Highways and Tunnels: Force Feedback Guidance for Visualisations

Sarah Fjelsted Alrøe[1] , Eve Hoggan[2] and Hans-Jörg Schulz[2]

[1]Department of Digital Design and Information Studies, Aarhus University, Denmark
[2]Department of Computer Science, Aarhus University, Denmark

Abstract

Non-visual methods of user guidance in visualisations are still relatively underexplored. This paper aims to address this, by establishing a foundation for appropriately using haptic force feedback in a pointing device to provide guidance, with a focus on pulling and constraining. To explore these guidance methods, a force feedback enabled mouse was constructed, along with a force feedback enabled data visualisation. A user study was conducted, subjecting the participants to different degrees of pulling and constraining guidance, helping them solve navigation tasks. The study found significant quantitative and qualitative changes in behaviour and experience across conditions. We conclude that these two modes of feedback can be used for directing and prescribing guidance situations, provided they are used with restraint.

CCS Concepts

• **Human-centered computing** → *Empirical studies in visualization;* • **Hardware** → *Tactile and hand-based interfaces;*

1. Introduction

A significant challenge in interactive visualisation is traversing, exploring, and understanding large datasets. In these situations, the dataset may be too large to be presented on a single screen, and areas of interest may be located far apart. To address these problems, various forms of guidance can be applied to help users get an overview of relevant data points, choose an appropriate visualisation, or complete the correct steps to reach a desired goal. Appropriately applied guidance has been demonstrated to have both positive performance and experiential effects [CGM19b].

In interactive visualisations, guidance is used to help users navigate the system and data, and overcome situations in which the usage flow is stalled because the next action is uncertain [CGM19a]. Significant research has been conducted to create novel guidance systems for particular tasks and evaluate their effect. The forms of guidance demonstrated vary from helping navigate networks through pointers in various directions [MSDK12] and finding appropriate zoom levels for multiscale data [LMS*12] to helping pick an appropriate data visualisation [WMA*16] and following a complex analytic workflow [SSL*12]. Sophisticated frameworks have been proposed to design and create appropriate guidance mechanisms for visual-analytic systems [CAA*20, HS23, SCEA23].

Worthy of note, the vast majority of research on guidance for interactive visualisations is exclusively in the visual modality, with only few exceptions like vibrotactile feedback [HS20, MKB05] and robot-assisted feedback [DSD*23]. This stands in contrast to the well-established field of *haptic rendering* – i.e., the haptic representation of material properties – for making data more tangible in volume visualisation in general [PCT*07, AS96, WS11] and in

medical applications in particular [LLCY06, Pal11]. In doing so, these approaches provide additional cues beyond the mere visual feedback to the user, allowing for intuitive and precise view manipulations, selections, and placement of probes even in occluded 3D visualisations [BYK*21, Sec.3.2].

In particular, force feedback has several advantages over other forms of haptic feedback: It can be used to present sophisticated feedback through a large number of design parameters, and even provide feedforward, where it can be used to guide the user's hand. Previous research on multimodal interaction has established that providing auditory or haptic feedback can improve performance through offloading of the visual channel [BPG*06]. It should follow that the same gains may be obtained by extending guidance beyond the visual modality.

This paper investigates the question of how to provide appropriate guidance in interactive visualisation using a force feedback enabled pointing device. We answer this by (1) introducing CASTOR, a computer mouse with software-configurable force feedback, (2) using CASTOR to conduct a user study to evaluate the different parameters for force feedback guidance, (3) reporting on pertinent observations from that user study, and (4) concluding about appropriate forms of force feedback guidance.

2. The CASTOR and its Guidance Feedback

Various force feedback pointing devices have been developed in the past [KHS95, KCKS04]. Common to these mice is that they are anchored to the desk, as a means of applying force. These mice have a limited range of movement and may be experienced signifi-

Figure 1: *The* CASTOR *mouse with its two omni-wheels protruding from the bottom and the BLDC motors inside the shell.*

cantly differently from regular mice. One early example exists of an unanchored mouse [AS94], that provides friction through an electromagnet inside the mouse and a solenoid to prod the user's finger. Since then, some advancements have been made, first by Kudo et al. [KSK07] who created an unanchored 2dof ball mouse with force feedback, and later Kianzad and MacLean [KM18] who created a similarly functioning, miniaturised mechanism in a pen-like device.

