



City Research Online

City St George's, University of London

Citation: Neate, T., Maffra, S. A., Frier, W., You, Z. & Wilson, S. (2023). Using Mid-Air Haptics to Guide Mid-Air Interactions. Paper presented at the 19th IFIP TC13 International Conference, 28 Aug - 1 Sep 2023, York, UK. doi: 10.1007/978-3-031-42280-5_3

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/35694/>

Link to published version: https://doi.org/10.1007/978-3-031-42280-5_3

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Using Mid-Air Haptics to Guide Mid-Air Interactions

Timothy Neate¹, Sergio Alvares Maffra², William Frier³,
Zihao You¹, and Stephanie Wilson⁴

¹ Department of Informatics, King's College London, UK

² Department of Mathematics, King's College London, UK

³ Ultraleap Ltd., Bristol, UK

⁴ Centre for HCI Design, City, University of London, UK

Abstract. When users interact with mid-air gesture-based interfaces, it is not always clear what interactions are available, or how they might be executed. Mid-air interfaces offer no tactile affordances, pushing systems to rely on other modalities (e.g. visual) to guide users regarding how to interact with the interface. However, these alternative modalities are not always appropriate or feasible (e.g. eyes-free interactions), meaning that they are not easy to learn from touch alone. Despite the possibility of conveying contactless haptic information in mid-air through ultrasound phased arrays, this technology has been limited to providing feedback on user interactions. In this paper, we explore the feasibility of using mid-air haptics to guide gestures in mid-air. Specifically, we present approaches to guide the user's hand in cardinal directions, execute a hand gesture and navigate a 2D mid-air plane, which we tested with 27 participants. After, reporting encouraging results which suggest good accuracy and relatively low workload, we reflect on the feasibility and challenges of using haptic guidance mechanisms in mid-air.

Keywords: Mid-air haptics · guidance · gesture · mid-air interfaces

1 Introduction

Mid-air, touchless, gesture-based interfaces are increasingly common [29]. They are a means for users to interact with augmented and virtual reality (AR/VR) [70], vehicles [61], home environments [60] and public displays [49]. Research has considered how complex interactions, such as continuous control with interface elements [30] and text input [45], can be done in mid-air. This is often a challenge because mid-air interactions provide no tactile feedback, and visual and auditory feedback is often undesirable or insufficient (e.g. when driving [61, 62] or flying [26]). In response to this, researchers have turned to mid-air haptics. Mid-air haptic devices provide tactile sensations on the user's skin, without the need for physical connection with a user's body [11, 36, 53]. This is commonly accomplished by focusing ultrasound – sound > 20kHz and beyond the human hearing range – onto a user's hand [53]. Mid-air haptics have been used to provide feedback on mid-air, eyes-free control tasks when driving [32, 38], and in situations

where it is unhygienic to touch a device physically, e.g. medical settings [54] or public displays [34, 40].

A remaining challenge with mid-air interfaces is their lack of explicit information on how to use them as they are not, as Delamare et al. put it, *self-revealing* [14]. Mid-air interactions are innately ambiguous as they have no visual or physical representation. The available interactions, and how they should be executed, are not clear [56]. Currently, no work has considered how these cues might be provided through mid-air haptics, which are increasingly considered as an output modality for mid-air interactions. In this paper, we suggest that information about how to interact with a mid-air interface might be provided by mid-air haptics, addressing the potential ambiguity of mid-air interfaces. We explore how mid-air haptic sensations can be used to guide users in specific interactions, namely directional movement, hand gestures and path-based interactions. Our key contributions are:

1. Methods for mid-air haptics that guide directional movement, hand gestures and path-based interaction from users
2. An investigation of the efficacy and user experience of the proposed methods

2 Background

2.1 Mid-Air Interactions and Guidance

A key challenge for touchless mid-air interactions is the lack of tactile information as users complete gestures or control tasks. It is therefore not clear what gestural interactions are available to a user, or how successfully they are interacting. Information about possible interactions, or the user’s ongoing interactions can be provided in visual form. However, this can add cognitive load to the interaction and is undesirable in many eyes-free contexts (e.g. driving [32]).

A common method for supporting users in understanding the ways they might interact with a mid-air interface is to tell them explicitly what interactions are available. For instance, some systems provide textual or spoken instructions about the relation between gestures and outcomes [62]. Others provide training for learning gestural input [2], or methods to simply suggest that mid-air interaction is possible in the first place [27]. Explicit tuition and feedback on gestures has also been considered in language therapy [55], but with the goal of communicating with others gesturally, rather than with a system. A body of research has also investigated what gestures are intuitive to people. Elicitation studies (see [66] for a systematic review), where participants respond via gesture to a given stimulus and preferences are aggregated, have been used to understand expected gestures in various contexts [16, 68, 72]. Research has also considered how we might best inform users when their interactions are out of the range of a sensor system. For instance, Morrison et al. [50] use audio to inform users when their body is out of range of a computer vision system.

Research has also been conducted to explore the notion of providing *guides* for users to support their interactions – e.g. giving information about the path

which the user must traverse to complete a gesture. All prior approaches we are aware of are visual or offer only binary haptic feedback (e.g. off/on). OctoPocus [5] provides 2D visual indicators of the next steps in a given path to show where the user should move their gesture next in 2D space. Similarly, Delamare et al. [15] consider how these approaches might support interactions in 3D space. Anderson et al. [4] considered, through their ‘YouMove’ system, how users might record and then learn custom gestures for full-body interaction with augmented reality mirrors. Visual information has also been used to support users in continuous mid-air interactions by showing the next points in a gesture path in AR [10]. Freeman et al. [19] explored a feedforward method for conveying what interactions are possible with a mid-air interface via wrist-worn rhythmic vibration information and light displays. Finally, Lopes et al. [42] considered the concept of dynamic affordance, i.e. providing feedback on possible interactions to users via muscle stimulation.

Freeman et al. [20] augmented a mid-air haptic device with LEDs. The authors then showed that the combination of mid-air haptic and visual cues could help guide the user to a given position above the device. Van den Bogaert et al. [6] argued that using mid-air alone could convey information to the users *before* they executed an action. For instance, they suggested using mid-air haptics to produce a “force field” around potentially harmful home appliances (e.g. a hot oven). However, this idea was neither implemented nor tested. Taking the idea one step further, Brown et al. [9] showed that associations between patterns and given meanings could be conveyed through mid-air haptic patterns, which suggests the feasibility of employing mid-air haptics as a guidance mechanism. However, the guidance itself is something that remains to be implemented and tested – an endeavour that we undertake with the present study.

2.2 Mid-Air Haptics

Mid-air haptics have been considered as a means for providing tactile feedback without the need for physical contact with an interface. This has involved the use of air jets [64], electric arcs [63] and focused ultrasound [24, 36]. Focused ultrasound approaches have been commercialised, based on the work of Carter et al. [11], in the form of [Ultraleap](#) (formerly Ultrahaptics). Typically, these systems use optical hand-tracking and an array of ultrasonic transducers to produce focused ultrasound on the user’s palm, which induces a tactile sensation. Localisation and movement of such induced haptic stimuli have been demonstrated [71]. Building on this, modulation techniques have been developed to render 2D planar patterns [22, 31, 57], as well as 3D objects [37, 41, 47].

