BiasBuzz: Combining Visual Guidance with Haptic Feedback to Increase Awareness of Analytic Behavior during Visual Data Analysis

Jamal Paden*
jpaden@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Arpit Narechania*
arpitnarechania@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Alex Endert
endert@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

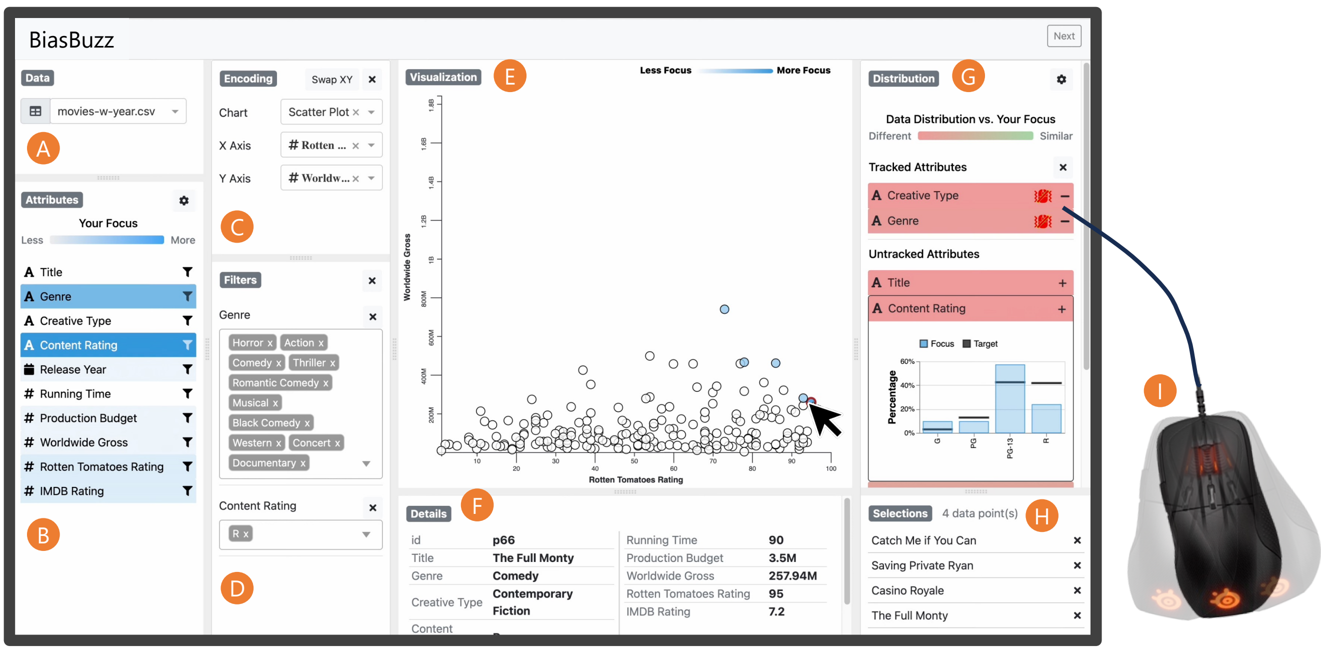

Figure 1: An existing visual data analysis tool, Lumos [18] (A)-(H), enhanced by wiring it to a gaming mouse [26] (I) to increase awareness of exploration biases. Our enhanced system (BiasBuzz) provisions visual guidance by highlighting a user's prior interactions (blue) and deviations from expected behavior (red, green) along with haptic feedback via mouse vibrations when the deviation is high.

ABSTRACT

During visual data analysis, users may inadvertently focus more on certain aspects of data, affecting analysis outcome(s). Existing tools primarily rely on visual cues (e.g., highlight already visited data) to increase user awareness of such analytic behaviors. We believe this single, visual modality is a passive form of guidance that adds to users' cognitive load already engaged in analysis. We investigate how a dual modality (visual guidance and haptic feedback) can capture users' attention and more actively guide them in their pursuits. We interface an existing visual data analysis tool with a gaming mouse. This enhanced system tracks user interactions and communicates biases by vibrating the mouse (haptic) and simultaneously displaying contextual information in the tool (visual). A formative study with nine users revealed that this dual modality increased analytical awareness in some cases but some users found the haptic mouse vibrations to be distracting and disturbing, informing the design of future multimodal user interfaces.

*Both authors contributed equally to this research.

CHI EA '24, May 11–16, 2024, Honolulu, HI, USA

ACM ISBN 979-8-4007-0331-7/24/05
https://doi.org/10.1145/3613905.3651064

CCS CONCEPTS

• Hardware → Haptic devices; • Human-centered computing → Empirical studies in visualization.

