

RoboHapalytics: A Robot Assisted Haptic Controller for Immersive Analytics

Shaozhang Dai, Jim Smiley, Tim Dwyer, Barrett Ens, and Lonni Besancon

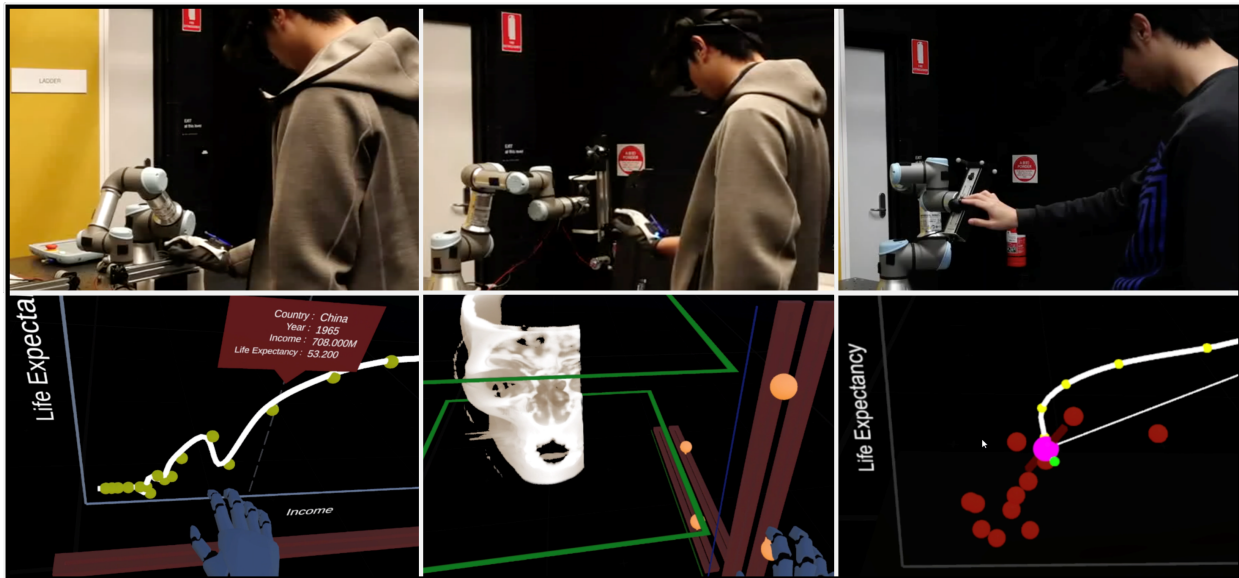


Fig. 1. Use Cases of the dynamic slider system: a time-series chart selector (left top and left bottom); a 3D medical image sectioning tool; and a time-varying scatter plot (middle top and middle bottom) (right top and right bottom)

Abstract— Immersive environments offer new possibilities for exploring three-dimensional volumetric or abstract data. However, typical mid-air interaction offers little guidance to the user in interacting with the resulting visuals. Previous work has explored the use of haptic controls to give users tangible affordances for interacting with the data, but these controls have either: been limited in their range and resolution; were spatially fixed; or required users to manually align them with the data space. We explore the use of a robot arm with hand tracking to align tangible controls under the user's fingers as they reach out to interact with data affordances. We begin with a study evaluating the effectiveness of a robot-extended slider control compared to a large fixed physical slider and a purely virtual mid-air slider. We find that the robot slider has similar accuracy to the physical slider but is significantly more accurate than mid-air interaction. Further, the robot slider can be arbitrarily reoriented, opening up many new possibilities for tangible haptic interaction with immersive visualisations. We demonstrate these possibilities through three use-cases: selection in a time-series chart; interactive slicing of CT scans; and finally exploration of a scatter plot depicting time-varying socio-economic data.

Index Terms—Haptic Feedback, Human Centred Interaction, Robotic Arm

1 INTRODUCTION

In 1965, at the dawn of interactive computer graphics technology, Ivan Sutherland remarked that “*If the task of the display is to serve as a looking-glass into the mathematical wonderland constructed in computer memory, it should serve as many senses as possible*” [65]. He went on to speculate that in “The Ultimate Display” the computer would be able to control the existence of matter. While that precise

technology remains science fiction, virtual reality coupled with active tangible devices allows us to begin to simulate the look and feel of computer-controlled matter. There have been various systems proposed to provide immersive data visualisations with various levels of haptic feedback (see Sect. 2). However, in this paper we present a novel approach to immersive haptic data interaction which uses a robot arm to automate the position of tangible controls (an actuated slider and a rotary encoder), and offers interaction affordances congruent to the common navigation and selection tasks of many visualisation applications. To our knowledge we are the first to use a desktop robot arm for tangible interaction and “encountered haptics” (Sect. 2.3) in an immersive analytics context. The robot allows us to:

- Shaozhang Dai is with Monash University. E-mail: shaozhang.dai@monash.edu
- Jim Smiley is with Monash University. E-mail: jim.smiley@monash.edu
- Tim Dwyer is with Monash University. E-mail: tim.dwyer@monash.edu
- Barrett Ens is with Monash University. E-mail: barret.ens@monash.edu
- Lonni besancon is with Linköping University. E-mail: lonni.besancon@gmail.com

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.xxx/TVCG.201x.xxxxxxx

- position and orient physical controls for direct interaction with data with six Degrees of Freedom (6 DoF);
- extend the range of a small physical slider control to the full reach of the robot, i.e. where past work exploring slider controls for immersive analytics had a 10cm range (Sect. 2), ours has an effective range of up to 76cm (Sect. 3);
- change the orientation of the slider to control any of the principal orthogonal (x, y, z) dimensions, or any other orientation (e.g., to align with the axes of a rotated model rather than world coordinates); and

- simulate a non-linear slider control with a linear physical slider (Sect. 7.3).

The other key technology that enables our system is precise hand-position tracking. Hand-position tracking allows us to:

- place the physical controller at the location of a VR affordance just-in-time as the user reaches for it (e.g., to place the slider knob at the position of a data point or at the filter control of an axis);
- provide physical feedback for a multiplicity of virtual affordances using the motor of the actuated slider for precise haptics; and
- get the physical controller out of the way when not required for hands-free mid-air interaction.

For immersive analytics, this means that we can:

- provide tangible interaction affordances for data elements anywhere within the 3D volume of reach of the robot;
- allow users to perform slider selections and filtering operations along arbitrary (even non-linear) axes.

There are questions about how effective our robot-assisted solution is in achieving the above criteria. In particular, does using the robot to extend the range of a small slider actually provide a selection device that is as effective as a large physical slider? Furthermore, how do these tangible slider solutions compare to a more standard mid-air selection gesture? To address these questions we report on a comparative study evaluating three slider conditions: virtual, physical, and dynamic (robot-assisted). We find that the robot-assisted slider approaches the physical slider in terms of speed and provides a similar accuracy, while both tangible controls are significantly faster and more accurate than mid-air interaction, as reported in Sect. 4.

Thus, while the robot-assisted slider has similar capability to a large physical slider in terms of speed and accuracy, it enables the other possibilities, mentioned above, of being arbitrarily movable and reorientable while the physical slider is too large and unwieldy to allow this. We further demonstrate the capabilities offered by integration of the robot controlled slider device with hand-position tracking through exploration of three data visualisation use-cases:

- A time-series chart selector, demonstrating one-to-one scale direct interaction with a time-axis slider to read off specific data points (Sect. 7.1);
- A 3D medical image sectioning tool, demonstrating hand-tracking and one-to-one scale range filtering across three orthogonal axes and integrating a second tangible control mounted on the device to perform rotations of the model (Sect. 7.2);
- A VR implementation of the DimpVis technique [40] for direct interaction with time-varying scatter plots (Sect. 7.3).

2 RELATED WORK

The presented work pertains to immersive analytics [13, 50], human-robot interactions in virtual reality, as well as haptic controllers and tangible interactions in the context of visualisation and data analysis.

2.1 Tangible Interaction for visualisation

Tangible interaction aims to leverage natural human interaction with everyday 3D objects to facilitate interacting with digital content [24, 34]. For 3D manipulations, tangible interaction has been shown in past research to be fast and precise [8], entertaining [8, 75], as well as fostering and assisting collaborations [51, 53]. Researchers have thus tried to develop tangible devices to assist visual exploration of complex and usually spatial data [4, 12]. A pioneering example is from Hinckley et al. [33] who used tracked props for neurosurgeons that allowed them to manipulate cutting planes and data in order to explore the internal structure of their datasets. While some research prototypes make use of generic tangible devices, to increase versatility, other research projects have investigated specific tangible props. Such devices range from a 3D gun to explore DNA data [58], a paper roll to select thin fiber structures [36], or balls to interact with geo-data [56]; to pen-like props for 3D brushing [28], selection of regions of interest [19] or manipulate cutting planes [35].