Due to its relative maturity in this field, we consider focused ultrasound as a means to convey mid-air haptic stimuli in the remainder of this section and in our work. Several researchers have explored the possibilities and limitations of mid-air haptics using ultrasound. Wilson et al. [71] reported that participants could localise mid-air haptics focus points on their hand within 8.5mm accuracy, finding motion perception best when the focus point travelled longer distances

across the hand. Further studies [39, 52] investigated oblique motion (e.g. diagonally across the hand), finding that it was easier to recognise cardinal (i.e. left/right/up/down) rather than oblique motion. Hajas et al. [31] showed that tracing a given 2D pattern with a slowly moving focal point on a user’s hand produced better recognition performance than tracing the pattern rapidly and repeatedly [22], or with several focal points [57].

Applications of ultrasound mid-air haptics are diverse. For instance, the technology has been used for in-vehicle infotainment systems [25, 32, 38], to add haptic properties to objects in VR/AR [43, 54], for interaction with public displays [40], and to augment audiovisual media [1]. In all these scenarios, mid-air haptics are employed as feedback mechanisms, using haptic information to convey the state of the interface following user actions. The presence of a feedback mechanism improves the user’s perception of control in their action (i.e. sense of agency) [12, 17]. However, it remains unclear whether mid-air haptics could help users in planning or learning a given action. Indeed, using mid-air haptics as a *guidance* mechanism has received little attention within HCI and haptic communities.

3 Mid-Air Haptic Guidance Stimuli Development

In this paper, we consider the use of mid-air haptics to guide users or suggest given actions. In doing this, we focused on two main types of interaction with mid-air interfaces:

- **Directional movement** in a 2D plane, which is common for the control of cursor-like elements and mid-air widgets – e.g. sliders, knobs;
- **Gestures**, due to the large number of mid-air interfaces which use gestures as input.

We began with the assumption that there is a finite number of directions and gestures that the user can use with a system. In future studies, larger and more complex sets of gestures might be considered via approaches such as elicitation (see [73] and Section 6). We developed mid-air haptic stimuli for conveying a limited number of directional, gestural and path-based interactions through 3 rounds of piloting with 8 researchers and students in City, University of London (4 male, 4 female), along with various intermediary piloting to tune parameters, also including self-piloting (by the authors of this paper). All mid-air haptic stimuli were developed using the Ultrahaptics Evaluation Kit (UHEV1) developed by Ultraleap (formerly Ultrahaptics). We also developed a custom Python front-end that would allow us to change the parameters of the stimuli for piloting, which linked to the Ultrahaptics SDK for driving the ultrasonic array (as in Figure 3). The Python front-end, and usage instructions, can be found [here](#).

3.1 Mid-Air Haptic Stimuli to Convey Direction

We started with the aim of investigating how to suggest a notion of direction to a user – i.e. that they should move their hand in a specific direction. This

is analogous to the visual guides provided in prior work [4, 5, 14]. Lerusso et al. [39, 52] have shown that motion is better perceived in cardinal directions, rather than oblique directions. We, therefore, considered cardinal direction as an initial focus. Building upon research such as [31, 71], we chose to explore stimuli which traverse the whole hand as this research, and our initial piloting, found that larger movement differences are perceived better than smaller. Through piloting we also found that to convey a given direction, it was better to move the stimuli in the direction of intention – e.g. to convey ‘left’ it was better to move from right to left. This is analogous to ‘default’ rather than ‘inverted’ in videogame controls’ [23] and reverse scrolling, rather than natural, with a computer mouse. This is also similar to the methods used in similar work on (visual) guides [4, 5, 14], e.g. showing visual movement in the direction that the user should travel. The speed of the focus point’s traversal and its frequency were also piloted until we found a good ‘average match’ for the population. We began with traversal values around 100ms, up to around 1000ms (maxima and minima from [71]), finding 500ms to be the best fit overall. For modulation frequency, we arrived at 200Hz. The final stimuli are shown in Figure 1.




Fig. 1. All stimuli used in the experiment relative to a user’s right hand. The mid-air sensation interpolated across all points on the user’s palm – e.g. ‘Left’ interpolates from $[1, 0]$ to $[-1, 0]$. Parameters set after piloting were: sensation duration of 500ms; modulation frequency of 200Hz; a distance of 50mm (most effective distance travelled on the palm from piloting – we did not fit ‘per participant’). Darker dots indicate later in time (e.g. left goes from centre to far left).

3.2 Mid-Air Haptic Stimuli to Convey Gesture

We then explored specific gestures. Our aim here was to understand if we could convey the gestures available to a user, given some known qualities of the interface and its interactions – e.g. cardinal direction to convey translation (movement of a UI element in a given direction). In addition, we saw the potential for not

just suggesting gestures available but using guidance information as a way to train a user in completing specific physical motions. This might be applied to learning a new mid-air interface and its inputs, but also in other digital, or remote learning environments – for instance guiding a user towards the correct note on a piano (e.g. extending [35]’s mid-air piano), or for specifically *training gestures*. A motivator for this work was also the importance of gestures in functional communication for people with complex communication needs. Gesture is a vital pathway for people who face challenges with verbal output. Augmented communication technologies should consider this, but currently do not [13]. Finally, on training gestural response for access needs, Marshall *et al.* [46] and Roper *et al.* [55] have also considered how technology can train gestures for those with complex communication needs.

We selected four commonplace gestures – “Swipe Left”, “Swipe Right”, “Tap”, “Pinch” – based on their frequency of use in mid-air interactions literature – including those shortlisted by Van den Bogaert *et al.* [7] in their elicitation study, and those explored by Ameer *et al.* [3] and Marin *et al.* [44]. The final gestures are shown in Figure 2. Guidance stimuli for these gestures were based on the [Ultraleap Sensation Editor](#) demos. Again, these stimuli were fine-tuned through piloting with participants. Stimuli are shown in Figure 1.




Fig. 2. Gestures used in the study – ‘Swipe Left’, ‘Swipe Right’, ‘Tap’, ‘Pinch’. Images use ‘motion blur’ to indicate movement.

3.3 Mid-Air Haptic Stimuli to Convey Paths

Finally, we investigated conveying 2D paths to users as a form of guidance (i.e. the computer continuously ‘controlling’ a user in space). This is analogous to Burny *et al.* [10]’s work – previously discussed – who use visual information to support users in completing paths. We also sought to understand if we might use this strategy to position users anywhere above a mid-air haptic board (in 2D space) so that they can begin interactions – this might be useful, for instance, for positioning a user’s hand correctly to interact with a specific part of the mid-air UI. As discussed in the Background section, previous work has considered how to position the user’s hand in the centre of the board and at the correct height for stimulation with static mid-air haptics and light-based information [20], however, we saw this as an opportunity for positioning the user at *any* point (in a 2D

plane) above the board. In piloting, we explored how we might support users in navigating paths. We found that a fruitful approach was to use the same stimuli as for cardinal direction to move a user across a discretized (3x3) grid (as shown in Figure 3). We implemented it so that the cardinal direction stimuli discussed previously were linked to the keyboard – up to the ‘up’ key, down to the ‘down’ key, etc – and controlled by a researcher. The directional stimuli played repeatedly and looped when the keys were held down, then stopped immediately (e.g. half-way through a stimuli) as soon as the key was released. We chose to use a human-controlled interaction to explore this proof of concept, rather than implement a control algorithm to guide the user on the path. We accept that this might have introduced a human-based bias into the interaction, but chose to do this after finding that error handling (e.g. a user straying from the path) introduced confounds beyond the scope of our present. We also chose to do this to explore human-in-the-loop dynamics for remote interaction (further delineated in the discussion).