KEYWORDS

haptic, visualization, guidance, visual data analysis

ACM Reference Format:
Jamal Paden, Arpit Narechania, and Alex Endert. 2024. BiasBuzz: Combining Visual Guidance with Haptic Feedback to Increase Awareness of Analytic Behavior during Visual Data Analysis. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3613905.3651064

1 INTRODUCTION AND BACKGROUND

Interactive visualizations can succinctly present large information, facilitating knowledge discovery during data-driven analysis and decision-making processes. However, interacting with visualizations can sometimes result in insular data analysis practices that may cloud the user's judgment, negatively influencing the subsequent outcome(s) [6]. For example, Cho et al. [7] demonstrated the anchoring effect in visualizations – the tendency to focus too heavily on one piece of information when making decisions. Dimara et al. [8] studied the attraction effect in visualizations – where one's choice between two alternatives is influenced by the presence of an irrelevant (dominated) third alternative – and mitigated it by allowing the user to delete data points during a task or by altering the visual representation of the data altogether. Wall et al. introduced four definitions of the term "bias" in data visualization, relating to human cognitive, perceptual, and societal biases, and a fourth usage as a model mechanism [29]. Wall et al. also formalized a framework for measuring bias and presented six metrics that model bias from user interactions with data during visual data analysis [28]. Wall et al. [30] introduced "interaction traces" (visual traces of a user's interaction history) in a visualization as a means to increase awareness of and mitigate social biases in politics. Narechania et al. [18] presented interaction traces to make users more aware of their own unconscious biases wherein they emphasize certain parts of the data while neglecting others, also known as exploration biases [11].

Existing visual data analysis tools [10, 16–19, 30] primarily rely on visual cues (e.g., coloring visited data points darker than others [18]) to enhance user awareness of (biased) analytic behaviors. However, we believe this singular, visual modality is a passive form of guidance, potentially adding to users' cognitive load already engaged in analysis. Narechania et al. [18] found that visually presenting interaction traces increase user awareness but can sometimes cause confusion and go unnoticed. **So we ask:** *"How can we use alternate modalities, e.g., haptic feedback, as a stimulus to the existing visual feedback, to strengthen and reinforce the overall guidance?"*

Haptics refers to the science and technology involving the sense of touch, particularly focusing on the creation and study of tactile sensations and feedback [13, 23]. Haptic devices have been used in many applications such as remote systems for visually impaired people [25], anxiety and depression treatment [3], assistive communication technologies for children with autism [5], wrist-mounted devices for alerting users of warnings in a cybersecurity context [9], affecting the state of mind of users watching the news [22], and gaming [1, 32]. Akamatsu et al. showed that with a haptic mouse, users move faster and click targets within a wider area than users with a typical mouse [1]. Kyung et al. studied how a unique mouse with "force feedback" was more effective than a normal mouse at helping users recognize shapes in a task [15]. Terry et al. found that haptic mice reduce the response time spent on visual tasks on a computer [27]. Han et al. used an off-the-shelf haptic mouse for a study related to guidance in visualization and participants who used the haptic features performed better than users that did not [12]. In this work, we explore how a combination of visual guidance and haptic feedback can help users be *even more* aware of their analytic behavior during a visual data analysis task.

We enhanced an existing visual data analysis tool, Lumos [18], by interfacing it with a mouse capable of generating haptic feedback. This enhanced system (**BiasBuzz**) tracks user interactions with data, measures exploration biases, and communicates them to the user in the form of mouse vibrations (haptic feedback) and simultaneous display of contextual information in the user interface (visual guidance). This combination of visual and haptic elements seeks to create a more engaging experience for users during data analysis. We conducted a formative study with nine users to investigate the effectiveness of this dual (haptic feedback plus visual guidance) modality in increasing awareness of (biased) analytic behaviors. Our findings indicate that this dual modality can sometimes capture users' attention and actively guide them to 'fix' potentially 'biased' analytic behaviors. However, the haptic mouse vibrations, while effective, can also be distracting and/or disturbing, putting into context their usage during visual data analysis. We discuss implications of our study design to inform the design of future multimodal guidance-enriched user interfaces for visual data analysis.

2 HAPTIC FEEDBACK: DESIGN CHOICES AND CONSIDERATIONS

To design a visual data analysis system combining both visual and haptic feedback modalities, we identified a number of aspects to consider and choices to make. We illustrate these through a scenario. Consider a visual data analysis tool (e.g., Lumos [18]) where users upload a tabular dataset and perform analysis by creating different visualizations and applying relevant filters. To help the user not exhibit exploration bias, i.e., emphasize certain attributes and records more than others, the system tracks the user's interactions and visually presents any bias back to the user, in real-time. This tool achieves this by (1) highlighting already visited data attributes and records and (2) presenting the deviation of user's interaction pattern from the underlying distribution, computed as the AD metric by Wall et al. [28].

We intend for the feedback in the tool to "guide" the user to exhibit less exploration biases in their interactions than they did prior to the feedback. This means discouraging "biased" exploration methods and reinforcing "unbiased" exploration methods without highlighting specific data points in the interface.

The AD metric characterizes how a user's interactive behavior deviates from expected behavior and ranges from 0 (no bias) to 1 (high bias). By default, the system chooses a **proportional** baseline of expected behavior, wherein interactions with any given data point are equally likely and also reflecting the true underlying distributions of data attributes. For instance, if a user interacts primarily with 'Drama' movies among a dataset of movies that contains predominantly 'Action' movies, the AD metric for the *Genre* attribute will be high (more emphasis). If the user instead spent more time interacting with 'Drama' movies, proportional to the distributions in the dataset, the AD metric value for *Drama* would be low (less emphasis).