Other research projects have instead studied the use of existing, ‘generic’ devices (some with built-in tracking solutions) to increase their versatility. Tablet-like devices have been used for this purpose,

whether self-tracked or relying on external tracking. Cassinelly and Matasoshi [15] used a screen combined with external tracking to allow experts to explore and annotate medical data. Similarly, Spindler et al. [63, 64] used tracked screen-like props to augment an existing visual representation with colours or to switch between multiple levels of abstraction. Pushing the concept further, some researchers proposed to use mobile devices which, combined with their tactile screens, provide more versatility for data exploration or selection through a combination of touch and tangible interaction [7, 11, 47, 60, 62]. Other approaches have also relied on commercial devices such as the 3D mouse [73].

Particularly related to our work and on the spectrum between more specialised and generic devices, is the work from Cordeil et al. [17] investigating a novel approach to controller design with the concept of *spatio-data coordination* [18]. In their work, they used an actuated slider which they mounted on three orthogonal axes to provide a volumetric range selector. This work was recently extended by Smiley et al. [61] who decoupled the three axes from Cordeil et al. [17] to provide universal and versatile axes controls that are handheld and can be combined. In their work, the slider constrains interaction to a single axis. Such constrained interactions with tangibles have been shown in the past to be beneficial to data manipulation [31]. While Smiley et al. [61] used their axes for collaborative immersive analytics and by proposing to explore multiple axes, we instead focus on using a single axis for a single user and explore how it can be used to provide immersive haptic data interaction facilitated by a robot arm to automatically position the axis under the user’s hand.

2.2 Haptic and actuated devices for visualisation

Our concept relies on an actuated slider to provide users with a versatile tangible device for visualisation tasks. Numerous research projects have previously investigated actuated devices in the context of visualisation. One category of such devices lies in actuated physicalisation of datasets. While most of the early physical visualisation prototypes were static, more recent prototypes have been developed to be dynamic and potentially interactive. The majority of these systems function with arrays of motorised bars [26, 45, 54, 66], but some also propose a more fluid control of the final physical shape [25, 55]. In these actuated systems, the physical coordinates represent the digital dataset and their actuation allows users to visualise different datasets or to, for instance, experience how a dataset changes over time [26, 44, 66]. Such actuated systems usually rely on expensive and specific hardware and only rarely allow for interaction (see e.g., [66]). To provide a more versatile control of the final geometry and facilitate richer interactions, Le Goc et al. developed Zooids [42]: small robots that, thanks to external tracking, rearrange themselves on a table to provide several visual representations and adapt to different datasets. They also serve as versatile controllers to interact with the data they display. It should be noted that compared to the robot system we present here, Zooids only operate on a 2D surface and are far too small to provide force feedback.

Our work takes more inspiration from the possibilities offered by other haptic and actuated systems provided by some commercial devices. One such device is the PHANTOM and other haptic pens that have been extensively used in visualisation applications. They have been used, for instance, to allow users to feel and explore through haptics scalar and vector fields [48, 49, 69] or to simulate medical examinations [68]. However, those devices and associated research projects mostly focus on providing haptic forces to allow for exploration and help users feel with only a pen stylus as a tangible tool. In contrast, our approach is particularly inspired by the work of Lischke et al. [46] who used a motorised slide potentiometer as an interactive device for virtual environments. Our prototype, however, investigates such a solution for visualisation tasks and proposes to dynamically adjust the position of the device to follow the user’s hands to provide seamless data interaction mechanisms in immersive contexts.

2.3 Encountered-type Haptics

In the field of virtual reality, researchers have tried to introduce realistic haptic feedback through different approaches. Particularly related to

our work is the idea of “encountered-type haptics” [3, 29, 39, 72]. The idea is simple: an “active” device is used to provide different contact points under various situations. In other words, the device will adjust its position and orientation to match the user’s actions in order to provide them with realistic haptic feedback in a variety of tactile situations. The idea to use robotic arms as active interaction devices has received a lot of attention in the literature through the study of human-robot interactions [30]. Many studies have involved human-robot interactions in the VR field as well. For example, Lee et al. [43] developed a remote control approach to manipulate the robot end effector movement with hand gestures through virtual reality controllers. Robot arms are also widely used in encountered-type haptics to continuously follow and match the user’s actions [3, 22, 52, 72]. For example, Araujo et al. [3] used a robot arm to provide feeling of textures through a tactile device on a robot arm and simple interactions through buttons. Mercado et al. [52] used a robot arm to provide an on-demand tangible surface by placing a physical plane under the contact area when the user is putting an object on a surface in the virtual world. Such approaches have been evaluated and the results seem to indicate that they enhance immersive experiences as they provide real-time “being there” sensation when interacting with virtual objects [3, 29, 52].

While other devices such as drones [1, 76] or other custom robots [29, 67] could be used for haptics, they tend to restrict the direction or angles that the device can eventually reach. Robot arms, however, are more versatile with a high load-bearing ability and have thus even been used to simulate large objects such as walls or tables in virtual experiences [52]. The speed at which they operate allows for the representation of several objects (see, e. g., [52]). In addition, the detachable head design of commercial robot arms increases the variety of haptic feedback available for a single encountered haptic feedback system [3].

However, most of the past encountered haptic systems focus on providing support to users and increase their immersive experience. In contrast, we aim at using the possibility to provide users with a system that allows them to interact seamlessly with visualisation independent of their hands’ positions. We thus combine a tangible axis with a haptic slider on the robot arm to explore the possibilities that such a system can offer in immersive data analysis.

3 DYNAMIC SLIDER SYSTEM

We designed a robot assisted dynamic slider system (Dynamic Slider) which can automatically align itself with its model in the virtual world (see Fig. 2). It is able to align to different virtual sliders of various lengths and orientation in the same scene. The sliders are built from open-source CNC gantry kits.¹ These were modified by mounting a 3D printed knob to the gantry plate, and replacing the stepper motor with a DC motor to allow constant torque. A 600 pulse per revolution industrial optical encoder was also installed, and we 3D printed 40-tooth, 2 mm pitch timing pulleys for a distance of 0.13 mm per encoder pulse. This was found to be too precise to easily select an exact value manually, and the code adjusted to give a discrete increment spacing at 0.4 mm. An additional rotary encoder can be mounted at a perpendicular angle at the end of the axis.

To help negate cogging torque of the DC motor, a small haptic pulse is integrated directly at the microcontroller encoder interrupt call, which delivers a small amount of torque per discrete value to the H bridge motor driver in the opposite direction of slider travel, assisting the user to not overshoot. An alternative to this could be investigating the use of low cogging torque slotless DC motors, however these are expensive and not easily available.

The slider assembly is mounted on the end of a UR3 (Universal Robot 3.0 series) robotic arm. For the version attached to the robot, we configured the sliders to have a 170 mm selectable range. Longer than this was considered to be impractical for robot mounting, as it would mean that extremities of the slider assembly would move very quickly (potentially dangerously so) when rotated. The UR3 has a 500 mm

¹<https://openbuildspartstore.com/v-slot-mini-v-linear-actuator-bundle/>




Fig. 2. The dynamic slider system is composed of our self-built slider (from a DC motor, an optical encoder, a 3D printed slider knob, a rotary encoder, and a modified belt driven CNC gantry) and a robotic arm. A hand tracking glove is used to capture hand position and posture in VR, and a red button is attached on the work desk for the experiment.

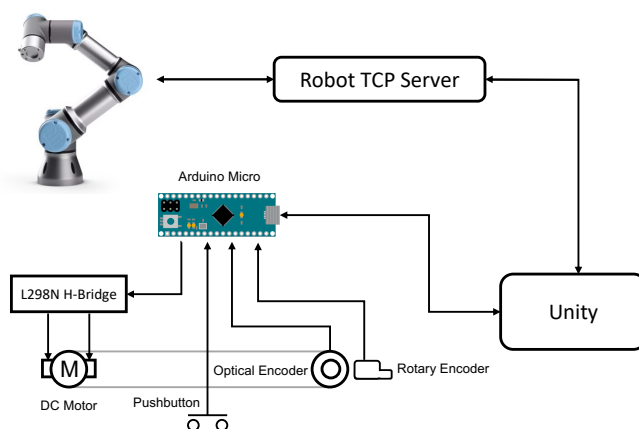


Fig. 3. System connection diagram. A TCP server is used for communication between robot and Unity, and an Arduino is used for communication between physical pops and Unity.

radius of reach, meaning that the total extended slider range available for user selection can be up to 1170 mm, at base, allowing a potential 2,925 discrete values per axis, at full extension.

