortable allowing this system to make all decisions.

Lyons, Joseph B., and Svyatoslav Y. Guznov. "Individual differences in human–machine trust: A multi-study look at the perfect automation schema." *Theoretical Issues in Ergonomics Science* 20, no. 4 (2019): 440-458.

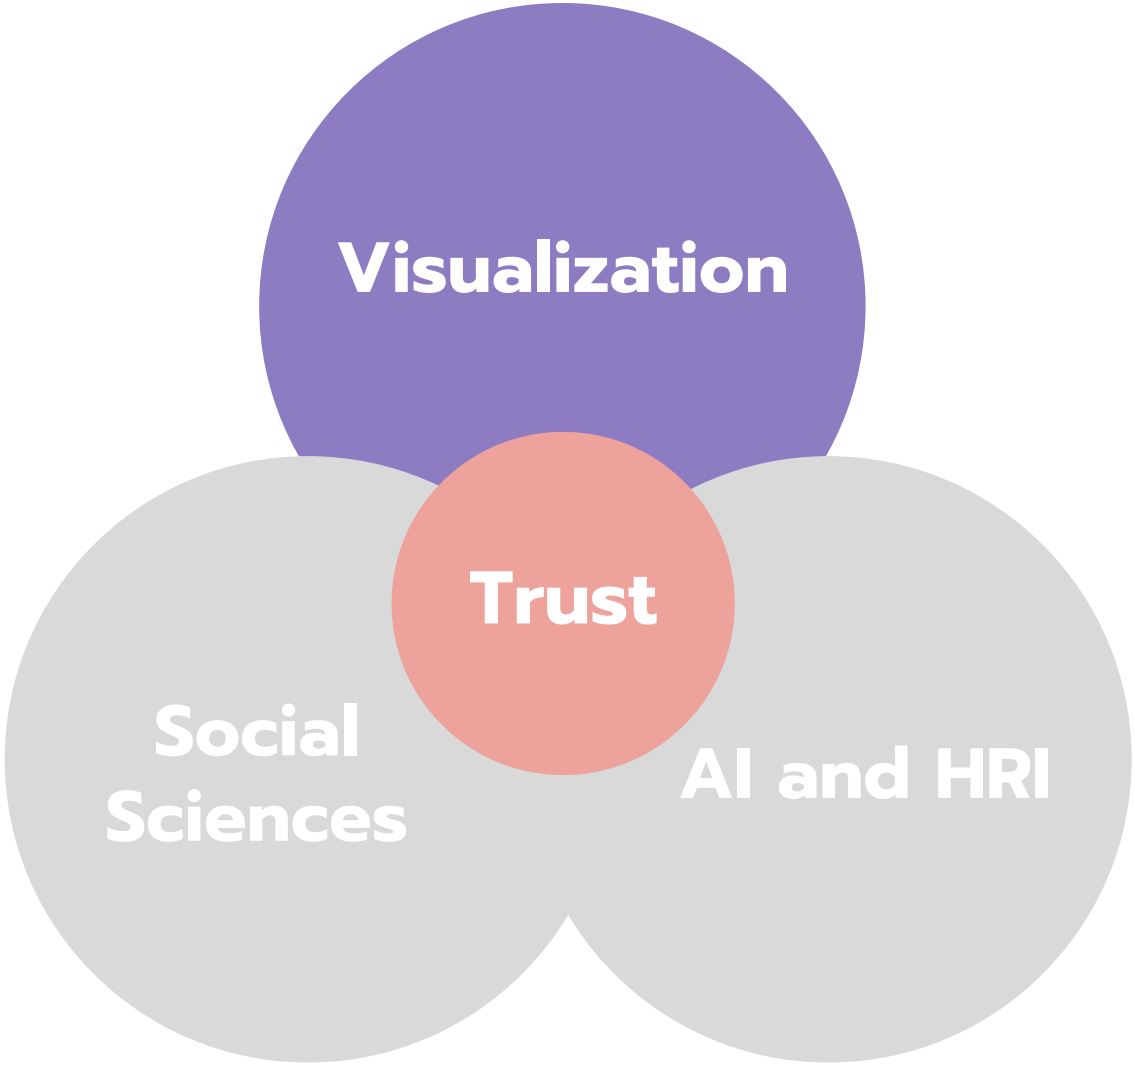
Visualization
Trust
Social Sciences
AI and HRI

What is “Trust” in Data Visualization?