Next, consider this tool is connected to a haptic-enabled gaming mouse (e.g., SteelSeries 710 [26]) that appropriately vibrates from time to time to capture the user's attention. Next, we describe some of our considerations while designing the timing, duration, intensity, and pattern of these vibrations.

2.1 Vibration Timing: When to vibrate?

We considered triggering a mouse vibration every time exploration bias is detected for an attribute. The AD metric [28], used to quantify the deviation of user's interaction pattern of an attribute from its underlying distribution ranges from [0, 1], where 0 implies less deviation and 1 implies more deviation. Based on our own testing and pilot studies, we set a threshold of AD=0.7 above which an attribute is said to have exhibited exploration bias. This, however, can still result in multiple vibrations depending on the number of attributes whose AD values are above the threshold. Thus, we decided to allow the user to select the attributes they wish to track and only consider this subset for vibration.

2.2 Vibration Duration and Cooldown Period: How long should the vibration last?

Gaming mice have vibration motors built into them to provide haptic feedback during gameplay. These motors often generate heat when they are in use for extended periods or at high intensity. To prevent these motors from overheating, as a protective mechanism, these mice *cooldown* for a short time period before vibrating again. One of the implications of this behavior in our visual data analysis scenario is that if the user interacts twice in quick succession, and both times bias is detected, the mouse would still only vibrate once. Only after the cooldown period, if the detected bias is still active, will the mouse vibrate again. Between this timeframe, the vibrations can be considered 'lost', necessitating an alternative modality (e.g., visual) to communicate the same information.

2.3 Vibration Intensity: How strongly to vibrate?

Gaming mice often enable customization of the vibration intensity (or strength) and pattern during gameplay. In our visual data analysis scenario, we can map intensity to the amount of exploration bias (e.g., less bias is **'z'** whereas more bias is **'Z'**), where **'z'** and **'Z'** represent one vibration pulse. During our own testing and pilot studies, we noticed less variance between different vibration intensities, making it difficult for users to differentiate between them. Thus, we decided to set the default vibration intensity at a constant, highest level (**'Z'**).

2.4 Vibration Pattern: How to vibrate?

Gaming mice often enable customization of the vibration pattern. To design our visual data analysis scenario, we reviewed the design space of haptics [21, 24, 31] and considered mapping different vibration patterns to different attributes (that are exhibiting bias). For example, given a movies dataset, **'Z'** represents one vibration pulse. A biased *Genre* attribute would make the mouse vibrate as **'ZZ..ZZ..ZZ'**, wherein the mouse vibrates for a short duration two times every time bias is detected in the *Genre* attribute. Also in this example, a biased *Budget* attribute may vibrate as **'ZZZ.ZZZ.ZZZ'**, wherein the mouse vibrates three times every time bias is detected in the *Budget* attribute.

During our own testing and pilot studies, we found that keeping track of different vibration patterns can become confusing for the user. Thus, we set the default vibration pattern to a single, long pulse set to the highest strength and show additional contextual information, e.g., the attribute(s) name and its AD value, visually in the user interface (UI). Note that a gaming mouse generally does not have its own display to show this information, hence we have chosen to utilize the tool's UI.

3 EVALUATION

We conducted a formative study to understand how visual and haptic feedback can together increase user awareness and guide them to mitigate exploration biases during visual data analysis.

3.1 Study Design

Participants: We recruited nine participants enrolled in a bachelors degree program in a computing or related field at a public university in the United States. We screened these participants based on their self-reported visualization literacy (≥3/5). Demographically, our participants comprised seven men and two women, all in the age range of 21 to 32 years.
Dataset. 709 movies with 9 attributes: *Production Budget* (#), *Worldwide Gross* (#), *Running Time* (#), *IMDB Rating* (#), *Rotten Tomatoes Rating* (#), *Release Year* (📅), *Content Rating* (**A**), *Genre* (**A**), and *Creative Type* (**A**).
Task. "*Create a list of 10 movies that you would like to watch. These movies should reflect the underlying dataset as it relates to Release Year, Genre, and Content rating. Feel free to use the tracking feature to help you achieve your goal.*"
Study Session. We conducted the study in-person in a controlled lab environment. After providing consent, participants saw a video tutorial that demonstrated the features of the visual data analysis tool and the gaming mouse (5 minutes). Participants then performed a practice task on a dataset of cars to get acquainted with the study interfaces (5 minutes) before starting the actual task on the dataset of movies (20 minutes). After the task, participants provided feedback via a post-study questionnaire and a short debriefing interview (5 minutes). Each study session lasted about 60 minutes for which we compensated each participant with a $15 gift card. We encouraged participants to think aloud during this task recorded the screen and audio for subsequent qualitative analysis.