For our study, we limited the range of our dynamic slider to 760 mm, allowing a total of 1900 discrete points along the slider’s axis, and the full length physical slider adjusted to match this exactly. (see Sect. 4). The robot is controlled through its API,² and all the scripts involved are uploaded to GitHub.³

We used a Vicon motion capture system with five cameras for accurate position tracking and alignment of the slider, hand, and the VR environment. The encoder is read by an Arduino Micro, which also drives the motor via an L298N H-Bridge driver board. The Arduino is connected to the local PC via a USB serial connection. Information from the Vicon is collected and processed from another local device and transmitted through UDP wireless communication. We use a local server to exchange information between the robotic arm and the local PC with a constant communication delay of 5~6 ms. Fig. 3 illustrates the interconnection between the main components of our system.

²<https://www.universal-robots.com/download/manuals-cb-series/script/script-manual-cb-series-sw33/>

³<https://github.com/szdai2021/Universal-Robot.git>

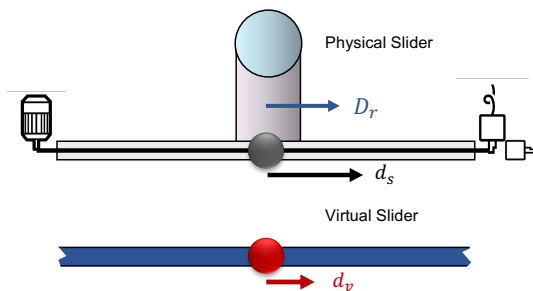


Fig. 4. Displacements during interaction: D_r , Displacement of the robotic arm with regard to physical world; d_s , Displacement of the physical slider knob with regard to the slider itself; and d_v , Displacement of the virtual slider knob with regard to the virtual world.

3.1 Position Synchronisation

Position synchronisation between the physical world and the virtual world is crucial as we are using a short slider that is supposed to mimic, in real time, the interaction that a long slider could provide. The position change of the virtual slider knob, d_v , regarding a fixed coordinate system (the global coordinate in the virtual space) depends on the displacement of the robotic arm, D_r , and the position of the slider knob referring to the physical slider itself, d_s (see Fig. 4). According to equation 1, by having a correct measure of the robotic arm displacement and the physical slider change, an accurate virtual slider displacement can be obtained, where α and β represent the coefficient of converting the physical displacement of the robotic arm and the physical slider knob into the virtual environment coordinates. Therefore, with the tracking of the slider knob position and the position of the robotic arm end effector, the position change of the slider knob is presented in the virtual environment within the offset of the Vicon tracking tolerance and the encoder accuracy.

$$d_v = \alpha D_r + \beta d_s \quad (1)$$

During development, we tested several camera-based hand tracking techniques including Leap Motion and built-in hand tracking from the VR headset. Although camera-based hand tracking is easy to implement, and can capture the hand pose with acceptable accuracy, the reliability is highly dependant on the distance from the camera and the tracking range is substantially limited by the view angle of a single device. Therefore, to have a reliable and robust tracking result, we chose the combination of a dedicated VR glove (Manus Prime II) with a Vicon motion tracking system for pose and position tracking respectively. Several tracking markers are mounted on the back of the glove for overall position tracking, and flex sensors within each finger of the glove provide real-time update of hand posture in VR. An initial calibration step automatically aligns the VR and the physical worlds by comparing the vector angle and position of two VR controllers, tracked both by the Vicon and the VR headset, and then applies a position and rotation offset to the VR rig root position in the Unity Engine.

3.2 Slider and Robotic Arm Behaviour

To extend the perceived slider range and simulate the experience of interacting with a longer slider, both the mobility and the force feedback of the slider should be made as smooth as possible. After some pilot testing, we decided to assign a buffer at either end of the slider, which activates motion of the robotic arm. Outside of these buffer regions, interaction with the slider knob is identical to a normal slider, however once the knob enters a buffer region, the robotic arm starts to move with the slider to extend its range. We considered three different methods of using buffers to achieve a realistic long slider:

- Large-Buffer: a large buffer assignment which activates the robotic arm with constant speed and displacement;
- Small-Buffer: a small buffer assignment to maximize the region not affected by the robotic arm movement; and




Fig. 5. Interaction scenario for the dynamic slider. The robot arm moves in the same direction as the slider knob movement, while the slider motor is pulling in the opposite direction. As the knob gets close to the end of the slider, the speed of the robotic arm and motor output increases.

- Dynamic-Buffer: a large buffer assignment which affects the robotic arm dynamically. The speed of the robotic arm varies based on the relative knob position inside the buffer region;
- We conducted a pilot controlled experiment with simple pointing task with six participants to evaluate these three methods. The results of this pilot experiment are detailed in our supplementary materials at <https://osf.io/puzch/>. Overall, our data indicate that the Dynamic-Buffer technique had noticeably better results in both speed and accuracy. We therefore decided to focus on the Dynamic-Buffer design only in the remainder of this manuscript.

When the slider knob enters the buffer region (green area in Fig. 5), the robotic arm is activated and starts moving in the same direction as the slider knob. To adapt to the resulting change in speed, the speed of the robotic arm is inversely proportional to the distance between the slider knob and the end of the physical slider. Therefore, the robot's speed matches the user's hand movement unless the user is moving the knob faster than the implemented speed limit, which will result in hitting the end of the physical slider. This will, however, not stop the robot's operation which would eventually catch up.

When the robot is moving in the same direction as the user's hand, the slider friction normally felt when pushing the slider will be reduced. This may cause an unnatural experience that the slider is 'pulling' the user's hand. To mitigate this, we activate the DC motor on the slider simultaneously. This generates a force acting on the slider knob in the opposite direction of the robotic arm movement to provide an illusory pushing sensation. The magnitude of the generated force is proportional to the acceleration of the robotic arm, which is tuned to provide the 'correct' amount of opposing force against the hand.

3.3 Safety Consideration

The UR3 has a maximum output setting of 250 N Force; 1000 W Power; 5000 mm/s Speed; and 100 kgm/s Momentum. Although the robotic arm has a collision detection system to automatically stop the robot operation when obvious resistance is detected, for safety we have limited the maximum speed of the robot to minimise the potential for damage and injury. The robot speed was limited to 300 mm/s during the experiment, which is much less than the safety guidelines [71]. We have also set limitations to each of the individual joints so that the robotic arm can only operate in the specifically permitted area. We have also eliminated unexpected movement by resetting the robot position near the base of the robot whenever there is no active interaction in the virtual world and designing the path of the robot to ensure it would




Fig. 6. Diagram of the coverage of the robotic arm (blue area), the base position for reset (yellow point), and the permitted area for user study (green area) from top-down view (Left), and the side view of the robotic arm with slider and button (Right).

always approach the user from far to near rather than laterally. Since we are using a self-built physical slider mounted on the robot, all visible edges on the physical slider are covered with a layer of soft foam to further limit risks.

4 CONTROLLED EXPERIMENT

We conducted a controlled experiment in which users manipulate a slider in an immersive context to evaluate the performance difference between three conditions: *Virtual* – without any tangible or haptic feedback during interactions; *Physical* – with haptic feedback from a 1:1 ratio physical slider; *Dynamic* – with haptic feedback from the combination of a short physical slider and a robotic arm. We measured task completion time and error rate during positioning of a virtual control at a specific location, as well as subjective feedback from the users. Such tasks are common to evaluate new interactive techniques or devices in the HCI and visualisation literature (see, e. g., [37]).

4.1 Task and Study Design

We now describe the design of our user study and the motivation behind it. The aiming task is conducted by generating a random target and asking the participant to move the knob and point it to the centre of the target. For precise visual feedback, a red arrow is attached to the virtual slider knob and the random target is designed as a yellow arrow with a thin red line indicating the centre of the target (see Fig. 7). There is a red button in both the virtual world and the physical world for the participants to claim their completion. Participants are asked to perform the task with their right hand only, which provides a real-world-like experience where small position adjustment is required to approach each object. The target position is generated randomly based on the current position of the virtual slider knob, and the new position is restricted in three ranges: *small range* where wrist movement only can cover the target range (19 cm to 21cm); *middle range* where elbow movement only is sufficient (34 cm to 36 cm); and *large range* where waist rotation may be required (51 cm to 53 cm). The task repeats 10 times for each range and each condition includes all three ranges. Therefore, the total number of responses per participant is 3 conditions \times 3 ranges \times 10 repeated tasks = 90 responses.

The maximum speed and reachable area of the robotic arm is restricted for safety as described above. These limitations reduce the flexibility of the dynamic slider and the influence of this restriction is expected in the data collection and the subjective feedback.

4.2 Procedure

The order of conditions for each participant and the order of the random target for each condition are counter-balanced to reduce the impact of learning effects. Participants were instructed to put on the headset prior to starting each condition, to prevent them from seeing the physical slider. For each condition, there are 10 training trials of the pointing




Fig. 7. VR scene of the study. Users have to use their hand to grab the virtual knob and point to the yellow marker with the help of a red-arrow.

task. Before the start of the actual tasks, participants are asked to move the virtual knob to the centre of the slider.