Mayr et al. (2019) split into two:

- Trustworthiness
- **Trust Perception**

Do You Trust What You See?
Toward A Multidimensional Measure of Trust in Visualization

Saugat Pandey, Oen G. McKinley *†
Washington University in St. Louis

R. Jordan Crouser‡
Smith College

Alvitta Ottley§
Washington University in St. Louis

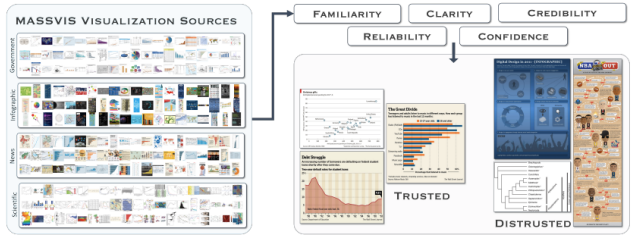

Figure 1: An illustrative overview of our experiments. Participants rated a diverse set of visualizations based on their FAMILIARITY, CLARITY, CREDIBILITY, RELIABILITY, and CONFIDENCE. We examined how these ratings align with visual features and trust rankings.

Abstract

Few concepts are as ubiquitous in computational fields as trust. However, in the case of information visualization, there are several unique and complex challenges, chief among them: defining and measuring trust. In this paper, we investigate the factors that influence trust in visualizations. We draw on the literature to identify five factors likely to affect trust: credibility, clarity, reliability, familiarity, and confidence. We then conduct two studies investigating these factors' relationship with visualization design features. In the first study, participants' credibility, understanding, and reliability ratings depended on the visualization design and its source. In the second study, we find these factors also align with subjective trust rankings. Our findings suggest that these five factors are important considerations for the design of trustworthy visualizations.

Index Terms: Human-centered computing—Visualization—Visualization design and evaluation methods

1 Introduction

Trust is an amorphous concept but is also an integral part of human interactions with machines. Many fields in Computer Science and Engineering have explored the topic of trust in recent decades, including Machine Learning [43], Automation and Robotics [1], and Information and Communication Technology [34]. Research in information visualizations has recently adopted a similar focus, with work aimed at finding the context where trust can be best applied to the field [16, 25], often pulling from other disciplines such as psychology [5] or related fields such as cartography [13].

These prior works highlight the need to better understand the role of trust in visualization. They suggest that the effectiveness of visualization depends on the user's trust in the data and representation. When users perceive a visualization as untrustworthy, they may hesitate to rely on the information presented or take action based on it. This could lead to poor decision-making, missed opportunities, or dangerous outcomes for high-stakes decisions. However, little research is on the factors influencing users' trust or guidance for designing more trustworthy visualizations. Further, there is currently no standard approach for defining and evaluating the trustworthiness of visualizations, making it difficult to compare and evaluate different visualizations consistently and objectively [5, 16, 25]. This paper takes a step toward defining and evaluating trust in information visualization, drawing upon prior work in various fields (e.g., AI, human-robot interaction, and psychology). We focus on the *trust perception* and behavior of users instead of the "objective" trustworthiness of the data or visualization. As a result, we expect that the visualization design will likely substantially affect perceived trust.

We present the findings of two user studies. Experiment 1 (Exp. 1) examined the correlation between visualization design features and five trust dimensions: credibility, clarity, reliability, familiarity, and confidence. Excluding familiarity, the trust dimensions correlated with the visualization's design features and source. Experiment 2 (Exp. 2) was a follow-up study to understand how people describe trust in the context of visualization design. We asked participants to rank a collection of visualizations based on how much they trusted them and then explain their ranking. We then coded their responses and found that they often related to the five trust dimensions from Exp. 1. For example, the comments aligned with the credibility dimensions discussed the importance of visualizations labeled with the data source. Some participants talked about the importance of visualizations being easy to understand, mirroring clarity. These findings suggest a positive relationship between the five trust dimensions and perceived trust in visualizations. We discuss how these results may help designers improve trust in their visualizations.

2 Approaches to Defining and Measuring Trust

Trust has received a great deal of attention, with researchers proposing many definitions and conceptualizations of trust [19]. Although it is impossible to capture every existing definition of trust, we highlight some key perspectives.