3.2 Study Prototype Interfaces

3.2.1 Visual Interface. We enhanced an existing, open source visual data analysis tool, **Lumos** [18] (Figure 1). Lumos enables users to load a tabular dataset **(A)**, inspect its attributes and corresponding data distributions **(B)**, apply filter criteria **(D)**, and assign attributes to visual encodings **(C)** to eventually create visualizations **(E)** and inspect raw data records **(F)**. Lumos tracks users' interactions with data attributes and records and presents them back to the user in the form of visual highlights (e.g., by coloring visited data points in shades of blue **(E)**). Lumos also determines if the user has over- or underemphasized certain attributes and records and by how much by computing the AD (Attribute Distribution) metric [28] **(G)**. The AD metric values lie between 0 and 1; higher the AD metric, higher the deviation between the user's interaction with a certain category/quantile of data attribute and its underlying data distribution, implying higher exploration bias. Lastly, the Selections panel **(H)** shows the list of selected data records (movies).

3.2.2 Haptic Interface. We used a **SteelSeries 710 gaming mouse** that can be made to vibrate programmatically [26] **(I)**. We chose this specific model because of its diverse vibration-related capabilities (e.g., timing, duration, intensity, pattern), ease of setup via an extensive API and documentation [26], and prior usage in a research study related to visualization [12].

3.2.3 Interfacing the Visual and Haptic Interfaces. We made the following enhancements to the Lumos UI to orchestrate the interactions with the haptic mouse, which is illustrated in Figure 2.

(Un)Tracking Attributes Not all attributes from a dataset might be important or relevant to the user's task (e.g., the *Age* attribute is irrelevant if the user's task is to ensure *Gender* diversity). Thus, we added the ability to "track" specific attributes, and only communicated AD bias for these "tracked" attributes. Lumos already supports the ability to bookmark one or more attributes, and we repurposed this to instead track one or more attributes **(G)**. When a user tracks one attribute, that attribute's AD metric value is compared to a preconfigured *high-bias threshold*=0.7, on a scale from 0 to 1. If the value is greater than the threshold, exploration bias is detected and subsequently communicated. When a user tracks multiple attributes, the *mean* of the AD metric values of the tracked attributes is computed and compared with the threshold, 0.7. If the value is greater than the threshold, exploration bias is reported.

Haptic Mouse Vibrations and Visual Icon Alerts. To report exploration bias(es) for the tracked attribute(s), we provided two modes: haptic mouse vibrations and visual icon alerts. Whenever exploration bias is detected, the mouse vibrates once for a split second. Note that our haptic mouse does not come with any kind of display; it just vibrates and lights up. Hence, it can only convey *when* there is bias but not *why or who is responsible for it.* Transmitting this information via Morse (or equivalent) code is out of scope for this study. Thus, to put the vibration into context, it is very important for the Lumos visual interface to show the corresponding attribute(s) and the AD metric values. To achieve this, we added visual alert icons next to each tracked attribute in the Distribution panel **(G)**.

Whenever the mouse vibrates, corresponding visual alert icons start flashing in a pulse animation (i.e., continuously increase and decrease in size), connecting the vibration to the corresponding attribute. When a user tracks multiple attributes and the mouse vibrates (i.e., when the mean AD metric value is greater than the threshold), the mouse also vibrates but the visual alert icon starts flashing only for those attributes whose individual AD metric value is greater than the threshold (i.e., who are, in a way, responsible for the overall exploration bias). This was a design choice to help the user formulate concrete *next step* interactions with specific attributes (e.g., the ones with the highest bias).

(Un)Muting Attributes We anticipated users wanting to stop experiencing the mouse vibrations either temporarily or permanently because of personal preference, distraction, or requirement of their ongoing analysis. Thus, we provided the capability to (un)mute attribute vibrations by toggle-clicking the visual alert icon. When a user mutes an attribute, the AD metric value of that particular attribute will not be used to trigger haptic feedback and visual alerts until the attribute is unmuted.

4 RESULTS

We report qualitative and quantitative findings from our user study. We transcribed participant audio recordings, divided the resultant transcripts into smaller sections, and two coders applied open coding [4]. All study material including participant interaction logs and detailed charts showing total number of attributes tracked and muted, total number of corresponding mouse vibrations, evolution of the AD metric, and usefulness scores are available for the interested reader in the supplemental material.

4.1 Qualitative Feedback

The overall feedback for the usage of haptic mouse and visual icon alerts was mixed. In the post-study questionnaire, participants scored their perceived utility of key aspects of the study on a Likert scale from 1 ("not useful at all") to 5 ("very useful"). All aspects including visual icon alerts (μ = 2.56), haptic mouse vibrations (μ = 3.33), the ability to track attributes (μ = 3.56), and the ability to mute attributes (μ = 2.67) received mixed scores. Findings from the qualitative analysis also resonated with the aforementioned sentiment, described next.