At the end of each condition, a set of questionnaires involving 14 virtual reality presence questions (IPQ) [59] and two confidence questions is completed by the participant based on their experience. After all three conditions are completed, participants need to finish a post-study survey giving their reflections about the three conditions.

In summary, our study consisted of four main steps: briefing, training, tasks, and post-study questionnaire. The briefing stage consisted of approximated five minutes of reading the task descriptions. We deliberately asked participants to perform the training trials slowly at first, to ensure they could get used to the interaction with the slider. Each participant performed a total of 100 trials (90 actual trials + 10 training trials) in a maximum of 20 minutes time. After each condition the participant was asked to fill the condition feedback survey and then, after the whole experiment, to complete a post-study questionnaire. Completing the feedback forms could take up to 15 minutes. In total, the duration of the study was around 45 minutes.

4.3 Participants and Data Collection

We recruited a total of 30 participants from our university (23 males and 7 females with mean age of 24.7 and median age of 25). The study was approved by the local ethics committee. We collected three types of data: i) task-related data, ii) interaction log data, iii) subjective measures data. For the task-related data, we recorded the error rate measured as the difference (in mm) between the current position of the virtual slider knob and the centre of the target after each of the button presses, as well as the completion time. Meanwhile, every position change in the virtual world during the experiment is recorded for later analysis as the interaction log data. In the post-study questionnaire, we asked participants to rank the sliders in all three conditions with respect to ease of use, confidence to use, and realism of the experience compared to a physical task (presence). We also encouraged participants to briefly describe the reasons for their choice.

5 RESULTS

Data from experiments has been historically analysed with null hypothesis significance testing (NHST). However, recent criticism of NHST-based analysis have highlighted its limitations. [2, 20, 27] and its tendency to lead to dichotomous inferences [6, 32] which can eventually lead to less robust findings in scholarly communications [16]. We thus decided to follow current APA recommendations [70] and report our results using estimation techniques with simple effect size and confidence intervals (CIs). The term effect size here refers to mean differences and ratios [21]. We interpret our CIs as providing different strengths of evidence about the mean [6, 9] and rather avoid the use of words such as “significant” which reflect more dichotomous interpretations. For completeness and although we only use estimation techniques, we provide *p*-values in each figure and in our additional materials, but a method to obtain a *p*-value reading of our results is detailed by Krzywinski and Altman [41]. Following current best practices, we pre-registered the




Fig. 8. Mean completion time for each condition. Errors bars are 95% Confidence intervals (CIs).




Fig. 9. Pair-wise completion time ratios. Error bars: 95% CIs

analysis of our empirical data [10, 16] and share our experimental materials, data, pre-registration, and code at <https://osf.io/puzch/>. We detail below the results of our pre-registered analysis.

5.1 Completion time

We report the absolute mean values of task completion time in seconds for each condition in Fig. 8 and the pair-wise ratios between each condition in Fig. 9. To correct for the positive skewness of completion-time data, we analysed log-transformed measurements and present anti-logged results, which is, considering the log-normal distribution of time data, a standard analysis [5, 57] that is commonly used in HCI and visualisation work (e. g., [8, 38, 73]). We consequently present geometric means which dampens the effect of potential extreme completion times that are likely to bias arithmetic means. We can find in Fig. 8, from the complete lack of overlap between the confidence intervals of each conditions, very strong evidence that the virtual slider is slower than the dynamic slider which is in turn slower than the physical slider. This is confirmed by Fig. 9 which highlights that the virtual slider is almost 1.5 times (1.45, CIs[1.41;1.48]) slower than the physical slider but that the physical slider is 0.86 (CIs[0.844;0.874]) times faster than the dynamic slider. The size of the confidence intervals in Fig. 9 also highlight that the performances across participants are stable.

5.2 Error rate

We report the means of accuracy (in mm) for each condition in Fig. 10 and the pair-wise ratios between conditions in Fig. 11. We observe in Fig. 10 very strong evidence that the virtual slider has a higher error rate than the other two techniques for which our experimental data does not seem to show evidence of a difference. This is confirmed in Fig. 11 which highlights that the difference is of slightly less than 1 mm between the virtual slider and both the physical (0.77, CIs[0.652;0.894]) and dynamic slider (0.75, [0.876;0.635]). The accuracy of the physical and dynamic slider is, according to our data, approximately down to the millimeter (respectively 0.99 [0.923;1.06] and 1.01 [0.950;1.07]).

5.3 Ranking

We report on the ranking of each condition by the participants in Table 1. We can see in the table that the fully virtual slider was ranked as the least preferred technique more often than the other two (19 times ranked last), while the physical slider was ranked as the preferred technique the most (14 times). Overall, the mean ranking for the physical slider (1.67) are slightly better than the mean ranking for the dynamic slider (1.9) and much better than the mean ranking for the virtual slider (2.43).




Fig. 10. Average absolute error (in mm). Error bars: 95% CIs




Fig. 11. Pair-wise error differences (in mm). Error bars: 95% CIs

5.4 Qualitative Feedback

We report now the comments that the participants made about each of the three techniques that they tried in our controlled experiment.

Virtual Slider: The user experience of interacting with a slider with purely mid-air hand gestures is highly dependant on their background experience of VR. Indeed, the experimenter observed that most of the participants who did not have a lot of experience with VR systems were struggling with the virtual condition during the experiment. Four participants said they cannot grab the virtual slider knob properly even after the training, and one of them felt it is confusing to grab the virtual slider knob without actually touching something in hand. Only two participants who have a lot of experience in hand gesture interactions preferred this condition over the other two because of its flexibility in the interaction speed. Of six participants who preferred the virtual slider the most (Table 1), four had no prior VR experience. However, almost all participants (24 out of 30 including the four with no VR experience) agree that it is the most difficult condition to perform an accurate pointing task.

Physical Slider: The most common comment (21 out of 30) for the physical slider is that it is the easiest slider to use with minimum learning effort required as it resembles what participants do in real life. Additionally, the pointing task seems to be effortless in this condition. Participants also mentioned that they appreciated having the ability to “touch something”. Another comment (7 out of 30) on the physical slider is that a visible gap between the virtual hand and the virtual slider knob during interaction makes the experience seem less real.

Dynamic Slider: Most of the participants (20 out of 30) commented that the user experience of the dynamic slider is somewhere in between the virtual slider and the physical slider. Although the generated haptic feedback and the presence of the tangible controller (the physical slider knob) provides a better interaction experience, there are still tiny but noticeable difference between the dynamic slider and the physical slider (comments from 15 participants). This kind of difference might be negligible if a user is focusing on something other than the slider, but it was logically more problematic in our controlled experiment which solely

Table 1. Ranking of each condition with **best** and **worst** highlights.

Technique	Mean	Median	SD	1st	2nd	3rd
Virtual	2.43	3	0.817	6	5	19
Physical	1.67	2	0.711	14	12	4
Dynamic	1.9	2	0.759	10	13	7

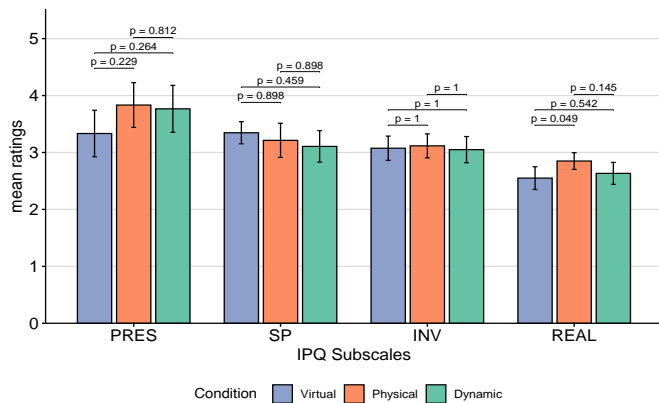


Fig. 12. Subscales of the IPQ questionnaires. IPQ: General Presence (GP), spatial presence (SP), involvement (INV), and realism (REAL).

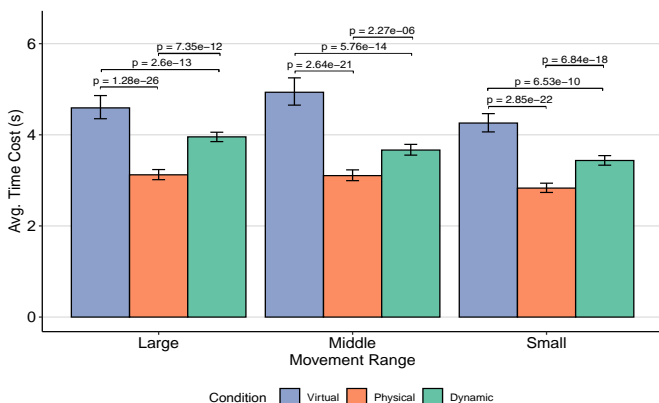


Fig. 13. Mean completion time for each range. Error bars: 95% CIs

focused on a pointing task. Another common comment is that when participants were moving the knob and approaching the destination, there was a small vibration-like experience and a small unexpected movement of the slider knob. Therefore, participants have to make an extra adjustment in order to point to the target properly. Despite the above comments, an interesting fact is that although participants noticed the difference between dynamic slider and the physical slider, many of them (around 18) did not realise they were actually interacting with a short slider on a moving robotic arm, as participants were intentionally asked to wear the VR headset at the beginning of each condition so that their feedback would not be affected by the physical environment change (the robotic arm with the short slider will be right in front of their hand).