*contributed equally
†e-mail: p.saugat@wustl.edu, m.oen@wustl.edu
‡e-mail: jcrouser@smith.edu
§e-mail: alvitta@wustl.edu

a Multidimensional Framework and Study of Trust in Data Visualizations

Stefkovics, Johanna Beyer, Eric Moerth, Hanspeter Pfister,
Cindy Xiong Bearfield, Carolina Nobre

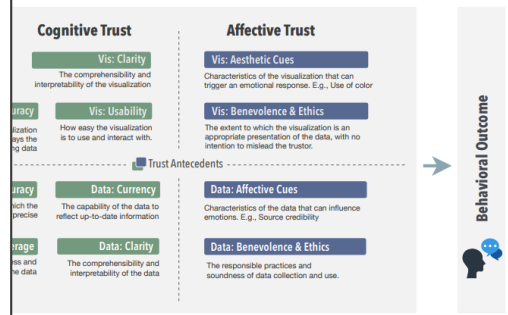

h outlines the development of trust in visualizations. The framework defines the different onents of trust (cognitive and affective trust). Both cognitive and affective trust can relate to a. *Individual characteristics* can play a role in shaping one's level of trust in visualizations, s a results of trust judgements.

data visualization, as it plays a crucial role in the interpretation and decision-making processes nces outlines the multi-dimensional factors that can play a role in trust formation, most data a single-item scale to measure trust. We address this gap by proposing a comprehensive, operationalization of trust in visualization. We do this by applying general theories of trust izing and extending earlier work and factors identified by studies in the visualization field. We st in visualization, to distinguish between cognitive and affective elements, as well as between cedents. We use our framework to design and run a large crowd-sourced study to quantify the rust in science visualizations. Our study provides empirical evidence for several aspects of our tably the impact of cognition, affective responses, and individual differences when establishing

nce, framework

is more widely adopted
f trust in visualizations

re with UMass Amherst.
@cs.umass.edu
Pfister are with Harvard
rd.edu,
mail:
E-mail:

201x. Date of Publication
For information on
to: reprints@ieee.org.
Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

becomes increasingly important, particularly when the data presented is urgent (e.g., climate change, Covid-19, etc.) For example, recent research exploring trust in Covid forecast visualizations showed how differences in visual encodings can significantly affect viewers' trust in the information and willingness to incorporate the information in their decision-making process [43].

The concept of trust and the factors that can play a role in its formation has long been explored in the field of social sciences. From that research, two parallel paradigms of trust emerge: trust defined by Mayer et al. as "the willingness of a party to be vulnerable to the actions of another party" [37] and trust defined by McAllister et al. as composed from cognition-based "rational" and affect-based "emotional" factors [40]. More specifically, *cognitive trust* is defined as trust based on the knowledge and evidence of someone's ability and achievements, while *affective trust* is defined as trust based on the emotional bond with someone [40]. Research in this field also defines the factors that precede trust formation as *antecedents* of trust. Examples of trust antecedents include benevolence and behavioral integrity [37]. In this

Building and Eroding: Exogenous and Endogenous Factors that Influence Subjective Trust in Visualization

R. Jordan Crouser, Syrine Matoussi, Lan Kung, Saugat Pandey, Oen G. McKinley, and Alvitta Ottley

Abstract

Trust is a subjective yet fundamental component of human-computer interaction, and is a determining factor in shaping the efficacy of data visualizations. Prior research has identified five dimensions of trust assessment in visualizations (credibility, clarity, reliability, familiarity, and confidence), and observed that these dimensions tend to vary predictably along with certain features of the visualization being evaluated. This raises a further question: how do the design features driving viewers' trust assessment vary with the characteristics of the viewers themselves? By reanalyzing data from these studies through the lens of individual differences, we build a more detailed map of the relationships between design features, individual characteristics, and trust behaviors. In particular, we model the distinct contributions of **endogenous** design features (such as visualization type, or the use of color) and **exogenous** user characteristics (such as visualization literacy), as well as the interactions between them. We then use these findings to make recommendations for individualized and adaptive visualization design.

Keywords: Trust, data visualization, individual differences, personality

1 Introduction

In today's information age, trust plays a critical role in influencing decision-making processes across different fields. Furthermore, the widespread dissemination of misinformation and disinformation has highlighted the urgent need to assist individuals in distinguishing truth from falsehood, presenting a significant societal challenge. As data visualizations continue to become indispensable tools for conveying complex information in accessible forms, it is important to acknowledge that they, too, are susceptible to manipulation, distortion, and misinterpretation like any other mode of communication. Thus, understanding the factors that influence people's trust in specific data visualizations is critical for designing effective reasoning aides.