P1-P9 refer to the nine participants. P1, P5, P9 were positive about both the haptic mouse and the visual icon alerts. P5 said, *"the mouse vibrations and visual alerts are [both] very good at drawing your attention towards data points you've been missing out on."* P1 said, *"I think [the vibrations] reminded me of my goal, so they changed my attention to focus on the tracked attributes."* P1 also noted they *"didn't notice the visual attribute alerts as much compared to the haptic feedback, but [they] like that it shows red when [they] haven't looked at data proportionate to that attribute."*

On the contrary, P2, P3, P6 disliked both. P2 said, *"I barely spent any time [with the Distribution Panel] near the beginning of the task and I already feel punished [due to high AD values]."* They projected *"[they] might get immune to it eventually and discard it as a nuisance rather than something that's giving helpful information."* P2 said, *"I'd prefer a post-facto email with suggestions rather than instant haptic punishments."* P6 did not understand the mouse vibrations or visual icon alerts very well. According to them, *"There was [high] latency between the event and the vibration so [they] had a hard time linking the vibration to its meaning."* Because *"[they] did not figure out how the mouse vibration worked [they] did not understand the [visual] icon [alerts] either."* P3 were more hopeful, suggesting *"the [haptic and visual alerts for attributes] would be more useful if they were more relevant to the way I was looking through the dataset [instead of comparing with the underlying data distribution as the baseline],"* suggesting alternate baselines to be employed [18]. These sentiments indicate a strong rejection of the mouse's vibrations, putting things in context.

P8 did not like the mouse vibrations but liked the visual icon alerts. They said, *"[the visual icon alerts] affected my data exploration*

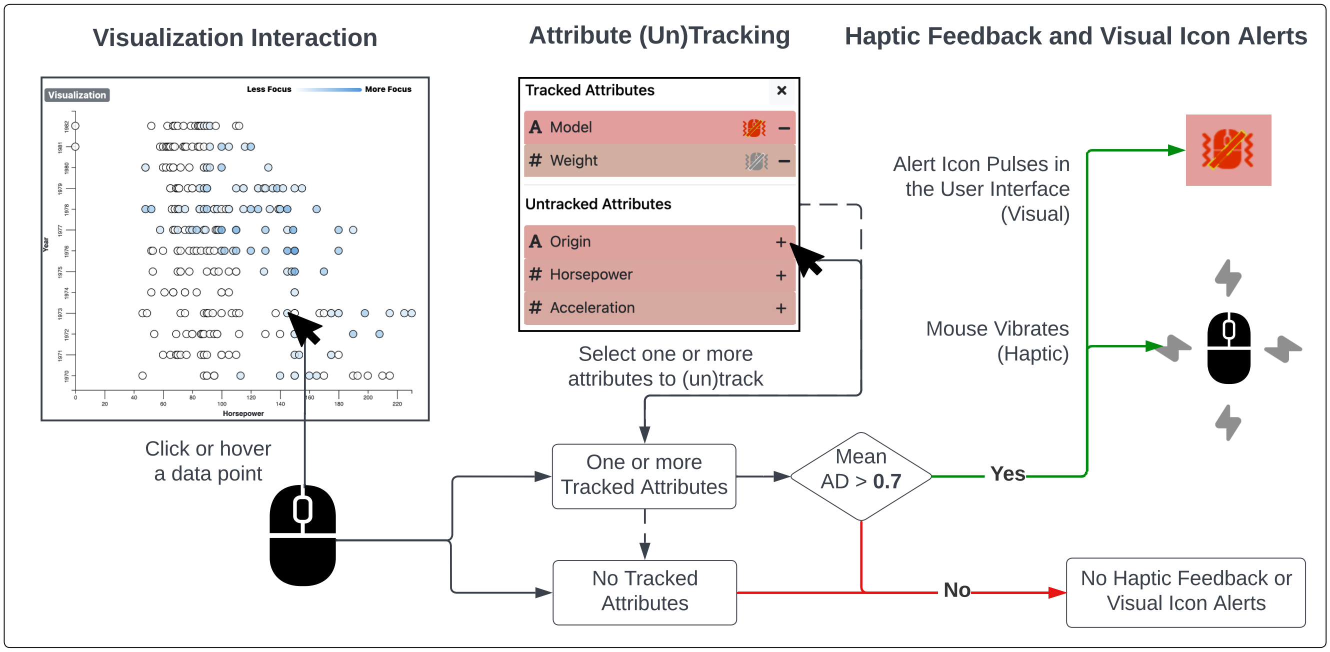

Figure 2: The interaction sequence diagram to trigger arhaptic feedback and visual icon alerts in BiasBuzz. When a user interacts with a datapoint, tracks one or more attributes (for bias mitigation), and if the mean AD metric value for these tracked attributes is greater than a predetermined threshold of 0.7, the mouse vibrates and the corresponding visual alert icons pulse in the user interface. In all other scenarios, there is no haptic feedback or visual guidance.

because it made me want to avoid that data point that it vibrated on." P4 and P7 liked the mouse vibrations but not the visual icon alerts. On the mouse vibrations, P4 said, *"There was one time [the vibrations] went off, and I was like 'ok we need to look at action thriller' and another time I was like 'hey you need to get a drama'."* P7 said, *"When I felt the vibration, I switched my focus to the [Distribution] panel and get some additional information."* On the visual icon alerts, P4 said, *"I feel like because the window was really small I had to scroll to find exactly what attribute was setting it off."* This issue can potentially be mitigated with interface enhancements. P7 said, *"I think the problem I had is that I am not familiar with the [visual icon alert] meaning. I thought the red icon indicates I am making some errors or mistakes, so I am thinking to 'fix' it."* All of these observations demonstrate the wide range of reactions to our enhancements.