4 Study Method

We undertook a study to determine the potential of our prototype mid-air stimuli for guidance addressing the following research questions:

- RQ1** How intuitive are the mid-air haptic stimuli for guiding *directional movement*?
 - RQ1.1** How accurately can we guide users with mid-air haptic stimuli in specific directions *without prior experience*?
 - RQ1.2** How accurately can we guide users with mid-air haptic stimuli in specific directions *with prior experience* of the stimuli?
 - RQ1.3** What is the user’s *subjective experience* when guided in directions using mid-air haptics?
- RQ2** How intuitive are the mid-air haptic stimuli for guiding *gestural responses*?
 - RQ2.1** How accurately can we guide users with mid-air haptic stimuli to perform specific gestures *without prior experience*?
 - RQ2.2** How accurately can we guide users with mid-air haptic stimuli in performing specific gestures *with experience* of the stimuli?
 - RQ2.3** What is the user’s *subjective experience* of being guided to perform specific gestures with mid-air haptic stimuli?
- RQ3** How effective are the developed stimuli for *continuous, real-time guidance* along a given path?
 - RQ3.1** How accurately can users *follow a given mid-air path* when guided by the developed mid-air haptic stimuli?
 - RQ3.2** What is the *subjective experience* of being guided along a mid-air path with guidance from the mid-air haptic stimuli?

4.1 Participants

We recruited a convenience-based student sample of 27 participants. No participants had been involved in the earlier piloting stages. All participants self-reported good or corrected visual acuity and no challenges with experiencing sensations in their hands. 19 identified as male and 8 as female. Ages ranged from 20–33 (average = 25.4; SD = 4.0). All participants, except one, were confident with technology – agreeing or strongly agreeing with the statement “*I am technologically proficient – e.g. I use tablets, smartphones, computers regularly*”. Only 2/27 participants had experience with mid-air haptics before the study.

4.2 Study Procedure

The study was conducted in a usability testing room at King’s College London (Figure 3). The study was run by one experimenter who had practised using the developed stimuli (discussed in the previous section), and the experimental technique, through piloting as described previously at City, University of London. Participants were consented by the experiment coordinator on the day through an informed consent procedure in line with our ethical approval process, with authorisation from the ethics committee at King’s College London.

We used the Evaluation Kit developed by Ultraleap – an ultrasound phased array which provides tactile sensations in mid air (pictured top left, Figure 3). Participants sat in front of the ultrasound phased array board with a large TV display (Samsung QM55R-B) directly in front of them for visual information when required. Each participant was instructed to place their hand 20cm above the board. Stimuli were played from a laptop (Lenovo IdeaPad 3) connected to the board, through a custom front-end interface built in Python. Participants were positioned so that they had no view of the experiment coordinator’s laptop display. The experiment coordinator also used the directional keys on the laptop, again linked to the stimuli. Response data were collected via a spreadsheet by the experiment coordinator (on a separate laptop – a MacBook Pro) and subjective data were captured via an iPad. Pauses were introduced between trials and participants were free to rest when they wanted to mitigate fatigue.

Part 1.1: Conveying Directional Movement: No Prior Experience with Developed Stimuli To answer RQ1, in Part 1 of the study, we aimed to find out if participants could determine cardinal direction from the developed mid-air haptic stimuli, with no prior experience with the stimuli (RQ1.1). We explained to participants that the device would produce a sensation and they should respond in the way they saw appropriate. On the display, we provided a visual aid showing the four possible responses the device was trying to convey to them – 4 videos of a hand moving in the given directions (up/down/left/right). Each participant received all four direction stimuli once (i.e. 4 trials), but in a random order. After each sensation was played, the participant was asked to respond by moving their hand in the direction they thought the stimuli was

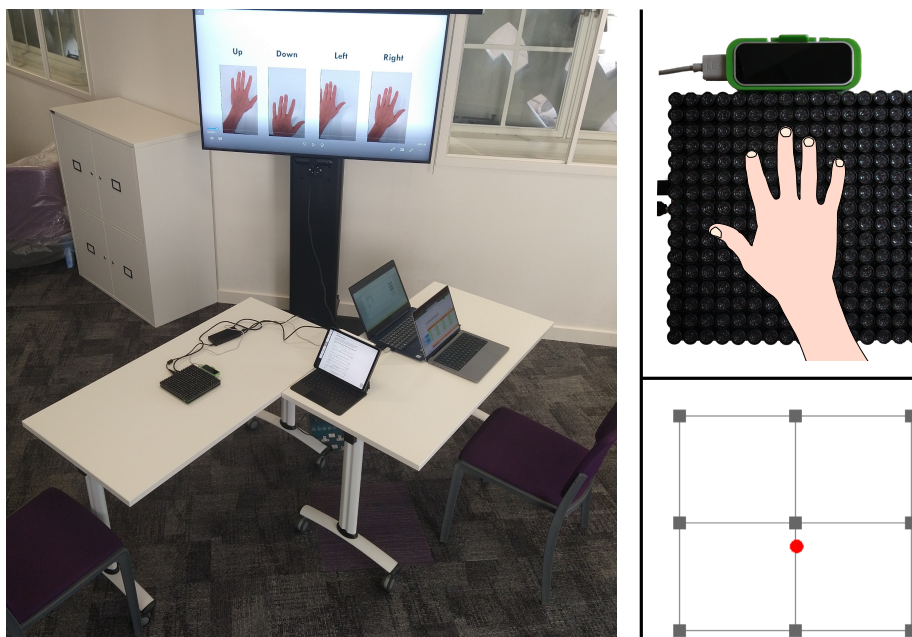


Fig. 3. Study setup shown left. Participants were seated on the left chair, the experiment coordinator on the right. Top right shows the UltraLeap UHEV1 board (later referred to as ‘the board’) used in the study, with example of hand position overlaid. Bottom right shows the UI for guiding the participant in Part 3 of the study – the nine grey points correspond to points in 2D space above the board – top left being the top left corner, for instance. The red point shows the position of the centre of the participant’s hand.

conveying, and verbally confirming the direction – e.g. “*up*”. The experimental coordinator then entered the response into the spreadsheet.

Part 1.2: Conveying Directional Movement: With Knowledge of Developed Stimuli After completing Study Part 1.1, regardless of success in determining the stimuli, participants were instructed with the ‘correct’ correspondence between given stimuli and intended response: they were given the chance to experience each stimulus again and the experiment coordinator told them its correspondence. To answer RQ1.2, participants then undertook 10 trials in which they were tasked with determining the cardinal direction from a given stimulus. Stimuli were played randomly across the four cardinal directions. As in Part 1.1, participants gave their response through movement and verbal confirmation.

Part 1.3: Subjective Response to Conveying Directional Movement To answer RQ1.3 – the users’ subjective experience – we administered a NASA

TLX to understand workload, and asked participants to rate how hard/easy they found it to understand each cardinal direction on a 5-point Likert scale, ranging from *Strongly Disagree* to *Strongly Agree* – e.g. *I found it easy to understand when the device was indicating to go ‘up’*.

Part 2.1: Conveying Gestures: No Prior Experience with Developed Stimuli To answer RQ2.1 we adopted the same strategy as was employed in Part 1.1, but this time we asked participants to respond from a set of common gestural responses (rather than from a set of directions). Again, the chosen four gestures – ‘Swipe Left’, ‘Swipe Right’, ‘Tap’, ‘Pinch’ were shown on the display and participants were asked to respond without prior knowledge of the gesture-stimuli correspondence.

Part 2.2: Conveying Gestures: With Prior Experience of Developed Stimuli As in Part 1.2, regardless of their success in determining the gestures, participants were instructed with the ‘correct’ correspondence between given stimuli and intended gestures. We then conducted 10 randomised trials in a manner similar to Part 1.2 to answer RQ2.2.