As individuals engage with data visualizations, they are implicitly working to evaluate the accuracy and credibility of the presented information in order to inform their judgments and actions. Trust inspires confidence, and when the data underlying the visualization is sound, this confidence can support well-informed decision-making. However, trust doesn't exist by itself – it is relational. The same degree of trust in a *misleading* visualization or mistrust in a faithful or unbiased visualization can propagate erroneous conclusions and misguided actions.

Furthermore, the significance of trust in visualizations and technology extends to its profound implications for public perception and engagement across various sectors. In contexts such as organizational integration of technologies like artificial intelligence [9] and web services security [18], trust is critical for user acceptance and adoption [2]. In domains such as journalism, science communication, and public health, visualizations serve as powerful tools for engaging audiences and conveying complex information. However, the effectiveness of these visualizations hinges on their perceived trustworthiness [7]. A nuanced understanding of the factors that shape trust can empower practitioners to craft visualizations that resonate with their audiences, thereby fostering heightened engagement and comprehension.

Pandey et al. conducted two studies to explore the relationships between various visual design features and five interrelated facets of trust: credibility, clarity, reliability, familiarity, and confidence [15]. The first study asked participants to rate various visualizations along these dimensions and identified several design features that have a significant correlation with participants' subjective perceptions of trust. They observed that colorful visualizations and visual embellishments garnered greater favor among participants. Moreover, visualizations from news media were perceived as *more credible* and *more reliable* than those from scientific or governmental agencies, even in cases where information regarding the source of the visualization was not explicitly available. This suggests that participants may be picking up on disciplinary norms around data-driven communication. It has been hypothesized that scientific and government entities' tendency toward technical and data-dense designs may render them less accessible (and therefore less trustworthy) to everyday viewers.

The second experiment investigated how individuals weigh these trust dimensions within the context of visualization design. First, stimuli from the first experiment were sampled to retain visualizations that had both high response rates and low variance and then were further down-sampled to only those examples with the highest and lowest scores along each dimension. Participants were then asked to assess each example along the original 5 dimensions, as well as their *overall trust* in the visualization. In this experiment, factors such as source credibility, content familiarity, and type of visualization emerged as significant correlates with overall trust rankings.

These findings underscore the complex interplay between visualization design and perceptions of trust. Moreover, we know that patterns of interactions with data visualizations are not universal, but are instead modulated by individual differences between users [13]. Building on these insights, this paper explores the relationships between factors *endogenous* to the visualization (i.e. visual metaphor, use of color, source, etc.) and *exogenous* factors (such as individual differences in personality or cognitive ability, educational background, and cultural influences), and how these factors combine to affect perceived trust. Striking a balance among these factors is paramount for effectively communicating information, fostering accurate comprehension, and helping decision-making across diverse audiences.

1.1 Contributions

This work makes the following contributions:

1. We conducted a supplemental reanalysis of data collected by Pandey et al. [15] through the lens of **individual differences**.
2. We identified *visualization type* and *visualization literacy* as key **endogenous** and **exogenous** predictors of trust.
3. We observed that endogenous and exogenous factors have **nontrivial interactions** in how they influence trust.

By considering these factors, we aim to deepen our understanding of the nuanced dynamics between visualization design and trust.

Do You Trust What You See?
Toward A Multidimensional Measure of Trust in Visualization

Saugat Pandey, Oen G. McKinley *†
Washington University in St. Louis

R. Jordan Crouser‡
Smith College

Alvitta Ottley§
Washington University in St. Louis

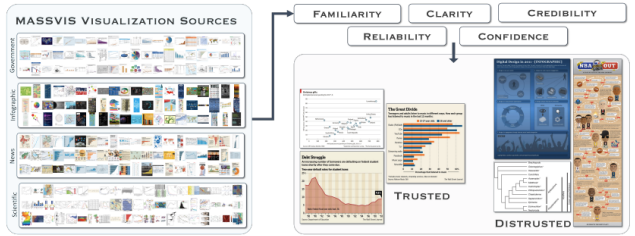

Figure 1: An illustrative overview of our experiments. Participants rated a diverse set of visualizations based on their FAMILIARITY, CLARITY, CREDIBILITY, RELIABILITY, and CONFIDENCE. We examined how these ratings align with visual features and trust rankings.