For our study, we built a force feedback mouse called CASTOR (see Figure 1), whose design remains as close as possible, in shape and functionality, to a traditional computer mouse and can be operated in the same way. The CASTOR is a 2-button, 2-axis mouse that can provide force feedback on both axes. This is achieved through two wheel assemblies that are mounted at a 90° angle, protruding slightly from the bottom of the mouse. Each wheel assembly consists of an omni-wheel mounted to a BLDC motor. This design was inspired by work on force feedback rotary knobs [vOH20]. The 90° angle lets each wheel encode movement and provide force feedback in their respective axis without affecting each other. A Raspberry Pi Pico controls both motors through two L6234D driver circuits. All materials to reproduce the CASTOR mouse incl. circuit schematics and 3D print files for the casing are available at https://vis-au.github.io/highways/. The mouse can be configured to pull in a given direction, to retain a set position, to simulate a customisable bumpy texture on movement, or to apply friction proportional to its velocity. We use these configurations to guide users by pulling and by constraining.

Pulling actively drags the pointer towards a target. We use the metaphor of a *haptic highway*, because the pulling feedback can lead the user along a track connecting multiple targets and can be traversed at high speeds. The CASTOR is set to provide a maximum force of 1.9N with a transfer function mapping some state of the system (e.g., position or time) to force intensity. All examples given here use either a constant force intensity or one that is determined by the distance between the pointer and the target.

Constraining restricts the pointer from moving into or out of a bounding area. In Dennerlein et al., this takes the shape of "tunnelling", where the pointer is constrained in one axis to improve steering performance [DMH00], hence our metaphor of a *haptic tunnel*, pushing the pointer towards the middle of a "tunnel" between the origin and target. The bounding walls may be rigid, applying full force as the pointer touches them, or be soft, slowly ramping up the force as the pointer moves further outside them.

3. User Study Setup

Ceneda et al. established three degrees of guidance [CGM*17]: Orienting, which informs of possible options; Directing, which suggests a specific option; and Prescribing, which forces the choice of a specific option. We hypothesise that the strength of force feedback will affect the perceived degree of guidance – e.g., a strong pull will feel like prescribing guidance, while a weak pull will feel like directing guidance. We further expect continuousness and directness of feedback to have similar effects, increasing the perceived degree of guidance, yet also causing undesirable strain.

The experiment was conducted using Cartograph's 2D embedding of Wikipedia articles [SSL*17]. The visualisation of the embedding was created in a way that

- maximises the *data-ink ratio* [Tuf01] to not confound the haptic guidance with non-essential, potentially distracting visuals;
- uses a high-contrast encoding with ample whitespace to reduce *simultaneous contrast*, known to cause large errors [KHI*03];
- does not rely on colour, removing the impact of *colour blindness*, which can affect guidance in wayfinding scenarios [LLC20];

An example of the interface is shown in Figure 2. A small arrow pointing towards the target was provided for all studied scenarios to make the case without force feedback guidance a meaningful baseline, while at the same time maintaining comparability between scenarios – i.e., the only aspect changing in the setup was the presence/absence of pulling and/or constraining force feedback.
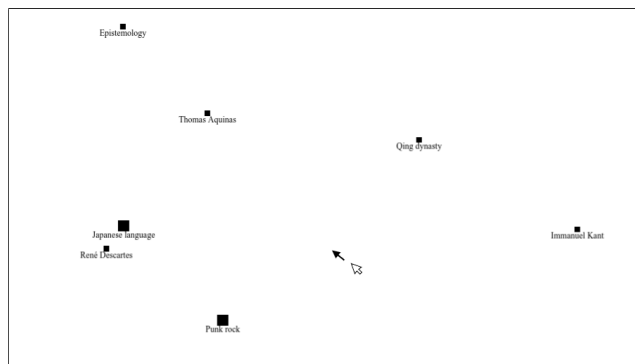
Figure 2: *Example of the interface shown to the study participants.*