Part 2.3: Subjective Response to Conveying Gestures As before, subjective response (RQ2.3) was determined via NASA TLX and a Likert questionnaire with 5-point Likert scale questions to understand the ease of identifying each gesture.

Part 3.1: Path Tracing We created four paths in a 2D plane across 9 points. Paths were analogous to a pattern unlock screen and the task analogous to teaching a user how to unlock a mid-air interface. Four paths, to be given to all participants, were chosen by choosing a random start and end coordinate four times (using a random number generator). Final generated paths are shown in Figure 8.

To address RQ3.1, the experiment coordinator ‘guided’ the participant around each of the four paths indicated in Figure 8. The experimental coordinator viewed the participant’s hand position on their laptop screen (see Figure 3, bottom right) and then used the mid-air haptic stimuli to guide the users’ hand in the four cardinal directions, towards a target. The experiment coordinator operated the stimuli with the arrow keys which corresponded to the intended direction of travel. The stimuli were the same as in Part 1 (e.g. focus point traversing from right to left to indicate left travel), however, they were repeated as long as one of the arrow keys was held down. If the participant deviated from the path, the experiment coordinator aimed to bring the participant back to the point from which they diverged. Analysis was conducted by plotting each path and classifying it as ‘correct’ (all edges and points followed), ‘partially correct’ (over 50% of edges followed with correct start/end points) or incorrect (anything else), with 2 independent coders classifying the data. Plots of combined paths (i.e. all participants) were also made for aggregated analysis.

Part 3.2: Subjective Response to Path Tracing As with the previous parts of the study, we administered a NASA TLX and a 5-point Likert questionnaire, asking about the ease of following the path, after the four paths were completed.

5 Results

5.1 Part 1.1: Conveying Directional Movement: No Prior Experience

As shown in Figure 4, given no prior experience of cardinal direction stimuli, participants were able to correctly identify the stimuli-direction correspondence in a total of 95 out of the 108 trials (4×27) – 88% accuracy, with 2 no movement or a negative verbal response (accounting for the rows of Figure 4 not summing to 100%). A confusion matrix is shown in Figure 4 (left) – suggesting the most common confusion was right/left and up/down.

		Observed User Response							
		Down	Left	Right	Up	Down	Left	Right	Up
Actual output	Down	89%	0%	0%	7%	94%	1%	0%	4%
	Left	0%	89%	7%	4%	0%	93%	5%	2%
	Right	0%	11%	85%	0%	2%	3%	92%	3%
	Up	7%	0%	4%	89%	7%	7%	0%	87%
Part 1.1 (without prior knowledge)					Part 1.2 (with prior knowledge)				

Fig. 4. Results of Parts 1.1 and 1.2, comparing the mid-air haptic stimuli for cardinal direction and user responses.

A χ^2 test of independence was conducted to determine the relationship between stimuli and observed user response to stimuli without prior experience. There was a significant relationship – Pearson- χ^2 (9, $N=106 = 240.125$, $p < 0.001$), with a strong effect size ($\Phi = 1.50$) suggesting that users could identify the stimuli for cardinal directions significantly better than chance.

5.2 Part 1.2: Conveying Directional Movement: with Experience

With experience about how the stimuli related to the cardinal directions, participants were able to correctly determine the correspondence in 247 out of a total of 270 trials (10×27) – i.e. 91% accuracy). The random seed generated ‘Down’ 72 times, ‘Up’ 76 times, ‘Left’ 59 times and ‘Right’ 63 times. A confusion matrix is shown in Figure 4 (right). A χ^2 test of independence was conducted to determine the relationship between stimuli and user response *with* experience of the stimuli. There was a significant relationship (Pearson- χ^2 (9, $N=270 = 639.784$, $p < 0.001$), with a strong effect size ($\Phi = 1.54$), suggesting users could determine the correct response significantly better than chance.

		Observed User Response							
		Part 2.1 (without prior knowledge)				Part 2.2 (with prior knowledge)			
Actual output		Pinch	Swipe Left	Swipe Right	Tap	Pinch	Swipe Left	Swipe Right	Tap
	Pinch	58%	12%	4%	27%	73%	10%	3%	13%
	Swipe Left	7%	74%	11%	7%	4%	88%	3%	5%
	Swipe Right	0%	11%	85%	4%	4%	8%	80%	7%
	Tap	28%	12%	0%	60%	23%	3%	2%	72%

Fig. 5. Confusion matrix for Parts 2.1 and 2.2. of the study, comparing the mid-air stimuli for gestures and the user responses.

Part 1.3: Subjective Response Post-Part 1 Likert data results are shown in pink in Figure 6. The cardinal directions were universally perceived as “*easy to understand*” – all with a median value of four. A Friedman test across the four conditions suggests no significant difference in reported ease between cardinal directions: $\chi^2(3) = 1.0$, $p = 0.80$. As shown in Figure 7 NASA TLX for this was 29.2.

5.3 Part 2.1: Conveying Gesture Without Prior Experience

For the response to gesture-based stimuli, with no prior experience participants correctly determined the relationship between the stimuli and the intended gesture in 73 out of 105 trials (3 missing datapoints where participants responded that they did not know) – a total accuracy of 68%. A confusion matrix is presented in Figure 5. The most common mismatch was between the ‘Tap’ and ‘Pinch’ gestures. As with Part 1.1, a χ^2 test of independence was detected to determine if participant could determine stimuli without prior information better than chance. A significant relationship was (Pearson- χ^2 (9, N=105 = 123.61, $p < 0.001$), with a strong effect size ($\Phi = 1.09$).

5.4 Part 2.2: Conveying Gesture with Prior Experience

With experience about how the stimuli related to the intended gestures, the participants’ response accuracy increased to 79% – 212 out of 270 trials, with 2 missing data points. The random seed generated ‘Swipe Left’ 77 times, ‘Swipe Right’ 71 times, ‘Pinch’ 60 times and ‘Tap’ 62 times. Again, as shown in Figure 5, the most common confusion was ‘Tap’–‘Pinch’ (23%), however, this was smaller for ‘Pinch’–‘Tap’ (13%). There was a significant relationship (Pearson- χ^2 (9, N=268 = 425.9, $p < 0.001$), with a strong effect size ($\Phi = 1.26$), showing that participants could determine the gestures substantially better than chance.

5.5 Part 2.3: Subjective Experience of Gestural Response

Results of Part 2.3 are shown in Figure 6. ‘Swipe Left’ and ‘Swipe Right’ gestures were generally perceived as easy to understand, with median ratings of 4. ‘Pinch’

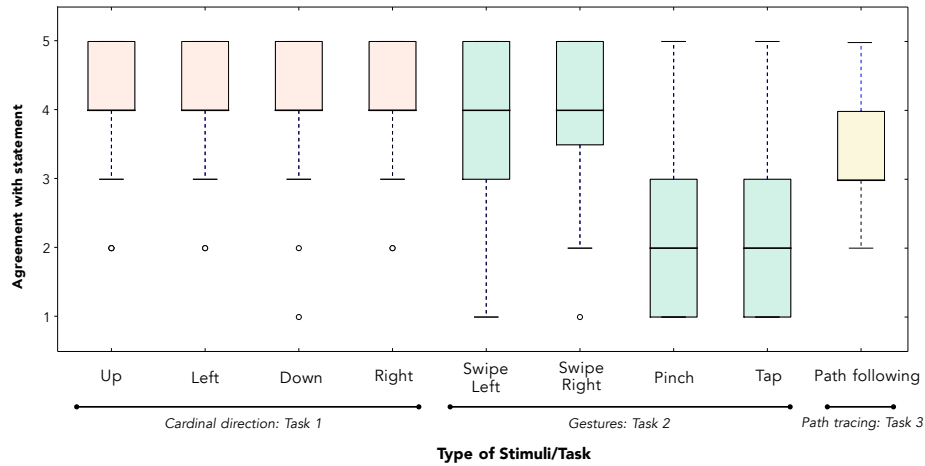


Fig. 6. Agreement with the statement “I found it easy to understand when the device was indicating given [direction/gesture/path], where 1 is “Strongly Disagree” and 5 is “Strongly Agree”.