5.5 Additional analyses

We detail below the additional analyses that we have conducted that were not specified in our pre-registration.

Presence questionnaire

Although we initially planned to report on the IPQ questionnaire for each condition, we neglected to specify this analysis in our pre-registration. The results of the questionnaire are visible in Fig. 12. Judging from the overlap of the confidence intervals for all conditions and subsets of the IPQ in Fig. 12, our experiment did not allow us to find evidence of a difference between the three tested conditions.

The impact of distance

In our controlled experiment we tested *small*, *medium* and *large* distance ranges to the target (Sect. 4.1). To confirm that the distance to




Fig. 14. Mean error for each range. Error bars: 95% CIs

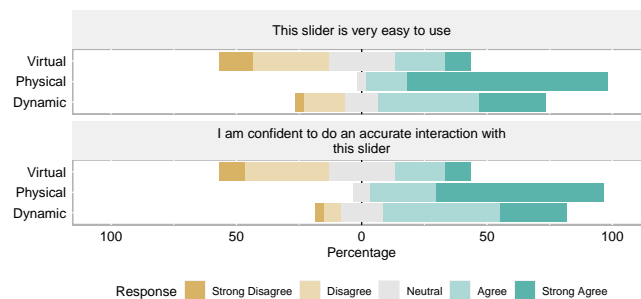


Fig. 15. Subjective feedback on usability and confidence levels.

the target does not influence the results, we conduct a supplementary, i. e., non pre-registered, analysis of the impact of range on the results. A breakdown per range is presented in Fig. 13 for completion time and in Fig. 14 for error rates. We see in Fig. 13 that the results found in the previous figures hold across all possible range of distance that we investigated. Similarly, Fig. 14 confirms that the tendencies observed above are consistent across all studied ranges.

Subjective assessments on usability and confidence

We report in Fig. 15 the Likert-scale evaluations of the participants. These quantitative results seem to confirm participants' comments and ranking that we previously detailed. The physical slider appears the easiest to use and the one providing the highest level of confidence, followed quite closely by the dynamic slider. The virtual-only slider is reported as the worst overall for both confidence and ease of use.

6 DISCUSSION

From our experiment and its data we can conclude that the dynamic slider that we developed seems to be as accurate as a real physical slider and much more precise than a virtual slider. This difference is likely to be partially explained by the exit-error [11]: it is impossible for participants to stay completely still after correctly positioning the virtual knob and before validating the placement. The accuracy difference between the virtual slider and the other two conditions is, however, relatively small. While such a small difference in accuracy might not matter in many use cases, past work (e. g., [11]) has argued that for visualisation cases it might, e. g., in the case of a surgeon performing a surgery planning task in an immersive context such as the one we used.

Our data also highlights that the physical slider is faster than the dynamic slider which is in turn faster than the virtual slider. From our observation of the participants, we argue that this difference in task completion time can be explained by the fact that participants tended to adjust the knob's position more times with the virtual slider than with the two physical sliders. It is likely that these extra adjustments




Fig. 16. Interacting with the time-series selector in both physical world (Left) and the virtual world (right).

were not necessary in the two physical conditions since users could physically hold a knob, thus echoing past research on the benefits of tangible devices in immersive contexts [23, 74]. The difference between the dynamic slider and the physical one is most likely explained by the current speed limitation of the dynamic slider. Theoretically, interaction with the dynamic slider could be made to match the physical slider as the robotic arm can accelerate extremely quickly. But for safety reasons the current implementation limits the robot speed.

Overall, it seems that our dynamic slider outperforms a fully virtual slider in all aspects (speed, accuracy and preferences) while being slightly slower than a physical slider and not quite as highly preferred. Since the physical slider was designed to be the best possible baseline for this particular task, these slight differences are not surprising and the performances we obtained with the dynamic slider are promising.

Indeed, it is unlikely that physical 1:1 sliders will be made for all possible scenarios and task types and our dynamic slider provides more versatility and modularity even for the simple task that we decided to focus on in this controlled experiment. But, regarding the advantages of the dynamic slider, we have to consider use-cases beyond the very simple, yet important one, studied in this controlled experiment. Our ultimate vision for a robot-facilitated interactive device that would follow the user's hand is within the context of immersive visualisation tasks for which we argue in the following section that the device can be useful and that we now illustrate in three use cases.

7 USE CASES

We explored three different use cases of the dynamic slider system. These use cases demonstrate the potential of the haptic feedback implementation in immersive data visualisation and are presented here in order of increasing sophistication and flexibility of the mapping of the physical control to the interaction affordances in the immersive visualisation. Please see the companion video to this paper for a detailed walk-through demonstration of the various interactions we implemented for each of these use cases.

7.1 Independent Variable Selection in a Time-Series Chart

Our first use case is a straight-forward application of the dynamic slider configuration of an extended-range but fixed-orientation axis slider selection task, similar to that tested in our study (Section Sect. 4). The 76 cm range of the dynamic slider can easily span a comfortable width for a detailed VR view of potentially complex data, hovering at waist-to-head height at about arms' length from the user.

To demonstrate this, we created a time-series data visualisation with a one-to-one mapping of slider knob position to position the x -axis, corresponding to the independent variable in a 2D line chart, e.g. time or (as shown in Figure Fig. 16) income. Following their design space for slider interaction with visualisation dynamics, Smiley et al. [61] we use haptic notching to snap the slider position to the nearest data point in the visualisation, and then display an infobox showing details for the data element at that point.

7.2 3D Medical Image Sectioning Tool

Cordeil et al. presented an "Embodied Axes" device featuring six small physical slider devices, paired along three orthogonal axes, to select a



Fig. 17. One robotic arm and one mounted slider represents multiple slider controls and reaches to the required position according to hand position (left). Select a slicing plane of a skeleton in the virtual world and move among in three axis direction (right).



Fig. 18. Robotic arm holding the rotary encoder for rotational interactions according to hand position (left). Select the rotary control in the virtual world and rotate the virtual skeleton around that axis (right).

box region within the $10 \times 10 \times 10$ cm volume of the axes [17]. The device was presented to domain experts in medical imaging in the context of slicing 3D CT imagery. The experts gave positive feedback about the utility of such an embodied interaction device, but with clear consensus that the biggest limitations were the small selection volume and the fixed orientation of the axes. The robot assisted haptic slider overcomes both of these limitations, extending the range of the selection to a reorientable $76 \times 76 \times 76$ cm volume, see Figure Fig. 17.

While only one physical slider is present, in VR the user sees six slider controls, paired along three orthogonal axes. As the user reaches towards a virtual slider control, the robot arm moves and reorients the physical slider such that the physical knob is in-situ with the virtual slider's knob by the time the user's fingers arrive at its location. Further, the user sees three rotation controls aligned at the top, front and side central axes of the image volume. When the user reaches for these controls the robot places the physical rotator encoder control at the end of the slider in-situ with the axis rotation affordance (see Fig. 18). They can then freely rotate the image around that axis. That is, in the former, the sliders remain fixed in world x, y, z coordinates to allow further sectioning of the image along arbitrary cutting planes, or the sliders are reoriented with the image such that sectioning remains aligned with the axial, sagittal and coronal planes of the rotated image volume.

The average haptic delay of this use case, i.e. time between the user initiating interaction in the virtual world and the dynamic robot/slider system moving the knob to the target location from its base position, is measured by initiating five evenly distributed points on each of the slider in the x, y, z axis direction. The average haptic delay for x and y are 1004 ms and 1296 ms respectively. It is 2881 ms for vertical axis, z . This is due to the fact that direct movement from the base position to any any point along this axis requires large joint movement and has high risk of collision with the table below and the user. Therefore, a sequence of movement is designed to avoid the risk, which requires more time. This delay could be significantly reduced by increasing the maximum speed at which we allow the robot to move, but for now we conservatively limit it for safety (Sect. 3.3).




Fig. 19. Selecting a point data in a scatter plot reveals its value changes with time. Robotic arm holding the slider to match that point data spatial position as well as the gradient of the value change (left) while virtual hand is in touch with point data (right).