4.2 Quantitative Analysis

Tracked Attributes. Participants tracked attributes for a total of 40 times (μ=4.44, η=4, σ=1.07, max=7, min=3). *Genre* was tracked the most (12 times) and *IMDB Rating*, *Worldwide Gross*, and *Running Time* were tracked the least (once). Participants scored (μ=3.55) the ability to track attributes relatively higher than other features like visual icon alerts, haptic mouse vibrations, and the ability to mute attributes. While P2 said, *"[the tracking] is a necessary piece for the whole design,"* P5 said, *"[the tracking] helped me notice the parts where my focus was deviating from expectation."*

Mouse Vibrations. The mouse vibrated a total of 142 times across all nine participants (μ=15.77, σ=8.72, max=36, min=3). Of these, *Genre* was above the exploration bias threshold (= 0.7) the most and vibrated 114 times. *Worldwide Gross*, *Rotten Tomatoes Rating*, *Running Time*, and *IMDB Rating* were all tracked by participants at one point or another, but none of these attributes were above the bias threshold to trigger vibrations. Many participants had interesting things to say about the mouse vibrations. P1 said, *"The only useful part about the haptic feedback is that it reminded me I hadn't reached my goal of selecting and viewing a proportionate amount of different movies with respect to the specific attributes I was tracking."* P4 *"could tell towards the end that the vibration is indicating that you need to work on something."*

Mouse Muteness. Participants muted the vibrations for a total of 56 times (μ=15.77, σ=9.40, max=32, min=0). *Genre* was muted the most (41 times). P8 said *"There are some attributes that I was not considering, so it was great to be able to mute these specific attributes."* P1 muted an attribute only once, and they did this because *"the haptic feedback wasn't too distracting, so I didn't see the need to mute the alert for specific attributes.*

Exploration Bias Mitigation: Did the AD metric values decrease? We plotted the total number of vibrations for each of the three attributes (*Genre*, *Content Rating*, and *Release Year*) against their corresponding final AD values. We observed no correlation suggesting that the vibrations did not reduce the AD values. P3, even though they experienced the most number of vibrations (n=36), said they did not feel they needed the vibrations to do well in the task. They said, *"[Vibration] definitely has a place for some tasks, but I didn't need it all that much for this one."* Similarly, P9 experienced the least number of vibrations (n=3) and did not find the vibrations useful, noting, *"[Vibrations] didn't affect my data exploration process because I was focused on the task of creating a list of 10 movies more than anything."*

Temporal Analysis: Did the vibrations nudge users to respond by interacting differently? Even though there was no overall decrease in AD values, there were instances when participants actually changed their subsequent interaction behavior after the mouse vibrated, either temporarily or permanently. P2 experienced several vibrations due to high AD values for *Genre*, but the AD value continued to drop throughout the session. Notably, P2 also muted *Content Rating* because of which its AD value remained high throughout the task. P4 experienced several vibrations towards the end of their session and the AD value of *Content Rating* dropped in tandem with those vibrations. They said, *"[Vibrations] helped me kinda narrow down genres toward the distribution."* Although the AD value of *Genre* did not drop significantly, their comment suggested that users can be made more aware and reflect on their choices during the task.

5 DISCUSSION AND TAKEAWAYS

5.1 Haptic vibrations can take some time to get used to.

Unlike games, visual data analysis systems often utilize a single visual modality. Thus, it is natural to expect some time before other modalities such as haptics are also accepted. For instance, P4 said they *"didn't notice [the vibrations] at first, but after some repetition, got used to looking at the distributions after feeling the vibrations."* P9 *"wondered if the mouse vibrating was a technical issue."* P8 *"found [the mouse vibrations] were very clear but [were] just not sure why the mouse was vibrating."* P7 even said, *"If you asked me to do it again (the task with the mouse), I could get more used to it."*

5.2 Haptic vibrations can be a positive stimulant to aid analysis.

Many participants were positively stimulated in one way or another directly after the mouse vibrated, lending credibility to the practicality of offering haptic modality as a more "aggressive", "active" form of guidance in visual data analysis. For example, the vibrations acted as a reminder of the analysis goal (P1), realization of missed out data points (P5) and attributes (P7), all of which nudged them to change subsequent interaction strategy.