18 participants (ages 21-28, avg=24.7; 13 men, 5 women) were given the task of finding a specific data point in the visualisation. The visualisation stretched far beyond the limits of the screen with the sought data point/target position being far off-screen and initially not visible. Starting and target positions were randomised, though with a constant distance (14,000px) between start and targets, and exclusively in cardinal or intercardinal directions. The view is panned using regular click-and-drag navigation, similar to that of mapping software. The force feedback was limited to pull intensity, and constraint rigidity, using changes in intensity for pulling and changes in rigidity for constraining. The constraining and pulling feedback were designed to fit this form of navigation as follows: For the pulling feedback, when the mouse is moving freely, the pointer is pulled directly towards the target. When the pointer is dragging the view, the pointer is instead pulled directly away from the direction of the target, moving the view closer to the target. For the constraining feedback, an invisible line is created, stretching from the origin at the start of the trial, to the target point of the trial. When the mouse is free, the pointer is forced towards the closest point on this line. When dragging, the pointer is instead forced in a direction, that moves the centre of the screen towards the closest point on this line. When pulling and constraining is combined, the two resulting forces of the feedback are simply added together. The specific settings and forces are given in Table 1.

This was a within-subjects study. The participants were first introduced to the task of finding a topic in the visualisation supported by navigational guidance. Then, in a sequence of 9 trials, all possible pairwise combinations of {no, weak, strong} pulling and {no, weak, strong} rigidity for constraining were tested in random order. When the cursor reached the target, a short questionnaire had to be completed, after which the next trial began. At the end, a short semi-structured interview was conducted.

We measured time taken to reach the target, total length of cursor movement, and cursor deviation from the ideal straight line from start to target, as well as user experience through the NASA TLX, three study-specific questions, and the semi-structured interview. The three questions were presented on a 5-step Likert-scale, while the NASA TLX questions were presented in 20 steps. The three questions were:

1. How much did you feel in control vs. the system? I was completely in control - The system was completely in control.
2. How confident did you feel that you were going in the right direction? Very uncertain - Very confident.
3. How much did you focus on getting to the target vs. focus on other elements? Entirely on target - Entirely on other things.

Level	Max parameter	Max force [N]
(0) No	0	0
(1) Weak	1	0.77
(2) Strong	2	1.9

Table 1: *Summary of factor levels and maximum force applied, measured in a single direction. Both pulling and constraining are capable of the same forces for the same level. Pulling and constraining combined may produce stronger forces than measured.*

4. Quantitative Study Results

The completion time, distance, and deviation from ideal line were tested for statistically significant differences across factor levels, signifying a change in behaviour/performance. The three variables were binned by each of the factors to identify which, if not both, of the factors affected behaviour. To establish statistical significance, the measurements were first determined to not have a normal distribution. For this reason, a Friedman test was applied, which showed a statistically significant difference in completion time and distance for pulling feedback, and significant difference in path deviation for constraining feedback. To identify which factor levels created this difference, a Wilcoxon signed-rank test was performed with p-values adjusted using the Bonferroni method.

While there is no statistically significant difference in completion distance or time across levels of constraining feedback, there is a statistically significant decrease in completion time ($p < 0.0001$) and a slight but significant decrease in completion distance ($p = 0.018$) across the levels of pulling feedback. The opposite holds true for path deviation, which only shows a statistically significant decrease ($p < 0.0001$) across levels of constraining feedback. This suggests that the two feedback types affect participant behaviour in different ways. Increasing the strength of pulling feedback increases the pace of the task and may lower the distance travelled. Increasing the strength of constraining feedback decreases the average deviation from the ideal path. It should be noted, that only the strong pulling feedback creates a statistically significant decrease in completion time, suggesting that weaker pulling configurations may not affect behaviour in this way. In contrast, only weak pulling creates a statistically significant decrease in completion distance.