7.3 DimpVis Multi-time-Series Curve Selector

The final use-case we consider is a virtual reality adaptation of the *DimpVis* scenario for direct interaction with time-series data curves, introduced by Kondo and Collins for large touch screen interaction [40]. An interesting aspect of this application is that, while in the other use cases presented above the slider control is used for interaction with the axes corresponding to up to three spatial dimensions, in *DimpVis* the primary interaction affordances are the datapoint marks themselves.

As seen in Figure Fig. 19, the data considered in this use case is from the GapMinder⁴ dataset of life-expectancy versus income curves for selected countries. While the user is not interacting, the view shown is a scatter plot of average income versus average life expectancy with a point for each of the countries for a single year. When the user moves their hand within the axis of the scatter plot, the data point nearest to the user's fingers is highlighted and a time curve is displayed showing the trajectory of life-expectancy versus income for the country corresponding to that point over the full time range of the data. Simultaneously with the nearest data point being highlighted in this way, the robot rotates the slider control to the tangent of the curve at that point, and moves the slider-knob such that the knob will be ready at the position of the data point by the time the user's fingers reach it. Moving the slider forward along the tangent will then advance the highlight to the next data point in the series, and simultaneously animate all the other data points to show the data for the same year.

The average haptic delay for this use case is measured in two ways: the delay between the occurrence of interaction in VR until the physical slider has reached its destination from base position to one of the data points (base-to-point delay) or from a previous selected point to another point (point-to-point delay). The average base-to-point delay is 2515 ms, and the average point-to-point delay is 1282 ms. However, both delays are theoretically linearly proportional to the distance travelled. Consequently, the average haptic delay of this use case is likely to vary for different data sets with different scatter point distribution.

8 LIMITATIONS

While our controlled study required participants to remove their hand and then reacquire the control after restarting the task timer, the study did not evaluate cases where the haptic controls move independently of the user's hand. Further research could explore in what circumstances (if any) such independent movement has an effect on users' ability to reacquire the control.

Our system relies on a marker-based infrared motion capture system which is currently the fastest and most accurate system for real-time hand tracking. While some headsets such as Oculus Quest2 feature in-built markerless tracking their capability is still limited in latency and hand pose accuracy. However, we expect that the headset-based tracking will improve to the point where an external system is no longer necessary. Similarly, the robot device we use (UR3) is relatively expensive and requires an external control and powersupply box. While the cost and set up for such equipment may be acceptable for certain

⁴gapminder.org

industrial and medical applications, simpler, lower-cost commodity devices would be required for wide adoption of encountered haptic visualisation applications.

9 CONCLUSION AND FUTURE WORK

While our controlled experiment and use-cases already demonstrate usability and applicability of the robot-assisted haptic controller we have developed, there are many avenues for future research in this vein.

A first obvious avenue for future work lies in the evaluation of the use cases we have presented in this manuscript. Indeed, we focused our evaluation on a 1D use case with a slider but the usability of the proposed interaction system should further be evaluated in other 3D use-cases. From our results, we can hypothesize that the very good accuracy we obtained would also be found in 3D use cases and that it is likely that participants would also feel like they are interacting the virtual slider with realistic haptic feedback. Concerning the observed delay measured in the use cases, it could be significantly reduced by increasing the robot speed or designing a unique and smoother path of movement from the base position to various destinations.

Second, we have so far tested only a slider and rotary encoder as controls attached to the robot arm. There are many other physical controls and actuators that could be attached, in order, for example, to simulate an entire cockpit of controls. Similarly, there may be application of a system like ours to provide access to data for blind or low vision (BLV) people. Past work has found benefit in static tactile graphics and 3D models for BLV people [14]. In future we would like to explore whether a system like ours could be a general solution to simulate actions with realistic haptic feedback for BLV people in their training or education.

Another exciting avenue for future work lies in extending the range of user interactions through different hardware and combinations of devices. Our UR3 robot is the most affordable model in that suppliers' series of "collaborative robot arms" (safe for direct human-robot interaction), but similar robots with up to approximately double the range are available, and one may envision combining sets of robots to fully surround the users. Mounting the robots on wheeled bases would extend the range even further, allowing the robots to follow users as they walk around the virtual environment. Of course, with greater range and complexity, safety may become a concern, or an opportunity for intelligent robot uprising.

Further study would also test "robohapalytics" in different application domains and use user-centred and collaborative design techniques to seek critical and formative feedback from domain experts. Further, while we are satisfied that the speed, accuracy and presence feedback collected in our controlled study demonstrates the viability and usability of our dynamic slider system to simulate a slider with range that is physically larger than the range of the slider track, a logical next step would be to test the dynamic reorientation of the slider in interactive scenarios such as those presented in Sections Sect. 7.2 and Sect. 7.3. In particular, it would be interesting to see if there are limits to the degree of curvature of non-linear trajectories that can be usefully tracked by the robot in scenarios similar to the *DimpVis* use case.

Finding a markerless alternative for the hand tracking system would also have significant improvement on the overall immersive experience if bare-hand interaction is allowed. Another interesting possibility for future study would be to test how robot maneuvered data physicalisations can support remote collaboration. For example, an assumption would be that haptic representations of data that is manipulated by other users would give a stronger sense of physical collocation than purely visual representations.

10 ACKNOWLEDGEMENTS

The authors wish to thank their participants. This research was supported under the Australian Research Council's Discovery Projects funding scheme (project number DP180100755) and the Knut and Alice Wallenberg Foundation (grant KAW 2019.0024).