5.3 Haptic vibrations can also be a source of distraction during analysis.

While the mouse achieved the desired effect of increased analytic awareness for some participants, there were multiple instances where it negatively affected the participant's analytic goals, which is undesirable. P7 said, *"It encouraged me to explore different movie attributes instead of the ones I am interested in."* P8 said, *"It affected my data exploration because it made me want to avoid that data point that it vibrated on. Am I supposed to avoid this data point?"*

6 LIMITATIONS, FUTURE WORK, AND CONCLUSION

The capabilities of the SteelSeries 710 mouse limited this study. As a common off-the-shelf mouse, its vibration intensity was not very high and due to its cooldown requirement, it could not vibrate for long time periods. As a result, even though this mouse supported different vibration patterns, we could not exploit this to different aspects of the user interface (e.g., unique pattern per attribute or a certain level of bias). Studying these via a custom-built mouse that is capable of stronger vibrations, more vibrations in rapid succession, and different types of vibrations, is future work. Furthermore, exploring additional modalities such as natural language [20], ambient display media (light, airflow, sound) [14], or squeeze-haptics [2] to communicate appropriate guidance is also future work.

In conclusion, we investigate how combining visual guidance with haptic feedback can help increase user awareness of and mitigate exploration biases during visual data analysis. We wired a gaming mouse to an existing visual data analysis tool. We enhanced this system to vibrate and reinforce its existing ability to detect and visually communicate exploration biases exhibited by the user. A formative study with nine users revealed that the dual guidance modality of visual and haptic feedback can sometimes increase user awareness of (biased) analytic behaviors but the mouse vibrations can also be distracting and disturbing, putting into context the design of future multimodal guidance-enriched user interfaces for visual data analysis.

ACKNOWLEDGMENTS

This work was supported by NSF grant IIS-1813281 and a Georgia Tech's President's Undergraduate Research Award. We thank our users and the Georgia Tech Visualization Lab for their feedback.

REFERENCES

[1] Motoyuki Akamatsu, Sigeru Sato, and I. Scott MacKenzie. 1994. Multimodal mouse: A mouse-type device with tactile and force display. *Presence: Teleoperators and Virtual Environments* 3, 1 (1994), 73–80. https://doi.org/10.1162/pres.1994.3.1.73

[2] Matthew A Baumann, Karon E MacLean, Thomas W Hazelton, and Ashley McKay. 2010. Emulating human attention-getting practices with wearable haptics. In *2010 IEEE Haptics Symposium*. IEEE, 149–156. https://doi.org/10.1109/HAPTIC.2010.5444662

[3] Leonardo Bonanni and Cati Vaucelle. 2006. A framework for haptic psychotherapy. *depression and anxiety* 12 (2006), 24.

[4] Richard E Boyatzis. 1998. *Transforming Qualitative Information: Thematic Analysis and Code Development.* Sage Publications.

[5] Gwénaël Changeon, Delphine Graeff, Margarita Anastassova, and José Lozada. 2012. Tactile emotions: A vibrotactile tactile gamepad for transmitting emotional messages to children with autism. In *Haptics: Perception, Devices, Mobility, and Communication: International Conference, EuroHaptics 2012, Tampere, Finland, June 13-15, 2012. Proceedings, Part I.* Springer, 79–90.

[6] Inga Chira, Michael Adams, and Barry Thornton. 2008. Behavioral Bias Within The Decision Making Process. *Journal of Business & Economics Research* 6, 8 (2008), 11–20. https://doi.org/10.19030/jber.v6i8.2456

[7] Isaac Cho, Ryan Wesslen, Alireza Karduni, Sashank Santhanam, Samira Shaikh, and Wenwen Dou. 2017. The anchoring effect in decision-making with visual analytics. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 116–126. https://doi.org/10.1109/VAST.2017.8585665

[8] Evanthia Dimara, Giles Bailey, Anastasia Bezerianos, and Steven Franconeri. 2018. Mitigating the attraction effect with visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 850–860. https://doi.org/10.1109/TVCG.2018.2865233

[9] Youngwook Do, Linh Thi Hoang, Jung Wook Park, Gregory Abowd, and Sauvik Das. 2021. Spidey Sense: Designing Wrist-Mounted Affective Haptics for Communicating Cybersecurity Warnings. In *Designing Interactive Systems Conference*

2021. ACM. https://doi.org/10.1145/3461778.3462027

[10] Mi Feng, Cheng Deng, Evan M Peck, and Lane Harrison. 2016. Hindsight: Encouraging exploration through direct encoding of personal interaction history. *IEEE transactions on visualization and computer graphics* 23, 1 (2016), 351–360.

[11] Sunwoo Ha, Shayan Monadjemi, Roman Garnett, and Alvitta Ottley. 2022. A unified comparison of user modeling techniques for predicting data interaction and detecting exploration bias. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 483–492. https://doi.org/10.48550/arXiv.2208.05021

[12] Wenkai Han and Hans-Jorg Schulz. 2020. Exploring vibrotactile cues for interactive guidance in data visualization. In *2020 International Symposium on Visual Information Communication and Interaction (VINCI)*. ACM, 1–10. https://doi.org/10.1145/3430036.3430042

[13] Vincent Hayward, Oliver R Astley, Manuel Cruz-Hernandez, Danny Grant, and Gabriel Robles-De-La-Torre. 2004. Haptic interfaces and devices. *Sensor review* 24, 1 (2004), 16–29. https://doi.org/10.1108/02602280410515770

[14] Hiroshi Ishii and Brygg Ullmer. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 234–241. https://doi.org/10.1145/258549.258715

[15] Ki-Uk Kyung, Heejin Choi, Kwon Dong-Soo, and Son Seung-Woo. 2004. Interactive mouse systems providing haptic feedback during the exploration in virtual environment. In *2004 International Symposium on Computer and Information Sciences*. ACM, 136–146.