Regarding the NASA TLX questions, the Performance, Effort, and Frustration questions were observed to be difficult to apply to the task, or seemed irrelevant. For this reason, only the Mental, Physical, and Temporal demand questions were examined, along the three experiment-specific questions. As the data is subjective and cannot be expected to follow any specific distribution, a Friedman test and following Wilcoxon test were performed.

There appears to be no statistically significant effect of constraining feedback on the NASA TLX results. Yet, increasing the strength of pulling feedback significantly lowers the task's mental demands ($p = 0.00048$) and physical demands ($p < 0.0001$), while significantly increasing the temporal demands ($p < 0.0001$). In other words, increasing the strength of pulling feedback makes the task generally easier but feels much more rushed.

The results of the control, confidence, and focus questions are more ambiguous than the TLX questions. The question about whether the user or system is in control shows a statistically significant difference for both pulling ($p < 0.0001$) and constraining ($p = 0.037$). This difference appears to correlate between levels of feedback in that increasing the level of pulling or constraining decreases the experienced level of participant control. For the question about confidence, pulling feedback creates a statistically significant difference in responses ($p < 0.0001$), while constraining does not ($p = 0.45$). Increasing the level of pulling results in greater participant confidence in getting to their target. For the question about focus on the target vs. other things, there appears to be no clear correlation between factor levels and responses.

5. Qualitative Study Results

We collated reoccurring experiences of the users as reported in the semi-structured interviews. Experiences that occurred for less than three participants were filtered out, and the remaining experiences were then further grouped into common themes. Here, we report specifically four themes that are also reflected in the quantitative study results.

Theme 1 – Constraining feedback is easy to overlook and may be difficult to understand: Common among all 18 interviews was a much larger focus on the pulling feedback compared to the constraining feedback. While the pulling feedback was readily mentioned and discussed, comments on the constraining feedback had to be prompted for almost every participant. When directly asked, six of the participants either did not at all or barely noticed this type of feedback, while three who did notice, did not regard the feedback as helpful. This lack of awareness may be explained by the fact that the pulling force is constant, whereas the constraining force gradually increases and decreases based on the pointer position. This is supported by the quantitative results as well as the movement data, where there, with few exceptions, were no significant differences between constraining feedback configurations.

Theme 2 – Pulling feedback is generally appreciated: When discussing the pulling feedback, most participants shared a positive impression in both utility and experiential regards. Adding the pulling feedback reduced the effort expended to solve the task: 11 of the participants reported that the addition of pulling feedback lowered the physical effort required to solve the task, while four of them further reported that it lowered the mental effort of finding the target. This is supported by the quantitative results, which show a large, statistically significant difference in physical and mental effort between each of the pulling feedback configuration.

Theme 3 – BUT only when not too strong: While the general reaction to the pulling feedback was positive, this was not the case for strong pulling specifically. Even though a few participants enjoyed the increased task efficiency provided by the strong feedback, most considered it as "too much". Common among the participants (13) was the sentiment that strong pulling feedback made them feel like they lost control, resulting in a loss of personal responsibility (7), a loss of attention (2), and a growing complacency towards the task (11). These experiences align with the concept of *automation complacency*, which posits that users exclusively monitoring a constantly reliable automated process grow complacent to their task [PMS93]. For other participants, the strong pulling instead resulted in discomfort and a loss of meaning. Five of the participants reported that the loss of control made them feel redundant or that their presence was meaningless. Multiple participants explicitly stated that they preferred the weaker pulling, in that it did not feel as controlling as the strong pulling (5), or was just helping instead of controlling (4).The differences in experience between weak and strong pulling are supported by the quantitative results, showing a significantly increased completion speed, temporal demand, and loss of control for the strong pulling feedback.