REFERENCES

- [1] M. Abdullah, M. Kim, W. Hassan, Y. Kuroda, and S. Jeon. Hapticdrone: An encountered-type kinesthetic haptic interface with controllable force feedback: Example of stiffness and weight rendering. In *IEEE Haptics Symposium (HAPTICS)*, pp. 334–339. IEEE, 2018. **Open Access**. doi: 10.1109/HAPTICS.2018.8357197
- [2] V. Amrhein, S. Greenland, and B. McShane. Scientists rise up against statistical significance. *Nature*, 567(7748):305–307, 2019. **Open Access**. doi: 10.1038/d41586-019-00857-9
- [3] B. Araujo, R. Jota, V. Perumal, J. X. Yao, K. Singh, and D. Wigdor. Snake charmer: Physically enabling virtual objects. In *Proc. TEI*, pp. 218–226, 2016. **Open Access**. doi: 10.1145/2839462.2839484
- [4] L. Besançon, W. Aigner, M. Boucher, T. Dwyer, and T. Isenberg. 3D mobile data visualization. In B. Lee, R. Dachsel, P. Isenberg, and E. K. Choe, eds., *Mobile Data Visualization*, chap. 4, pp. 111–150. CRC Press, Boca Raton, FL, USA, 2021. doi: 10.1201/9781003090823-4
- [5] L. Besançon and P. Dragicevic. La Différence Significative entre Valeurs p et Intervalles de Confiance. In AFIHM, ed., *29ème conférence francophone sur l'Interaction Homme-Machine*, p. 10. AFIHM, Poitiers, France, Aug. 2017. Alt.IHM.
- [6] L. Besançon and P. Dragicevic. The Continued Prevalence of Dichotomous Inferences at CHI. In *Proc. CHI Extended Abstracts*, p. 1–11. ACM, Glasgow, United Kingdom, May 2019. **Open Access**. doi: 10.1145/3290607.3310432
- [7] L. Besançon, P. Issartel, M. Ammi, and T. Isenberg. Hybrid Tactile/Tangible Interaction for 3D Data Exploration. *IEEE Trans. Vis. Comp. Graph.*, 23(1):881–890, Jan. 2017. doi: 10.1109/TVCG.2016.2599217
- [8] L. Besançon, P. Issartel, M. Ammi, and T. Isenberg. Mouse, tactile, and tangible input for 3D manipulation. In *Proc. CHI*, pp. 4727–4740. ACM, New York, 2017. **Open Access**. doi: 10.1145/3025453.3025863
- [9] L. Besançon, Y. Jansen, A. Cockburn, and P. Dragicevic. Definitely maybe: Hedges and boosters in the hci literature. 2021. doi: 10.31219/osf.io/mjg7h
- [10] L. Besançon, N. Peiffer-Smadja, C. Segalas, H. Jiang, P. Masuzzo, C. Smout, E. Billy, M. Deforet, and C. Leyrat. Open science saves lives: lessons from the covid-19 pandemic. *BMC Medical Research Methodology*, 21(1):1–18, 2021. **Open Access**. doi: 10.1186/s12874-021-01304-y
- [11] L. Besançon, M. Sereno, L. Yu, M. Ammi, and T. Isenberg. Hybrid Touch/Tangible Spatial 3D Data Selection. *Computer Graphics Forum*, 38(3):553–567, June 2019. doi: 10.1111/cgf.13710
- [12] L. Besançon, A. Ynnerman, D. F. Keefe, L. Yu, and T. Isenberg. The State of the Art of Spatial Interfaces for 3D Visualization. *Computer Graphics Forum*, 40(1):293–326, Feb. 2021. **Open Access**. doi: 10.1111/cgf.14189
- [13] W. Büschel, J. Chen, R. Dachsel, S. Drucker, T. Dwyer, C. Görg, T. Isenberg, A. Kerren, C. North, and W. Stuerzlinger. Interaction for immersive analytics. In K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, eds., *Immersive Analytics*, chap. 4, pp. 95–138. Springer, Berlin/Heidelberg, 2018. **Open Access**. doi: 10.1007/978-3-030-01388-2_4
- [14] M. Butler, L. M. Holloway, S. Reinders, C. Goncu, and K. Marriott. Technology developments in touch-based accessible graphics: A systematic review of research 2010–2020. In *Proc. CHI*, pp. 1–15, 2021. doi: 10.1145/3411764.3445207
- [15] A. Cassinelli and M. Ishikawa. Volume slicing display. In *SIGGRAPH ASIA Art Gallery & Emerging Technologies*, p. 88, 2009. **Open Access**. doi: 10.1145/1665137.1665207
- [16] A. Cockburn, P. Dragicevic, L. Besançon, and C. Gutwin. Threats of a replication crisis in empirical computer science. *Commun. ACM*, 63(8):70–79, July 2020. doi: 10.1145/3360311
- [17] M. Cordeil, B. Bach, A. Cunningham, B. Montoya, R. T. Smith, B. H. Thomas, and T. Dwyer. Embodied axes: Tangible, actuated interaction for 3d augmented reality data spaces. In *Proc. CHI*, pp. 1–12, 2020.
- [18] M. Cordeil, B. Bach, Yongchao Li, E. Wilson, and T. Dwyer. Design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 46–50, 2017. doi: 10.1109/PACIFICVIS.2017.8031578
- [19] G. de Haan, M. Koutek, and F. H. Post. Towards intuitive exploration tools for data visualization in VR. In *Proc. VRST*, pp. 105–112. ACM, New York, 2002. **Open Access**. doi: 10.1145/585740.585758
- [20] P. Dragicevic. Fair statistical communication in HCI. In J. Robertson and M. Kaptein, eds., *Modern Statistical Methods for HCI*, chap. 13, pp. 291–330. Springer International Publishing, Cham, Switzerland, 2016. **Open Access**. doi: 10.1007/978-3-319-26633-6_13
- [21] P. Dragicevic. A Mean Difference is an Effect Size. Research Report RR-9354, Inria Saclay Ile de France, July 2020.
- [22] A. Fedoseev, A. Tleugazy, L. Labazanova, and D. Tsetserukou. Teslamirror: Multistimulus encounter-type haptic display for shape and texture rendering in vr. In *ACM SIGGRAPH 2020 Emerging Technologies*, pp. 1–2. 2020.
- [23] M. Feick, S. Bateman, A. Tang, A. Miede, and N. Marquardt. Tangi: Tangible proxies for embodied object exploration and manipulation in virtual reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 195–206. IEEE, 2020. doi: 10.1109/ISMAR50242.2020.00042
- [24] G. W. Fitzmaurice. *Graspable User Interfaces*. PhD thesis, University of Toronto, Canada, 1996. **Open Access**.
- [25] S. Follmer, D. Leithinger, A. Olwal, N. Cheng, and H. Ishii. Jamming user interfaces: Programmable particle stiffness and sensing for malleable and shape-changing devices. In *Proc. UIST, UIST '12*, p. 519–528. ACM, New York, 2012. doi: 10.1145/2380116.2380181
- [26] S. Follmer, D. Leithinger, A. Olwal, A. Hogge, and H. Ishii. Inform: Dynamic physical affordances and constraints through shape and object actuation. In *Proc. UIST, UIST '13*, p. 417–426. ACM, New York, 2013. doi: 10.1145/2501988.2502032
- [27] G. Gigerenzer. Mindless statistics. *The Journal of Socio-Economics*, 33(5):587–606, 2004. doi: doi.org/10.1016/j.soc.2004.09.033
- [28] S. R. Gomez, R. Jianu, and D. H. Laidlaw. A fiducial-based tangible user interface for white matter tractography. In *Advances in Visual Computing*, pp. 373–381. Springer, Berlin, Heidelberg, 2010. **Open Access**. doi: 10.1007/978-3-642-17274-8_37
- [29] E. J. Gonzalez, P. Abtahi, and S. Follmer. Reach+ extending the reachability of encountered-type haptics devices through dynamic redirection in vr. In *Proc. UIST*, pp. 236–248, 2020.
- [30] M. A. Goodrich, A. C. Schultz, et al. Human-robot interaction: a survey. *Foundations and Trends® in Human-Computer Interaction*, 1(3):203–275, 2008. doi: 10.1561/1100000005
- [31] D. Hayatpur, S. Heo, H. Xia, W. Stuerzlinger, and D. Wigdor. Plane, ray, and point: Enabling precise spatial manipulations with shape constraints. In *Proc. UIST, UIST '19*, p. 1185–1195. ACM, New York, 2019. doi: 10.1145/3332165.3347916
- [32] J. Helske, S. Helske, M. Cooper, A. Ynnerman, and L. Besançon. Can visualization alleviate dichotomous thinking? effects of visual representations on the cliff effect. *IEEE Transactions on Visualization & Computer Graphics*, 27(8):3397–3409, apr 2021. **Open Access**. doi: 10.1109/TVCG.2021.3073466
- [33] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proc. CHI*, pp. 452–458. ACM, New York, 1994. **Open Access**. doi: 10.1145/191666.191821
- [34] H. Ishii and B. Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proc. CHI*, pp. 234–241. ACM, New York, 1997. **Open Access**. doi: 10.1145/258549.258715
- [35] P. Issartel, F. Guéniat, and M. Ammi. Slicing techniques for handheld augmented reality. In *Proc. 3DUI*, pp. 39–42. IEEE Computer Society, Los Alamitos, 2014. **Open Access**. doi: 10.1109/3DUI.2014.6798839
- [36] B. Jackson, T. Y. Lau, D. Schroeder, K. C. Toussaint, and D. F. Keefe. A lightweight tangible 3D interface for interactive visualization of thin fiber structures. *IEEE Trans. Vis. Comp. Graph.*, 19(12):2802–2809, Dec. 2013. **Open Access**. doi: 10.1109/TVCG.2013.121
- [37] Y. Jansen, P. Dragicevic, and J.-D. Fekete. Tangible Remote Controllers for Wall-Size Displays. In *ACM annual conference on Human Factors in Computing Systems. CHI '12*, pp. 2865–2874. ACM, Austin, TX, United States, May 2012. doi: 10.1145/2207676.2208691
- [38] Y. Jansen, P. Dragicevic, and J.-D. Fekete. Evaluating the Efficiency of Physical Visualizations. In *Proc. CHI*, pp. 2593–2602. ACM, ACM, Paris, France, Apr. 2013. doi: 10.1145/2470654.2481359
- [39] Y. Kim, H. J. Kim, and Y. J. Kim. Encountered-type haptic display for large vr environment using per-plane reachability maps. *Computer Animation and Virtual Worlds*, 29(3-4):e1814, 2018. doi: 10.1002/cav.1814
- [40] B. Kondo and C. Collins. Dimpvis: Exploring time-varying information visualizations by direct manipulation. *IEEE Trans. Vis. Comp. Graph.*, 20(12):2003–2012, 2014. doi: 10.1109/TVCG.2014.2346250
- [41] M. Krzywinski and N. Altman. Points of significance: Error bars. *Nature Methods*, 10(10):921–922, Oct. 2013. **Open Access**. doi: 10.1038/nmeth.2659