Abstract

Few concepts are as ubiquitous in computational fields as trust. However, in the case of information visualization, there are several unique and complex challenges, chief among them: defining and measuring trust. In this paper, we investigate the factors that influence trust in visualizations. We draw on the literature to identify five factors likely to affect trust: credibility, clarity, reliability, familiarity, and confidence. We then conduct two studies investigating these factors' relationship with visualization design features. In the first study, participants' credibility, understanding, and reliability ratings depended on the visualization design and its source. In the second study, we find these factors also align with subjective trust rankings. Our findings suggest that these five factors are important considerations for the design of trustworthy visualizations.

Index Terms: Human-centered computing—Visualization—Visualization design and evaluation methods

1 Introduction

Trust is an amorphous concept but is also an integral part of human interactions with machines. Many fields in Computer Science and Engineering have explored the topic of trust in recent decades, including Machine Learning [43], Automation and Robotics [1], and Information and Communication Technology [34]. Research in information visualizations has recently adopted a similar focus, with work aimed at finding the context where trust can be best applied to the field [16, 25], often pulling from other disciplines such as psychology [5] or related fields such as cartography [13].

These prior works highlight the need to better understand the role of trust in visualization. They suggest that the effectiveness of visualization depends on the user's trust in the data and representation. When users perceive a visualization as untrustworthy, they may hesitate to rely on the information presented or take action based on it. This could lead to poor decision-making, missed opportunities, or dangerous outcomes for high-stakes decisions. However, little research is on the factors influencing users' trust or guidance for designing more trustworthy visualizations. Further, there is currently no standard approach for defining and evaluating the trustworthiness of visualizations, making it difficult to compare and evaluate different visualizations consistently and objectively [5, 16, 25]. This paper takes a step toward defining and evaluating trust in information visualization, drawing upon prior work in various fields (e.g., AI, human-robot interaction, and psychology). We focus on the *trust perception* and behavior of users instead of the "objective" trustworthiness of the data or visualization. As a result, we expect that the visualization design will likely substantially affect perceived trust.

We present the findings of two user studies. Experiment 1 (Exp. 1) examined the correlation between visualization design features and five trust dimensions: credibility, clarity, reliability, familiarity, and confidence. Excluding familiarity, the trust dimensions correlated with the visualization's design features and source. Experiment 2 (Exp. 2) was a follow-up study to understand how people describe trust in the context of visualization design. We asked participants to rank a collection of visualizations based on how much they trusted them and then explain their ranking. We then coded their responses and found that they often related to the five trust dimensions from Exp. 1. For example, the comments aligned with the credibility dimensions discussed the importance of visualizations labeled with the data source. Some participants talked about the importance of visualizations being easy to understand, mirroring clarity. These findings suggest a positive relationship between the five trust dimensions and perceived trust in visualizations. We discuss how these results may help designers improve trust in their visualizations.

2 Approaches to Defining and Measuring Trust

Trust has received a great deal of attention, with researchers proposing many definitions and conceptualizations of trust [19]. Although it is impossible to capture every existing definition of trust, we highlight some key perspectives.

*contributed equally
†e-mail: p.saugat@wustl.edu, m.oen@wustl.edu
‡e-mail: jcrouser@smith.edu
§e-mail: alvitta@wustl.edu

5 Dimensions:
Credibility, Clarity, Familiarity, Confidence, and Reliability

“Affective”
Aesthetics,
Ethics,
Values

“Cognitive”
Clarity,
Accuracy,
usability

Vistrust: a Multidimensional Framework and Empirical Study of Trust in Data Visualizations

Hamza Elhamdadi, Adam Stefkovics, Johanna Beyer, Eric Moerth, Hanspeter Pfister, Cindy Xiong Bearfield, Carolina Nobre