and ‘Tap’ gestures were perceived as more challenging, with medians at 2, with only 5 (‘Tap’) and 4 (‘Pinch’) participants agreeing that this was easy. The four stimuli were compared via Friedman test, with a significant difference across the four gestures – $\chi^2(3) = 44.08$, $p < 0.001$. Post-hoc Wilcoxon Signed-Rank tests to determine differences in understanding between gestures were conducted. ‘Tap’ was perceived as significantly harder to understand than a ‘Swipe Left’ ($Z = -3.643$, $p < 0.001$) and ‘Swipe Right’ ($Z = -3.885$, $p < 0.001$). Similarly, ‘Pinch’ was significantly harder to understand than ‘Swipe Left’ ($Z = -4.016$, $p < 0.001$) and ‘Swipe Right’ ($Z = 4.274$, $p < 0.001$). There was no significant difference when comparing ‘Pinch’ and ‘Tap’ ($Z = -0.711$, $p = 0.477$), and also no significant difference between ‘Swipe Left’ and ‘Swipe Right’ ($Z = -1.134$, $p = 0.257$). As shown in Figure 7, NASA TLX for gestural stimuli was 40.37 – a higher workload than for the cardinal directions in Part 1 of the study.

	NASA-TLX Average						
	Mental Demand	Physical Demand	Temporal Demand	Performance	Effort	Frustration	Average
Part 1: Cardinal Direction	31.3	33.5	29.4	29.6	30.6	20.6	29.2
Part 2: Gesture	45.4	31.3	26.7	52.8	49.3	36.9	40.4
Part 3: Path Tracing	37.8	40.9	35.0	34.4	43.0	31.5	37.1

Fig. 7. Combined NASA-TLX averages for all 27 participants.

	Path 1	Path 2	Path 3	Path 4	Overall
Time (seconds)	26.1	9.5	37.4	15.1	22.0
Time (Standard Deviation)	15.9	15.9	16.1	16.1	16.0
Correct (%)	66.7	77.8	74.1	81.5	75.0
Partially correct (%)	22.2	3.7	25.9	7.4	14.8
Incorrect (%)	11.1	18.5	0.0	11.1	10.2

Table 1. Time to complete and the standard deviations are shown, along with the correctness of the paths overall.

5.6 Part 3.1: Path Tracing

Success of path completion was determined by two coders. Coder agreement was high (96.4%). Inter-rater Reliability (IRR) was calculated (Cohen’s $\kappa = 0.85$, $p < 0.001$) – i.e. ‘almost perfect’ agreement. The outcomes are shown in Table 1. Overall, the coders deemed that 75% of paths were completed correctly, 14.8% of paths were partially correct and 10.2% of paths were incorrect.

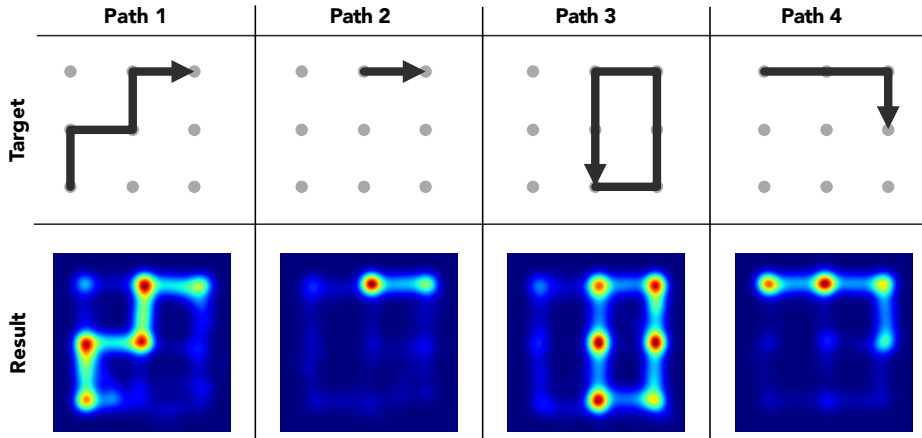


Fig. 8. Top shows the four target paths. Bottom shows combined paths of 27 participants as a heatmap for each path.

5.7 Part 3.2: Subjective Experience of Path Tracing

As shown in Figure 6, the path tracing task was judged as more challenging than determining cardinal directions in Part 1, but less challenging than determining non-directional gestures (‘Tap’, ‘Swipe’) in Part 2. This is also the same for the TLX score which, as seen in Table 7, falls between the two, with an overall score of 37.1.

6 Discussion, Applications and Limitations

6.1 Discussion of Directional Guidance

Regarding RQ1, while the use of mid-air haptic stimuli to guide directional movement has not been reported in the literature, the results of Part 1 suggest that this is possible and feasible. Participants were able to determine which direction the system was guiding them to move in without prior experience of the relationship between stimuli and direction (RQ1.1). Then, after experiencing the relationships between stimuli and directions, this accuracy increased (RQ1.2) – indicating that training is effective. As discussed in the Background section of this paper, Wilson et al. [71] conducted a study of directional recall – e.g. which direction users believed the *haptic sensation* was moving in. Importantly, we differentiate our work from this by considering which direction the stimuli is indicating the user *should* move in – analogous to the visual guidance approaches previously discussed (c.f. [4, 5, 14]). Given that our results are broadly comparable – although with varying setups – suggests that using mid-air haptics for directional guidance is reasonably intuitive to a user. In comparison with other work which has considered the determination of cardinal direction, our accuracy figures were higher than studies using finger-based sensations (e.g. [59]), likely due to the larger surface area of the whole hand, and prior work which found that larger distances of movement of sensation across the hand result in more accurate recall [71]. Our findings slightly contrast with Wilson et al.’s [71] who, when asking users to differentiate between longitudinal (up/down) and transverse motion (left/right), found that transverse was more accurately recalled. We found no significant difference with our methods.

Considering perceived workload and user experience (RQ1.3), our NASA TLX data suggest that workload for cardinal direction is relatively low compared to other computer-based tasks, and in the bottom 20% of all tasks measured by TLX [28]. To build upon our work, future studies might consider expanding mid-air haptic guidance to support users in learning gestures/inputs which incorporate oblique and curved motion, as in [39]. Our approaches could support users in learning potential ways to interact and which interactions might be available, rather than relying entirely on visual or auditory representations of gestures. This might be important UIs where eyes-free interactions are important – e.g. automotive UI [25, 32, 62]. Beyond supporting the guidance of directional movement-based interactions, these forms of mid-air haptic stimuli might be used to support other interaction paradigms such as error handling. Consider a scenario such as a user scrolling through a horizontal list (e.g. of songs) in an automotive UI through mid-air hand movements. When the user reaches the end of the list, we might go beyond considering just feedback methods (e.g. vibrating to tell the user there are no more interactions they can perform), to consider how we might design feedforward-based guidance interactions. We might, for instance, suggest what the user could do to *remedy their situation* – e.g. suggest an ‘escape gesture’ or that they should go in a given direction. We might also provide users with information about how much there is in each direction be-

fore the scroll – analogous to [8]’s sonically enhanced scrollbars. This might be implemented by varying the intensity of the stimuli – e.g. by changing the frequency of the modulation, or by using movement speed as a proxy for intensity, as explored by Freeman *et al.* [21].