- [42] M. Le Goc, L. H. Kim, A. Parsaei, J.-D. Fekete, P. Dragicevic, and S. Follmer. Zooids: Building blocks for swarm user interfaces. In *Proceedings of the 29th annual symposium on user interface software and technology*, pp. 97–109, 2016.
- [43] J. H. Lee, Y. Kim, S.-G. An, and S.-H. Bae. Robot telekinesis: application of a unimanual and bimanual object manipulation technique to robot control. In *2020 IEEE international conference on robotics and automation (ICRA)*, pp. 9866–9872. IEEE, 2020. doi: 10.1109/ICRA40945.2020.9197517
- [44] D. Leithinger, S. Follmer, A. Olwal, S. Luescher, A. Hogge, J. Lee, and H. Ishii. Sublimate: State-changing virtual and physical rendering to augment interaction with shape displays. In *Proc. CHI, CHI '13*, p. 1441–1450. ACM, New York, 2013. [Open Access](#). doi: 10.1145/2470654.2466191
- [45] D. Leithinger and H. Ishii. Relief: A scalable actuated shape display. In *Proc. TEI, TEI '10*, p. 221–222. ACM, New York, 2010. [Open Access](#). doi: 10.1145/1709886.1709928
- [46] L. Lischke, P. W. Woundefinedniak, S. Mayer, A. Preikschat, and M. Fjeld. Using variable movement resistance sliders for remote discrete input. In *Proc. ISS, ISS '17*, p. 116–125. ACM, New York, 2017. doi: 10.1145/3132272.3134135
- [47] D. López, L. Oehlberg, C. Doger, and T. Isenberg. Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3D data exploration. *IEEE Trans. Vis. Comp. Graph.*, 22(5):1616–1629, May 2016. [Open Access](#). doi: 10.1109/TVCG.2015.2440233
- [48] K. Lundin (Palmerius), M. Cooper, A. Persson, D. Evstedt, and A. Ynnerman. Enabling design and interactive selection of haptic modes. *Virtual Reality*, 2007. [Open Access](#). doi: 10.1007/s10055-006-0033-7
- [49] K. Lundin (Palmerius), M. Sillén, M. Cooper, and A. Ynnerman. Haptic visualization of computational fluid dynamics data using reactive forces. In *International Conference on Visualization and Data Analysis, San Jose, USA*, pp. 95–102, 2005. [Open Access](#). doi: 10.1117/12.587029
- [50] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas. *Immersive Analytics*, vol. 11190. Springer, Berlin, Heidelberg, Germany, 2018. doi: 10.1007/978-3-030-01388-2
- [51] P. Marshall, R. Fleck, A. Harris, J. Rick, E. Hornecker, Y. Rogers, N. Yuill, and N. S. Dalton. Fighting for control: Children's embodied interactions when using physical and digital representations. In *Proc. CHI*, pp. 2149–2152. ACM, New York, 2009. [Open Access](#). doi: 10.1145/1518701.1519027
- [52] V. R. Mercado, T. Howard, H. Si-Mohammed, F. Argelaguet, and A. Lécuyer. Alfred: the haptic butler on-demand tangibles for object manipulation in virtual reality using an ethd. In *Proc. WHC*, pp. 373–378, 2021. doi: 10.1109/WHC49131.2021.9517250
- [53] I. C. Olson, Z. Atrash Leong, U. Wilensky, and M. S. Horn. “It's just a toolbar!” Using tangibles to help children manage conflict around a multi-touch tabletop. In *Proc. TEI*, pp. 29–36. ACM, New York, 2011. [Open Access](#). doi: 10.1145/1935701.1935709
- [54] I. Poupyrev, T. Nashida, S. Maruyama, J. Rekimoto, and Y. Yamaji. Lumen: Interactive visual and shape display for calm computing. In *ACM SIGGRAPH 2004 Emerging Technologies, SIGGRAPH '04*, p. 17. ACM, New York, 2004. doi: 10.1145/1186155.1186173
- [55] A. Roudaut, A. Karnik, M. Löchtfeld, and S. Subramanian. Morpheus: Toward high “shape resolution” in self-actuated flexible mobile devices. In *Proc. CHI, CHI '13*, p. 593–602. ACM, New York, 2013. [Open Access](#). doi: 10.1145/2470654.2470738
- [56] K. A. Satriadi, J. Smiley, B. Ens, M. Cordeil, T. Czauderna, B. Lee, Y. Yang, T. Dwyer, and B. Jenny. Tangible globes for data visualisation in augmented reality. In *Proc. ACM CHI, 2022* (to appear).
- [57] J. Sauro and J. R. Lewis. Average task times in usability tests: What to report? In *Proc. CHI, CHI '10*, p. 2347–2350. ACM, New York, 2010. doi: 10.1145/1753326.1753679
- [58] S. Schkolne, H. Ishii, and P. Schroder. Immersive design of DNA molecules with a tangible interface. In *Proc. Visualization*, pp. 227–234. IEEE Computer Society, Los Alamitos, 2004. [Open Access](#). doi: 10.1109/VISUAL.2004.47
- [59] V. Schwind, P. Knierim, N. Haas, and N. Henze. Using presence questionnaires in virtual reality. In *Proc. CHI*, pp. 1–12, 2019.
- [60] M. Sereno, S. Gosset, L. Besançon, and T. Isenberg. Hybrid Touch/Tangible Spatial Selection in Augmented Reality. *Computer Graphics Forum*, 41(3):403–415, June 2022. doi: 10.1111/cgf.14550
- [61] J. Smiley, B. Lee, S. Tandon, M. Cordeil, L. Besançon, J. Knibbe, B. Jenny, and T. Dwyer. The made-axis: A modular actuated device to embody the axis of a data dimension. *Proc. ACM Hum.-Comput. Interact.*, 5(ISS), nov 2021. [Open Access](#). doi: 10.1145/3488546
- [62] P. Song, W. B. Goh, C.-W. Fu, Q. Meng, and P.-A. Heng. WYSIWYF: Exploring and annotating volume data with a tangible handheld device. In *Proc. CHI*, pp. 1333–1342. ACM, New York, 2011. [Open Access](#). doi: 10.1145/1978942.1979140
- [63] M. Spindler, S. Stellmach, and R. Dachselt. PaperLens: Advanced magic lens interaction above the tabletop. In *Proc. ITS*, pp. 69–76. ACM, New York, 2009. [Open Access](#). doi: 10.1145/1731903.1731920
- [64] M. Spindler, C. Tominski, H. Schumann, and R. Dachselt. Tangible views for information visualization. In *ACM International Conference on Interactive Tabletops and Surfaces, ITS '10*, p. 157–166. ACM, New York, 2010. doi: 10.1145/1936652.1936684
- [65] I. E. Sutherland. The ultimate display. In *Proc. IFIP Congress*, pp. 506–508, 1965.
- [66] F. Taher, J. Hardy, A. Karnik, C. Weichel, Y. Jansen, K. Hornbæk, and J. Alexander. Exploring interactions with physically dynamic bar charts. In *Proc. CHI, CHI '15*, p. 3237–3246. ACM, New York, 2015. doi: 10.1145/2702123.2702604
- [67] T. Tsujita, K. Sase, X. Chen, M. Tomita, A. Konno, M. Nakayama, A. Nakagawa, K. Abe, and M. Uchiyama. Development of a surgical simulator for training retraction of tissue with an encountered-type haptic interface using mr fluid. In *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 898–903. IEEE, 2018. doi: 10.1109/ROBIO.2018.8665269
- [68] S. Ullrich and T. Kuhlen. Haptic palpation for medical simulation in virtual environments. *IEEE Trans. Vis. Comp. Graph.*, 18(4):617–625, Apr. 2012. doi: 10.1109/TVCG.2012.46
- [69] T. van Reimersdahl, F. Bley, T. Kuhlen, and C. H. Bischof. Haptic rendering techniques for the interactive exploration of CFD datasets in virtual environments. In *Proc. Workshop on Virtual Environments*, pp. 241–246. ACM, New York, 2003. [Open Access](#). doi: 10.1145/769953.769981
- [70] G. R. VandenBos, ed. *Publication Manual of the American Psychological Association*. American Psychological Association, Washington, DC, 6th ed., 2009.
- [71] C. Vavra. Tips for improving safety, roi for collaborative robots: collaborative robots are becoming more common on the plant floor. deciding if they should be used requires considering safety and business goals and return on investment (roi). see new safety guidance in a technical specification, iso/ts 15066: 2016 robots and robotic devices—collaborative robots. *Control Engineering*, 63(3):26–30, 2016.
- [72] E. Vonach, C. Gatterer, and H. Kaufmann. Vrrobot: Robot actuated props in an infinite virtual environment. In *2017 IEEE Virtual Reality (VR)*, pp. 74–83. IEEE, 2017. doi: 10.1109/VR.2017.7892233
- [73] X. Wang, L. Besançon, M. Ammi, and T. Isenberg. Understanding differences between combinations of 2d and 3d input and output devices for 3d data visualization. *International Journal of Human-Computer Studies*, p. 102820, 2022.
- [74] C.-M. Wu, C.-W. Hsu, T.-K. Lee, and S. Smith. A virtual reality keyboard with realistic haptic feedback in a fully immersive virtual environment. *Virtual Reality*, 21(1):19–29, 2017. [Open Access](#). doi: 10.1007/s10055-016-0296-6
- [75] L. Xie, A. N. Antle, and N. Motamedi. Are tangibles more fun? Comparing children's enjoyment and engagement using physical, graphical and tangible user interfaces. In *Proc. TEI*, pp. 191–198. ACM, New York, 2008. [Open Access](#). doi: 10.1145/1347390.1347433
- [76] K. Yamaguchi, G. Kato, Y. Kuroda, K. Kiyokawa, and H. Takemura. A non-grounded and encountered-type haptic display using a drone. In *Proc. SUI*, pp. 43–46, 2016. [Open Access](#). doi: 10.1145/2983310.2985746