Theme 4 – Constraining and pulling feedback creates different forms of control: For the participants who managed to recognise and understand the constraining feedback, their conceptualisation of control for the two feedback types were quite different. Nine of the participants stated that they felt significantly more rushed by the pulling feedback, while the constraining feedback felt calmer. Two participants clearly reported about this distinction: Constraining feedback would have temporal freedom but positional control, while pulling feedback would have some positional freedom but temporal control. This is supported by the quantitative results showing loss of participant control for both factors, but a greater effect size for the pulling factor.

6. Discussion and Conclusions

Based on the results of the user study, answers to the question of appropriate force feedback can be approached: Using the pulling type feedback, it was possible to accomplish a spectrum of guidance ranging from directing to prescribing approaching full automation, though with significant concerns and caveats. Weak pulling is capable of producing an appropriate directing guidance towards a target without seeming to affect behaviour in any significant way and with minimal negative effects to experience. By applying strong pulling feedback, it is possible to create a strongly prescribing guidance, exerting more control over the user than would seemingly be possible through visual guidance. This however comes with a strong trade-off in that it creates significant negative experiences of being rushed, loss of control, and/or complacency with the system. In the form demonstrated in the user study, strong pulling is not appropriate for guidance. Using this form of feedback more sparingly, in short bursts, or through user activation may alleviate these issues.

For the constraining feedback, the question of appropriateness is difficult to evaluate. Its appropriateness is likely, since study participants generally found it inoffensive. Similar to the pulling feedback, however, too strong constraining feedback can result in negative experiences of loss of control, and it should therefore be applied with care. It should be noted that this mapping of constraining feedback to click-and-drag navigation is quite complex compared to the simple bounding box of a static UI. This may create particularly great differences in experience compared to other usage scenarios like traditional UI navigation. Further exploration of this may however prove fruitful as a way to increase guidance degree without incurring the negative aspects of strong pulling.

Finally, the user study revealed how force feedback guidance must make new distinctions of how control is exerted by the system. While traditional visual guidance is only capable of establishing sequential control over the user – e.g., by suggesting or enforcing a specific sequence of steps – the two force feedback methods exhibit a much more direct control over the user. As explained under Theme 5, the pulling feedback was capable of exhibiting a form of temporal control, suggesting or enforcing the pace at which the sequence would be carried out. In contrast, the positional control, as primarily exhibited by the constraining feedback, suggests or enforces where the pointer should move to carry out the sequence. The fact that these forms of control directly affect the user's physical abilities may help to explain the negative effects of increased control. Common to both feedback types is a concern that otherwise appropriate feedback may turn highly inappropriate the moment the goals of the system and of the user are no longer aligned. While this is a known problem in existing visual guidance theory, it becomes especially pertinent for these new forms of control.

References

[AS94] AKAMATSU M., SATO S.: A multi-modal mouse with tactile and force feedback. *International Journal of Human-Computer Studies 40*, 3 (1994), 443–453. doi:10.1006/ijhc.1994.1020. 2

[AS96] AVILA R. S., SOBIERAJSKI L. M.: A haptic interaction method for volume visualization. In *Proc. of IEEE Visualization* (1996), IEEE, pp. 197–204. doi:10.1109/VISUAL.1996.568108. 1

[BPG*06] BURKE J. L., PREWETT M. S., GRAY A. A., YANG L., STILSON F. R. B., COOVERT M. D., ELLIOT L. R., REDDEN E.: Comparing the effects of visual-auditory and visual-tactile feedback on user performance: a meta-analysis. In *Proc. of the 8th International Conference on Multimodal interfaces (ICIM'06)* (2006), ACM, pp. 108–117. doi:10.1145/1180995.1181017. 1

[BYK*21] BESANÇON L., YNNERMAN A., KEEFE D. F., YU L., ISENBERG T.: The state of the art of spatial interfaces for 3D visualization. *Computer Graphics Forum 40*, 1 (2021), 293–326. doi:10.1111/cgf.14189. 1

[CAA*20] CENEDA D., ANDRIENKO N., ANDRIENKO G., GSCHWANDTNER T., MIKSCH S., PICCOLOTTO N., SCHRECK T., STREIT M., SUSCHNIGG J., TOMINSKI C.: Guide me in analysis: A framework for guidance designers. *Computer Graphics Forum 39*, 6 (2020), 269–288. doi:10.1111/cgf.14017. 1