6.2 Discussion of Gestural Guidance

Regarding RQ2, guidance of gestural response via mid-air haptics has not been considered in the literature. Our study suggests that while directional gesture stimuli (‘Swipe Left’, ‘Swipe Right’) were effective, ‘Pinch’ and ‘Tap’ were less successful, yet still mostly determined correctly and significantly better than chance. When trained, this improved substantially – indicating that training is effective. There could be a number of reasons for this confusion. One reason is that the ‘average’ point of the stimuli was similar – and while ‘Pinch’ involved transverse motion (away from the centre), users were less able to make this differentiation. This might also relate to the *primacy effect*, a psychological effect where the first stimuli is remembered most prominently. In our case, as the starting points (in time) of the pinch and tap stimuli (see Figure 1) are the same (the centre), it means that these are most prominently felt and remembered at the end of the stimuli. This concurs with research by Wang *et al.* [69], who report a similar effect for haptic stimuli. Generally, this might mean that we consider avoiding stimuli with similar average points across the duration of the stimuli or similar beginning/end points when using haptic stimuli to convey gestural responses to users. However, more studies are needed to make conclusive recommendations. The workload, as measured by NASA TLX, was higher for gestures than for cardinal directions and path tracing. However, it was well within the lower 50% of workload for computer tasks, and in the lower 30% of tasks generally – again, suggesting promise in terms of general user experience and feasibility.

In future work, our gesture stimuli might be improved by offering a greater diversity of duration and frequency. Future stimuli might be designed through elicitation studies similar to Vogiatzidakis and Koutsabasis [67], except that the *referent* – i.e. the stimuli the user must respond to [73] – would be the mid-air haptic stimuli and the users’ responses would be the basis for the analysis. Related to this, to complement use cases such as the training of gesture for people with severe language impairments (e.g. [46, 55]), more stimuli are needed, with more considered design. For instance, to support the 30 pantomimic gestures used in Marshall *et al.* [46] would require diverse stimuli. Moreover, some gestures would be challenging to support without more flexibility of the apparatus – gestures such as ‘Camera’ (see [46]) require a user to hold their hand to their face, with their thumb and index in front of their eye. Keeping one’s hand fixed in a 2D plane 20cm above the board makes this impractical. Moreover, mid-air haptics is limited by its planarity – present implementations (e.g. the board used in this study) affect a 2D plane; we cannot affect the top of the hand from below, or indeed closed elements (e.g. when the user closes their hands by connecting their thumb to fingers). To overcome these obstacles, we might consider approaches

such as Howard et al. [33], who consider the physical actuation of mid-air haptic boards, or approaches with transducer arrays above and below a user’s hand (as in acoustic levitation work [48]).

6.3 Discussion of Path Tracing Guidance

The data suggest that our stimuli could guide users along a specified path. When guided by another user (the experiment coordinator) this was done with high accuracy, following the exact path 75% of the time and only failing in about 10% of cases. Our work, therefore, could extend approaches such as Freeman et al. [19] – we might not only be able to guide a user to the ‘sweet spot’ on a mid-air haptic board, but indeed any spot. This might support users in discovering UI elements – for instance, encouraging users to move towards a slider/knob to begin a control task. Moreover, we might be able to indicate affordances to the user while completing this control task – e.g. that a slider is ‘slidable’ – an interaction that could be relevant to in-car settings such as [32]. Considering workload, it was higher for the path tracing task than for the simple cardinal direction, but lower than for gestures. Contrasting this with other TLX studies, it would be within the lowest 30% of tasks in general [28] – suggesting feasibility. We do acknowledge the limitation, however, that participants’ experiences from Part 1 of this study might have supported them in completing the path-tracing tasks as the stimuli were similar and they had become more accustomed to the mid-air haptic stimuli. It is also notable that Physical Demand for path following was the highest recorded TLX dimension of all. This was likely caused by the fact that the paths required more time to complete and therefore required the user to elevate their hand for longer – resulting in greater physical demand. This is supported by the fact that ‘Temporal Demand’ was also the highest for this task.

Our work also shows that, beyond guiding a user to a specific point, we can move them along a specific path. Practically, this might have benefits such as supporting users in avoiding UI elements in mid-air interfaces, or guiding/teaching users to complete specific path-based interactions. For example, a mid-air unlock screen, or mid-air motions for teaching dancing, surgery, or even games that use mid-air control. Finally, we envision scenarios like the ones explored in this paper, where a user is being guided by another person (the experiment coordinator in this case). In some ways, this was a limitation to our study – control algorithms might have been used to guide the user along a path, or to a point. For example, in future implementations, one could envision that a recogniser, such as the \$1 uni-stroke recognizer [74], could be used to provide instant feedback on whether the user has performed the correct path/gesture. Human-driven control via our approaches, however, are also scenarios we might envision. For instance, we might imagine using mid-air haptics for remote physicality – e.g. work such as [65], which considers how remote physical contact might be embodied through technology, or Neate et al. [51] who express the importance of tangibility in accessible communication (and indeed its loss during the pandemic for those who require it most). Finally, we might consider how mid-air haptics could be used

for the remote control of objects (e.g. similar to [48, 58]) – for remote presence and remote manipulation tasks such as those discussed by Follmer *et al.* [18] through actuation of objects in their inFORM project.

6.4 Limitations

Despite the general efficacy of our mid-air haptic approach to mid-air guidance, we acknowledge limitations in our work. One limitation is the convenience sample used for the study. As we recruited from a student population, our sample was biased towards a younger, ‘tech-savvy’ demographic. Moreover, the sample was skewed towards males in the main study (but with an equal male/female split in the piloting). Future work seeking to build on our proof of concept should ensure a more diverse sample of participants. Another limitation is the low number of stimuli explored for conveying gestures – this is a somewhat constrained discoverability task, with a small number of potential interactions. Further stimuli might result in more challenges in the disambiguation of gestures. Finally – as discussed in the background – we are constrained by the limitations of the Ultraleap board. While these have progressed since the UHEV1 board used in the study, further technical innovations in mid-air haptics (e.g. resolution, tactile sensation strength), and other future haptic approaches might make our approaches more feasible.

7 Conclusion

Mid-air interactions are increasingly commonplace, and are more frequently augmented by mid-air haptics. This paper has explored how we might leverage mid-air haptics to address one of the key challenges with mid-air interfaces – the lack of knowledge about how we might interact with them. We iteratively developed stimuli to guide users and to provide information about how they might interact. Our methods resulted in relatively low mental workload – suggesting the approach is feasible. While our stimuli set was limited, and providing guidance for the increasing set of interactions available with mid-air interfaces will be a major challenge, we have presented what we believe to be an interesting approach for supporting users with more intuitive mid-air interfaces. We invite future researchers to consider new guidance approaches to interactions with mid-air haptics, and to consider how they might be integrated into interaction scenarios and complement other modalities to make mid-air interactions more intuitive.

8 Acknowledgements

We would like to thank the participants for their involvement in the study and the anonymous reviewers of this work for their detailed feedback. We acknowledge funding from City, University of London’s Research Development Fund’s Pump Priming Scheme.