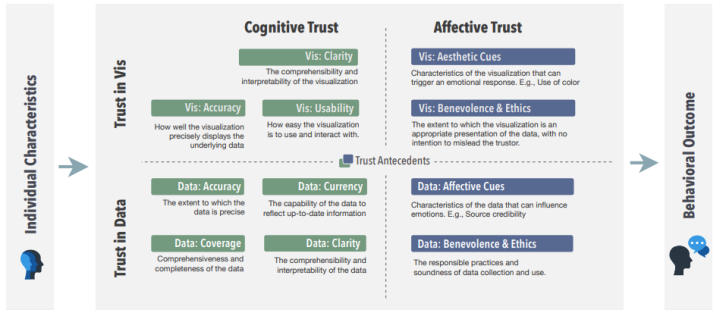

Fig. 1: An integrated framework, which outlines the development of trust in visualizations. The framework defines the different *trust antecedents* of the two basic components of trust (cognitive and affective trust). Both cognitive and affective trust can relate to the visualization and the underlying data. *Individual characteristics* can play a role in shaping one's level of trust in visualizations, and *behavioral outcomes* can emerge as a results of trust judgements.

Abstract—Trust is an essential aspect of data visualization, as it plays a crucial role in the interpretation and decision-making processes of users. While research in social sciences outlines the multi-dimensional factors that can play a role in trust formation, most data visualization trust researchers employ a single-item scale to measure trust. We address this gap by proposing a comprehensive, multidimensional conceptualization and operationalization of trust in visualization. We do this by applying general theories of trust from social sciences, as well as synthesizing and extending earlier work and factors identified by studies in the visualization field. We apply a two-dimensional approach to trust in visualization, to distinguish between cognitive and affective elements, as well as between visualization and data-specific trust antecedents. We use our framework to design and run a large crowd-sourced study to quantify the role of visual complexity in establishing trust in science visualizations. Our study provides empirical evidence for several aspects of our proposed theoretical framework, most notably the impact of cognition, affective responses, and individual differences when establishing trust in visualizations.

Index Terms—Trust, visualization, science, framework

1 Introduction

As the field of data visualization matures and is more widely adopted in public settings, understanding the role of trust in visualizations becomes increasingly important, particularly when the data presented is urgent (e.g., climate change, Covid-19, etc.) For example, recent research exploring trust in Covid forecast visualizations showed how differences in visual encodings can significantly affect viewers' trust in the information and willingness to incorporate the information in their decision-making process [43].

The concept of trust and the factors that can play a role in its formation has long been explored in the field of social sciences. From that research, two parallel paradigms of trust emerge: trust defined by Mayer et al. as "the willingness of a party to be vulnerable to the actions of another party" [37] and trust defined by McAllister et al. as composed from cognition-based "rational" and affect-based "emotional" factors [40]. More specifically, *cognitive trust* is defined as trust based on the knowledge and evidence of someone's ability and achievements, while *affective trust* is defined as trust based on the emotional bond with someone [40]. Research in this field also defines the factors that precede trust formation as *antecedents* of trust. Examples of trust antecedents include benevolence and behavioral integrity [37]. In this

- *Hamza Elhamdadi and Cindy Xiong Bearfield are with UMass Amherst. E-mails: helhamdadi@umass.edu, cindy.xiong@cs.umass.edu*
- *Adam Stefkovics, Johanna Beyer and Hanspeter Pfister are with Harvard University E-mails: adamstefkovics@fas.harvard.edu, jbeyer@g.harvard.edu, hpfister@g.harvard.edu*
- *Eric Moerth is with Harvard Medical School E-mail: ericmoerth@g.harvard.edu*
- *Carolina Nobre is with the University of Toronto E-mail: cnobre@cs.toronto.edu*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

“Endogenous” Factors

Visual metaphor, color choice, design, source

“Exogenous” Factors

Personality, cognitive ability, culture, education

Building and Eroding: Exogenous and Endogenous Factors that Influence Subjective Trust in Visualization

R. Jordan Crouser, Syrine Matoussi, Lan Kung, Saugat Pandey, Oen G. McKinley, and Alvitta Ottley

Abstract

Trust is a subjective yet fundamental component of human-computer interaction, and is a determining factor in shaping the efficacy of data visualizations. Prior research has identified five dimensions of trust assessment in visualizations (credibility, clarity, reliability, familiarity, and confidence), and observed that these dimensions tend to vary predictably along with certain features of the visualization being evaluated. This raises a further question: how do the design features driving viewers' trust assessment vary with the characteristics of the viewers themselves? By reanalyzing data from these studies through the lens of individual differences, we build a more detailed map of the relationships between design features, individual characteristics, and trust behaviors. In particular, we model the distinct contributions of **endogenous** design features (such as visualization type, or the use of color) and **exogenous** user characteristics (such as visualization literacy), as well as the interactions between them. We then use these findings to make recommendations for individualized and adaptive visualization design.