[CGM*17] CENEDA D., GSCHWANDTNER T., MAY T., MIKSCH S., SCHULZ H.-J., STREIT M., TOMINSKI C.: Characterizing guidance in visual analytics. *IEEE TVCG 23*, 1 (2017), 111–120. doi:10.1109/TVCG.2016.2598468. 2

[CGM19a] CENEDA D., GSCHWANDTNER T., MIKSCH S.: A review of guidance approaches in visual data analysis: A multifocal perspective. *Computer Graphics Forum 38*, 3 (2019), 861–879. doi:10.1111/cgf.13730. 1

[CGM19b] CENEDA D., GSCHWANDTNER T., MIKSCH S.: You get by with a little help: The effects of variable guidance degrees on performance and mental state. *Visual Informatics 3*, 4 (2019), 177–191. doi:10.1016/j.visinf.2019.10.005. 1

[DMH00] DENNERLEIN J. T., MARTIN D. B., HASSER C.: Force-feedback improves performance for steering and combined steering-targeting tasks. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems* (2000), ACM, pp. 423–429. doi:10.1145/332040.332469. 2

[DSD*23] DAI S., SMILEY J., DWYER T., ENS B., BESANÇON L.: RoboHapalytics: A robot assisted haptic controller for immersive analytics. *IEEE TVCG 29*, 1 (2023), 451–461. doi:10.1109/TVCG.2022.3209433. 1

[HS20] HAN W., SCHULZ H.-J.: Exploring vibrotactile cues for interactive guidance in data visualization. In *Proc. of the 13th International Symposium on Visual Information Communication and Interaction (VINCI'20)* (2020), ACM, pp. 14:1–14:10. doi:10.1145/3430036.3430042. 1

[HS23] HAN W., SCHULZ H.-J.: Providing visual analytics guidance through decision support. *Information Visualization 22*, 2 (2023), 140–165. doi:10.1177/14738716221147289. 1

[KCKS04] KYUNG K.-U., CHOI H., KWON D.-S., SON S.-W.: Interactive mouse systems providing haptic feedback during the exploration in virtual environment. In *Proc. of the International Symposium on Computer and Information Sciences (ISCIS'04)* (2004), Springer, pp. 136–146. doi:10.1007/978-3-540-30182-0_15. 1

[KHI*03] KOSARA R., HEALEY C. G., INTERRANTE V., LAIDLAW D. H., WARE C.: Thoughts on user studies: Why, how, and when. *IEEE Computer Graphics and Applications 23*, 4 (2003), 20–25. doi:10.1109/MCG.2003.1210860. 2

[KHS95] KELLEY A. J., HIGUCHI T., SALCUDEAN S. E.: An integrated haptographical user interface using a force feedback mouse. In *Proc. of the Intl. Conference on Human-Computer Interaction* (1995), Elsevier, pp. 713–718. doi:10.1016/S0921-2647(06)80299-6. 1

[KM18] KIANZAD S., MACLEAN K. E.: Harold's purple crayon rendered in haptics: Large-stroke, handheld ballpoint force feedback. In *Proc. of the IEEE Haptics Symposium (HAPTICS'18)* (2018), IEEE, pp. 106–111. doi:10.1109/HAPTICS.2018.8357161. 2

[KSK07] KUDO H., SUZUMORI K., KANDA T.: Force feedback mouse with differential mechanism for omni-traveling. In *Proc. of the Intl. Symposium on Micro-Nano Mechatronics and Human Science* (2007), IEEE, pp. 430–435. doi:10.1109/MHS.2007.4420894. 2

[LLC20] LEE H., LEE E., CHOI G.: Wayfinding signage for people with color blindness. *Journal of Interior Design 45*, 2 (2020), 35–54. doi:10.1111/joid.12169. 2