9 Authorship Statement

TN led the work with input from SAM, WF and SW. TN, WF, SAM and SW envisioned the experimental design. SAM built the Python front-end experimental apparatus, with input from TN and WF. ZY carried out the study with guidance from TN. The analysis was undertaken by TN, with support from WF, ZY and SAM. The paper writing was led by TN, with input from all authors.

References

1. Ablart, D., Velasco, C., Obrist, M.: Integrating mid-air haptics into movie experiences. In: ACM TVX (2017)
2. Ackad, C., Clayphan, A., Tomitsch, M., Kay, J.: An in-the-wild study of learning mid-air gestures to browse hierarchical information at a large interactive public display. ACM (2015)
3. Ameer, S., Khalifa, A.B., Bouhlel, M.S.: Hand-gesture-based touchless exploration of medical images with leap motion controller. In: Proceedings of SSD. IEEE (2020)
4. Anderson, F., Grossman, T., Matejka, J., Fitzmaurice, G.: Youmove: Enhancing movement training with an augmented reality mirror. In: Proceedings of UIST (2013)
5. Bau, O., Mackay, W.E.: Octopocus: A dynamic guide for learning gesture-based command sets. In: Proceedings of UIST. p. 37–46. Association for Computing Machinery, New York, NY, USA (2008)
6. Van den Bogaert, L., Geerts, D., Rutten, I.: Grasping the future: Identifying potential applications for mid-air haptics in the home. In: Extended Abstracts of ACM CHI. ACM (2019)
7. Bogaert, L.V.d., Geerts, D.: User-defined mid-air haptic sensations for interacting with an ar menu environment. In: Proceedings of EuroHaptics. pp. 25–32. Springer (2020)
8. Brewster, S.A.: Using non-speech sound to overcome information overload. Displays **17**(3-4), 179–189 (1997)
9. Brown, E., Large, D., Limerick, H., Frier, W., Burnett, G.: Validating the salience of haptic icons for automotive mid-air haptic gesture interfaces (2021)
10. Burny, N., Kieffer, S., Magrofuoco, N., Perez Medina, J.L., Roselli, P., Vanderdonckt, J., Vas, A.: Feedup, feedback, and feedforward in curve mid-air 3d gestures (2018)
11. Carter, T., Seah, S.A., Long, B., Drinkwater, B., Subramanian, S.: Ultrahaptics: multi-point mid-air haptic feedback for touch surfaces. In: Proceedings of UIST (2013)
12. Cornelio Martinez, P.I., De Pirro, S., Vi, C.T., Subramanian, S.: Agency in mid-air interfaces. In: Proceedings of ACM CHI (2017)
13. Curtis, H., You, Z., Deary, W., Tudoreanu, M.I., Neate, T.: Envisioning the (in)visibility of discreet and wearable aac devices. In: Proceedings of ACM CHI (2023)
14. Delamare, W., Coutrix, C., Nigay, L.: Designing guiding systems for gesture-based interaction. In: Proceedings of EICS. EICS '15, New York, NY, USA (2015)
15. Delamare, W., Janssoone, T., Coutrix, C., Nigay, L.: Designing 3d gesture guidance: Visual feedback and feedforward design options. In: Proceedings of AVI (2016)

16. Dim, N.K., Silpasuwanchai, C., Sarcar, S., Ren, X.: Designing mid-air tv gestures for blind people using user- and choice-based elicitation approaches. In: Proceedings of ACM DIS (2016)
17. Evangelou, G., Limerick, H., Moore, J.: I feel it in my fingers! sense of agency with mid-air haptics. In: Proceedings of IEEE WHC (2021)
18. Follmer, S., Leithinger, D., Olwal, A., Hogge, A., Ishii, H.: Inform: Dynamic physical affordances and constraints through shape and object actuation (2013)
19. Freeman, E., Brewster, S., Lantz, V.: Do that, there: An interaction technique for addressing in-air gesture systems. ACM (2016)
20. Freeman, E., Vo, D.B., Brewster, S.: Haptiglow: Helping users position their hands for better mid-air gestures and ultrasound haptic feedback. In: Proceedings of IEEE WHC (2019)
21. Freeman, E., Wilson, G.: Perception of ultrasound haptic focal point motion. In: Proceedings of ICMI. p. 697–701. ICMI '21 (2021)
22. Frier, W., Ablart, D., Chilles, J., Long, B., Giordano, M., Obrist, M., Subramanian, S.: Using spatiotemporal modulation to draw tactile patterns in mid-air. In: Proceedings of Eurohaptics. pp. 270–281. Springer (2018)
23. Frischmann, T.B., Mouloua, M., Procci, K.: 3-d gaming environment preferences: Inversion of the y-axis. *Ergonomics* **58**(11), 1792–1799 (2015)
24. Gavrilov, L.R.: The possibility of generating focal regions of complex configurations in application to the problems of stimulation of human receptor structures by focused ultrasound. *Acoustical Physics* **54**(2), 269–278 (2008)
25. Georgiou, O., Biscione, V., Harwood, A., Griffiths, D., Giordano, M., Long, B., Carter, T.: Haptic in-vehicle gesture controls. In: Proceedings of AutomotiveUI (2017)
26. Girdler, A., Georgiou, O.: Mid-air haptics in aviation—creating the sensation of touch where there is nothing but thin air. arXiv preprint arXiv:2001.01445 (2020)
27. Grace, K., Wasinger, R., Ackad, C., Collins, A., Dawson, O., Gluga, R., Kay, J., Tomitsch, M.: Conveying interactivity at an interactive public information display. In: Proceedings of PerDis (2013)
28. Grier, R.A.: How high is high? a meta-analysis of nasa-tlx global workload scores. In: Proceedings of HFES. SAGE Publications Sage CA: Los Angeles, CA (2015)
29. Groenewald, C., Anslow, C., Islam, J., Rooney, C., Passmore, P., Wong, W.: Understanding 3d mid-air hand gestures with interactive surfaces and displays: A systematic literature review. In: Proceedings of BCS. HCI '16 (2016)
30. Gupta, A., Pietrzak, T., Yau, C., Roussel, N., Balakrishnan, R.: Summon and select: Rapid interaction with interface controls in mid-air. In: Proceedings ACM ISS (2017)
31. Hajas, D., Ablart, D., Schneider, O., Obrist, M.: I can feel it moving: science communicators talking about the potential of mid-air haptics. *Frontiers in Computer Science* **2**, 534974 (2020)
32. Harrington, K., Large, D.R., Burnett, G., Georgiou, O.: Exploring the use of mid-air ultrasonic feedback to enhance automotive user interfaces. In: Proceedings of AutomotiveUI (2018)
33. Howard, T., Marchal, M., Lécuyer, A., Pacchierotti, C.: Pumah: Pan-tilt ultrasound mid-air haptics for larger interaction workspace in virtual reality. *IEEE Transactions on Haptics* (2020)
34. Huang, S., Ranganathan, S.P.B., Parsons, I.: To touch or not to touch? comparing touch, mid-air gesture, mid-air haptics for public display in post covid-19 society. In: SIGGRAPH Asia 2020 Posters (2020)