Keywords: Trust, data visualization, individual differences, personality

1 Introduction

In today's information age, trust plays a critical role in influencing decision-making processes across different fields. Furthermore, the widespread dissemination of misinformation and disinformation has highlighted the urgent need to assist individuals in distinguishing truth from falsehood, presenting a significant societal challenge. As data visualizations continue to become indispensable tools for conveying complex information in accessible forms, it is important to acknowledge that they, too, are susceptible to manipulation, distortion, and misinterpretation like any other mode of communication. Thus, understanding the factors that influence people's trust in specific data visualizations is critical for designing effective reasoning aides.

As individuals engage with data visualizations, they are implicitly working to evaluate the accuracy and credibility of the presented information in order to inform their judgments and actions. Trust inspires confidence, and when the data underlying the visualization is sound, this confidence can support well-informed decision-making. However, trust doesn't exist by itself – it is relational. The same degree of trust in a *misleading* visualization or mistrust in a faithful or unbiased visualization can propagate erroneous conclusions and misguided actions.

Furthermore, the significance of trust in visualizations and technology extends to its profound implications for public perception and engagement across various sectors. In contexts such as organizational integration of technologies like artificial intelligence [9] and web services security [18], trust is critical for user acceptance and adoption [2]. In domains such as journalism, science communication, and public health, visualizations serve as powerful tools for engaging audiences and conveying complex information. However, the effectiveness of these visualizations hinges on their perceived trustworthiness [7]. A nuanced understanding of the factors that shape trust can empower practitioners to craft visualizations that resonate with their audiences, thereby fostering heightened engagement and comprehension.

Pandey et al. conducted two studies to explore the relationships between various visual design features and five interrelated facets of trust: credibility, clarity, reliability, familiarity, and confidence [15]. The first study asked participants to rate various visualizations along these dimensions and identified several design features that have a significant correlation with participants' subjective perceptions of trust. They observed that colorful visualizations and visual embellishments garnered greater favor among participants. Moreover, visualizations from news media were perceived as *more credible* and *more reliable* than those from scientific or governmental agencies, even in cases where information regarding the source of the visualization was not explicitly available. This suggests that participants may be picking up on disciplinary norms around data-driven communication. It has been hypothesized that scientific and government entities' tendency toward technical and data-dense designs may render them less accessible (and therefore less trustworthy) to everyday viewers.

The second experiment investigated how individuals weigh these trust dimensions within the context of visualization design. First, stimuli from the first experiment were sampled to retain visualizations that had both high response rates and low variance and then were further down-sampled to only those examples with the highest and lowest scores along each dimension. Participants were then asked to assess each example along the original 5 dimensions, as well as their *overall trust* in the visualization. In this experiment, factors such as source credibility, content familiarity, and type of visualization emerged as significant correlates with overall trust rankings.

These findings underscore the complex interplay between visualization design and perceptions of trust. Moreover, we know that patterns of interactions with data visualizations are not universal, but are instead modulated by individual differences between users [13]. Building on these insights, this paper explores the relationships between factors *endogenous* to the visualization (i.e. visual metaphor, use of color, source, etc.) and *exogenous* factors (such as individual differences in personality or cognitive ability, educational background, and cultural influences), and how these factors combine to affect perceived trust. Striking a balance among these factors is paramount for effectively communicating information, fostering accurate comprehension, and helping decision-making across diverse audiences.

1.1 Contributions

This work makes the following contributions:

1. We conducted a supplemental reanalysis of data collected by Pandey et al. [15] through the lens of **individual differences**.
2. We identified *visualization type* and *visualization literacy* as key **endogenous** and **exogenous** predictors of trust.
3. We observed that endogenous and exogenous factors have **nontrivial interactions** in how they influence trust.

By considering these factors, we aim to deepen our understanding of the nuanced dynamics between visualization design and trust.

Future Work

- Incorporating visualization literacy measures
- Empirical validation of design guidelines
- How does interaction influence trust perception?
- Visual Analytics trust scale

Balancing Trust and Over-Reliance in Visual Analytics: The AI-in-the-Loop Dilemma

Alvitta Ottley

Visualization Interaction and Behavior Exploration

http://visualdata.wustl.edu/