[16] Qianyu Liu, Haoran Jiang, Zihao Pan, Qiushi Han, Zhenhui Peng, and Quan Li. 2024. BiasEye: A Bias-Aware Real-time Interactive Material Screening System for Impartial Candidate Assessment. *arXiv preprint arXiv:2402.09148* (2024). https://doi.org/10.48550/arXiv.2402.09148

[17] Arpit Narechania, Surya Chakraborty, Shivam Agarwal, Atanu R Sinha, Ryan A Rossi, Fan Du, Jane Hoffswell, Shunan Guo, Eunyee Koh, Alex Endert, et al. 2023. DataCockpit: A Toolkit for Data Lake Navigation and Monitoring Utilizing Quality and Usage Information. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 5305–5310. https://doi.org/10.1109/BigData59044.2023.10386133

[18] Arpit Narechania, Adam Coscia, Emily Wall, and Alex Endert. 2021. Lumos: Increasing awareness of analytic behavior during visual data analysis. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 1009–1018. https://doi.org/10.1109/TVCG.2021.3114827

[19] Arpit Narechania, Fan Du, Atanu R Sinha, Ryan Rossi, Jane Hoffswell, Shunan Guo, Eunyee Koh, Shamkant B Navathe, and Alex Endert. 2023. Datapilot: Utilizing quality and usage information for subset selection during visual data preparation. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18. https://doi.org/10.1145/3544548.3581509

[20] Arpit Narechania, Adam Fourney, Bongshin Lee, and Gonzalo Ramos. 2021. DIY: Assessing the correctness of natural language to sql systems. In *26th International Conference on Intelligent User Interfaces*. 597–607. https://doi.org/10.1145/3397481.3450667

[21] Marianna Obrist, Sriram Subramanian, Elia Gatti, Benjamin Long, and Thomas Carter. 2015. Emotions mediated through mid-air haptics. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 2053–2062.

[22] Simone Ooms, Minha Lee, Pablo Ceasar, and Abdallah Ali. 2023. FeelTheNews: Augmenting Affective Perceptions of News Videos with Thermal and Vibrotactile Stimulation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM. https://doi.org/10.1145/3544549.3585638

[23] Sabrina Paneels and Jonathan C Roberts. 2009. Review of designs for haptic data visualization. *IEEE Transactions on Haptics* 3, 2 (2009), 119–137. https://doi.org/10.1109/TOH.2009.44

[24] Bahador Saket, Chrisnawan Prasojo, Yongfeng Huang, and Shengdong Zhao. 2013. Designing an effective vibration-based notification interface for mobile phones. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 149–1504. https://doi.org/10.1145/2441776.2441946

[25] Stefano Scheggi, A. Talarico, and Domenico Prattichizzo. 2014. A remote guidance system for blind and visually impaired people via vibrotactile haptic feedback. In *22nd Mediterranean Conference on Control and Automation*. IEEE. https://doi.org/10.1109/MED.2014.6961320

[26] SteelSeries. 2015. SteelSeries Gamesense SDK. https://github.com/SteelSeries/gamesense-sdk

[27] J.A. Terry and H. S. Hsiao. 1988. Tactile feedback in a computer mouse. In *14th Annual Northeast Bioengineering Conference*. IEEE. https://doi.org/10.1109/NEBC.1988.19369

[28] Emily Wall, Leslie M Blaha, Lyndsey Franklin, and Alex Endert. 2017. Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *2017 IEEE conference on visual analytics science and technology (VAST)*. IEEE, 104–115. https://doi.org/10.1109/VAST.2017.8585669

[29] Emily Wall, Leslie M Blaha, Celeste Lyn Paul, Kristin Cook, and Alex Endert. 2018. Four perspectives on human bias in visual analytics. *Cognitive biases in visualizations* (2018), 29–42.

[30] Emily Wall, Arpit Narechania, Adam Coscia, Jamal Paden, and Alex Endert. 2021. Left, right, and gender: Exploring interaction traces to mitigate human biases. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 966–975. https://doi.org/10.1109/TVCG.2021.3114862

[31] Dangxiao Wang, Kouhei Ohnishi, and Weiliang Xu. 2019. Multimodal haptic display for virtual reality: A survey. *IEEE Transactions on Industrial Electronics* 67, 1 (2019), 610–623. https://doi.org/10.1109/TIE.2019.2920602

[32] Gyeore Yun, Minjae Mun, Jungeun Lee, Dong-Geun Kim, Hong Z Tan, and Seungmoon Choi. 2023. Generating Real-Time, Selective, and Multimodal Haptic Effects from Sound for Gaming Experience Enhancement. In *ACM CHI*. https://doi.org/10.1145/3544548.3580787