35. Hwang, I., Son, H., Kim, J.R.: Airpiano: Enhancing music playing experience in virtual reality with mid-air haptic feedback. In: Proceedings of IEEE WHC (2017)
36. Iwamoto, T., Tatezono, M., Shinoda, H.: Non-contact method for producing tactile sensation using airborne ultrasound. In: Proceedings of Eurohaptics (2008)
37. Kovács, P.T., Balogh, T., Nagy, Z., Barsi, A., Bordács, L., Balogh, G.: Tangible holographic 3d objects with virtual touch. In: Proceedings of ITS (2015)
38. Large, D.R., Harrington, K., Burnett, G., Georgiou, O.: Feel the noise: Mid-air ultrasound haptics as a novel human-vehicle interaction paradigm. *Applied ergonomics* **81**, 102909 (2019)
39. Lerusso, J., Rusanov, K., Perquin, M., Han, Y.j., Rossiter, H.E., Kolasinski, J.: Searching for the oblique effect in whole-hand tactile perception. *Cortex* (2019)
40. Limerick, H., Hayden, R., Beattie, D., Georgiou, O., Müller, J.: User engagement for mid-air haptic interactions with digital signage. *PerDis '19* (2019)
41. Long, B., Seah, S.A., Carter, T., Subramanian, S.: Rendering volumetric haptic shapes in mid-air using ultrasound. *ACM TOG* **33**(6) (2014)
42. Lopes, P., Jonell, P., Baudisch, P.: Affordance++: Allowing objects to communicate dynamic use. In: Proceedings of ACM CHI (2015)
43. Marchal, M., Gallagher, G., Lécuyer, A., Pacchierotti, C.: Can stiffness sensations be rendered in virtual reality using mid-air ultrasound haptic technologies? In: Proceedings of EuroHaptics. pp. 297–306. Springer (2020)
44. Marin, G., Dominio, F., Zanuttigh, P.: Hand gesture recognition with leap motion and kinect devices. In: IEEE ICIP. IEEE (2014)
45. Markussen, A., Jakobsen, M.R., Hornbæk, K.: Vulture: A mid-air word-gesture keyboard. In: Proceedings of ACM CHI (2014)
46. Marshall, J., Roper, A., Galliers, J., Wilson, S., Cocks, N., Muscroft, S., Pring, T.: Computer delivery of gesture therapy for people with severe aphasia. *Aphasiology* **27**(9), 1128–1146 (2013)
47. Martinez, J., Harwood, A., Limerick, H., Clark, R., Georgiou, O.: Mid-air haptic algorithms for rendering 3d shapes. In: IEEE HAVE (2019)
48. Marzo, A., Seah, S.A., Drinkwater, B.W., Sahoo, D.R., Long, B., Subramanian, S.: Holographic acoustic elements for manipulation of levitated objects. *Nature communications* **6**(1), 1–7 (2015)
49. Michelis, D., Müller, J.: The audience funnel: Observations of gesture based interaction with multiple large displays in a city center. *Intl. Journal of Human-Computer Interaction* **27**(6), 562–579 (2011)
50. Morrison, C., Smyth, N., Corish, R., O’Hara, K., Sellen, A.: Collaborating with computer vision systems: An exploration of audio feedback. In: Proceedings of DIS (2014)
51. Neate, T., Kladouchou, V., Wilson, S., Shams, S.: “just not together”: The experience of videoconferencing for people with aphasia during the covid-19 pandemic. In: Proceedings of ACM CHI (2022)
52. Perquin, M., Taylor, M., Lorusso, J., Kolasinski, J.: Directional biases in whole hand motion perception revealed by mid-air tactile stimulation. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior* **142**, 221–236 (2021)
53. Rakkolainen, I., Freeman, E., Sand, A., Raisamo, R., Brewster, S.: A survey of mid-air ultrasound haptics and its applications. *IEEE Transactions on Haptics* **14**(1), 2–19 (2020)
54. Romanus, T., Frish, S., Maksymenko, M., Frier, W., Corenthy, L., Georgiou, O.: Mid-air haptic bio-holograms in mixed reality. In: IEEE ISMAR-Adjunct (2019)
55. Roper, A., Marshall, J., Wilson, S.: Benefits and limitations of computer gesture therapy for the rehabilitation of severe aphasia. *Front. in Hum. Neuroscience* (2016)

56. Rovelo, G., Degraen, D., Vanacken, D., Luyten, K., Coninx, K.: Gestu-wan-an intelligible mid-air gesture guidance system for walk-up-and-use displays. In: IFIP Conference on Human-Computer Interaction. pp. 368–386. Springer (2015)
57. Rutten, I., Frier, W., Van den Bogaert, L., Geerts, D.: Invisible touch: How identifiable are mid-air haptic shapes? In: Proceedings of ACM CHI Adjunct (2019)
58. Sahoo, D.R., Nakamura, T., Marzo, A., Omirou, T., Asakawa, M., Subramanian, S.: Joled: A mid-air display based on electrostatic rotation of levitated janus objects. In: Proceedings of UIST (2016)
59. Sahoo, D.R., Neate, T., Tokuda, Y., Pearson, J., Robinson, S., Subramanian, S., Jones, M.: Tangible drops: a visio-tactile display using actuated liquid-metal droplets. In: Proceedings of ACM CHI (2018)
60. Samimi, N., von der Au, S., Weidner, F., Broll, W.: Ar in tv: Design and evaluation of mid-air gestures for moderators to control augmented reality applications in tv. In: Proceedings of MUM (2021)
61. Shakeri, G., Williamson, J.H., Brewster, S.: Novel multimodal feedback techniques for in-car mid-air gesture interaction. ACM (2017)
62. Shakeri, G., Williamson, J.H., Brewster, S.: May the force be with you: Ultrasound haptic feedback for mid-air gesture interaction in cars. In: Proceedings of AutomotiveUI (2018)
63. Spelmezan, D., Sahoo, D.R., Subramanian, S.: Sparkle: Hover feedback with touchable electric arcs. ACM (2017)
64. Tsalamlal, M.Y., Issartel, P., Ouarti, N., Ammi, M.: Hair: Haptic feedback with a mobile air jet. In: IEEE ICRA (2014)
65. Tsetserukou, D.: Haptihug: A novel haptic display for communication of hug over a distance. In: Proceedings of Eurohaptics. pp. 340–347. Springer (2010)
66. Villarreal-Narvaez, S., Vanderdonckt, J., Vatavu, R.D., Wobbrock, J.O.: A systematic review of gesture elicitation studies: What can we learn from 216 studies? In: Proceedings of DIS (2020)
67. Vogiatzidakis, P., Koutsabasis, P.: Gesture elicitation studies for mid-air interaction: A review. *Multimodal Technologies and Interaction* **2**(4), 65 (2018)
68. Vogiatzidakis, P., Koutsabasis, P.: Frame-based elicitation of mid-air gestures for a smart home device ecosystem. *Informatics* **6**(2) (2019)
69. Wang, C., Deng, H., Dong, Y., Zhang, X., Wang, D.H.: The capacity and cognitive processing of vibrotactile working memory for frequency. *Current Psychology* pp. 1–11 (2021)
70. Williams, A.S., Ortega, F.R.: Understanding gesture and speech multimodal interactions for manipulation tasks in augmented reality using unconstrained elicitation. *Proc. ACM Hum.-Comput. Interact.* (2020)
71. Wilson, G., Carter, T., Subramanian, S., Brewster, S.A.: Perception of ultrasonic haptic feedback on the hand: Localisation and apparent motion. In: Proceedings of ACM CHI (2014)
72. Wittorf, M.L., Jakobsen, M.R.: Eliciting mid-air gestures for wall-display interaction. In: Proceedings of NordiCHI (2016)
73. Wobbrock, J.O., Aung, H.H., Rothrock, B., Myers, B.A.: Maximizing the guessability of symbolic input. In: In Proceedings of ACM CHI Adjunct (2005)
74. Wobbrock, J.O., Wilson, A.D., Li, Y.: Gestures without libraries, toolkits or training: A \$1 recognizer for user interface prototypes. In: Proceedings of UIST (2007)