[LLCY06] LUNDIN K., LUNDSTRÖM C., COOPER M., YNNERMAN A.: Enabling haptic interaction with volumetric MRI data through knowledge-based tissue separation. In *Proc. of Volume Graphics* (2006), Eurographics Association. doi:10.2312/VG/VG06/075-078. 1

[LMS*12] LUBOSCHIK M., MAUS C., SCHULZ H.-J., SCHUMANN H., UHRMACHER A. M.: Heterogeneity-based guidance for exploring multiscale data in systems biology. In *Proc. of the IEEE Symposium on Biological Data Visualization BioVis* (2012), Roerdink J., Hibbs M., (Eds.), IEEE, pp. 33–40. doi:10.1109/BioVis.2012.6378590. 1

[MKB05] MCGOOKIN D. K., KILDAL J., BREWSTER S. A.: New views on haptic graph visualisation. In *Hands on Haptics: Exploring Non-Visual Visualisation Using the Sense of Touch (CHI'05)* (2005). 1

[MSDK12] MAY T., STEIGER M., DAVEY J., KOHLHAMMER J.: Using signposts for navigation in large graphs. *Computer Graphics Forum 31*, 3pt2 (2012), 985–994. doi:10.1111/j.1467-8659.2012.03091.x. 1

[Pal11] PALMERIUS K. L.: Combining vibrotactile and kinaesthetic cues in haptic volume visualization. In *Proc. of the 2011 IEEE World Haptics Conference* (2011), IEEE, pp. 511–515. doi:10.1109/WHC.2011.5945538. 1

[PCT*07] PERSSON P. B., COOPER M. D., TIBELL L. A., AINSWORTH S., YNNERMAN A., JONSSON B.-H.: Designing and evaluating a haptic system for biomolecular education. In *Proc. of the 2007 IEEE Virtual Reality Conference* (2007), IEEE, pp. 171–178. doi:10.1109/VR.2007.352478. 1

[PMS93] PARASURAMAN R., MOLLOY R., SINGH I. L.: Performance consequences of automation-induced 'complacency'. *The International Journal of Aviation Psychology 3*, 1 (1993), 1–23. doi:10.1207/s15327108ijap0301_1. 4

[SCEA23] SPERRLE F., CENEDA D., EL-ASSADY M.: Lotse: A practical framework for guidance in visual analytics. *IEEE TVCG 29*, 1 (2023), 1124–1134. doi:10.1109/TVCG.2022.3209393. 1

[SSL*12] STREIT M., SCHULZ H.-J., LEX A., SCHMALSTIEG D., SCHUMANN H.: Model-driven design for the visual analysis of heterogeneous data. *IEEE TVCG 18*, 6 (2012), 998–1010. doi:10.1109/TVCG.2011.108. 1

[SSL*17] SEN S., SWOAP A. B., LI Q., BOATMAN B., DIPPENAAR I., GOLD R., NGO M., PUJOL S., JACKSON B., HECHT B.: Cartograph: Unlocking spatial visualization through semantic enhancement. In *Proc. of the Intl. Conference on Intelligent User Interfaces (IUI'17)* (2017), ACM, pp. 179–190. doi:10.1145/3025171.3025233. 2

[Tuf01] TUFTE E.: *The Visual Display of Quantitative Information*, 2nd ed. Graphics Press, 2001. 2

[vOH20] VAN OOSTERHOUT A., HOGGAN E.: Reshaping interaction with rotary knobs: Combining form, feel and function. In *Proc. of the ACM Designing Interactive Systems Conference (DIS'20)* (2020), ACM, pp. 1973–1982. doi:10.1145/3357236.3395536. 2

[WMA*16] WONGSUPHASAWAT K., MORITZ D., ANAND A., MACKINLAY J., HOWE B., HEER J.: Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE TVCG 22*, 1 (2016), 649–658. doi:10.1109/TVCG.2015.2467191. 1

[WS11] WEI L., SOURIN A.: Function-based approach to mixed haptic effects rendering. *The Visual Computer 27*, 4 (2011), 321–332. doi:10.1007/s00371-011-0548-0. 1