Mid-Air Haptic Rendering of 2D Geometric Shapes with a Dynamic Tactile Pointer

Daniel Hajas[1], Dario Pittera[1], Antony Nasce[2], Orestis Georgiou[2], Marianna Obrist[1]

Abstract—An important challenge that affects ultrasonic mid-air haptics, in contrast to physical touch, is that we lose certain exploratory procedures such as contour following. This makes the task of perceiving geometric properties and shape identification more difficult. Meanwhile, the growing interest in mid-air haptics and their application to various new areas requires an improved understanding of how we perceive specific haptic stimuli, such as icons and control dials in mid-air. We address this challenge by investigating static and dynamic methods of displaying 2D geometric shapes in mid-air. We display a circle, a square, and a triangle, in either a static or dynamic condition, using ultrasonic mid-air haptics. In the static condition, the shapes are presented as a full outline in mid-air, while in the dynamic condition, a tactile pointer is moved around the perimeter of the shapes. We measure participants' accuracy and confidence of identifying shapes in two controlled experiments ($n_1 = 34, n_2 = 25$). Results reveal that in the dynamic condition people recognise shapes significantly more accurately, and with higher confidence. We also find that representing polygons as a set of individually drawn haptic strokes, with a short pause at the corners, drastically enhances shape recognition accuracy. Our research supports the design of mid-air haptic user interfaces in application scenarios such as in-car interactions or assistive technology in education.

Index Terms—Mid-Air Haptics; Touch; Geometry; Shape Perception; Memory Chunking; Haptic Controls; In-Car Interaction; Assistive Technology

I. INTRODUCTION

Mid-air haptics describes the technological solution of generating tactile sensations on a user's skin, in mid-air, without any attachment on the user's body. One way to achieve this is through the application of focused ultrasound, as first described by Iwamoto et al. in 2008 [1], and commercialised by Ultraleap in 2013. A phased array of ultrasonic transducers is used to focus acoustic radiation pressure onto the user's palms and fingertips. Modulating the focus points, such that it matches the resonant frequency of the cutaneous mechanoreceptors found in humans ($\sim$5 Hz to 400 Hz) [2], causes a localised tactile sensation to be perceived by the user. With the use of multipoint and spatiotemporal modulation techniques, it is possible to create more advanced tactile sensations such as lines, circles, and even 3D geometric shapes [3], [4], [5], [6], [7].

As ultrasonic mid-air haptic technology is being explored in more and more application areas such as in art [8], multimedia [9], virtual reality [10], [11], and in-car user interfaces

1. Sussex Computer Human Interaction (SCHI) Lab, Creative Technology Research Group, School of Engineering and Informatics, University of Sussex, Brighton, United Kingdom
2. Ultraleap Limited, Bristol, United Kingdom
Corresponding author: Daniel Hajas (dh256@sussex.ac.uk)

Fig. 1. Experimental set-up. An ultrasonic array is positioned inside an acrylic box. On top of the box there is an opening that allows participants' hand, specifically the palm, to be stimulated with mid-air touch.

[12], [13], several challenges have emerged regarding tactile interaction in mid-air. One such challenge is shape identification. In contrast to physical touch, we cannot explore the interaction space and acquire tactile information with the same set of exploratory procedures as those discussed by Lederman and Klatzky [14]. For example, we cannot push or squeeze the surface of a tactile cube displayed in mid-air to determine its stiffness, lift it to judge its weight, nor follow along its contours with our fingers to determine whether it is a cube or not, in the same way we would do with a physical object. While progress in perceiving material properties in mid-air, such as texture is being made [15], mid-air haptic technology faces some important challenges when geometric properties of haptic sensations are to be displayed and explored through mid-air touch. Namely, if the geometry of the displayed items remains ambiguous, e.g., if a circle were to easily be confused with a square, mid-air haptic technologies would be unsuitable for a wide range of applications that require accurate and reliable shape identification.

To address this important challenge, we have experimentally investigated new and existing approaches to displaying 2D geometric shapes in mid-air. Specifically, we distinguish between two ways of rendering 2D tactile shapes either as *static* or *dynamic*. In the former case, the stationary outline of a shape (e.g., a circle, square, or triangle) is displayed in mid-air, while in the latter case, a slowly moving pressure point traces the outline of the shape. In the following, we will measure the performance of these two haptic rendering approaches: 1) stationary shapes, and 2) dynamic tactile points, with regards to their ability to accurately convey 2D geometric information

to the user.

To that end, and based on our own prior observations and experiences of people interacting with mid-air haptic technology, we have hypothesised that geometric shapes are recognised more accurately and more confidently when they are presented as dynamic stimuli. For instance, a circle is more likely to be recognised when a tactile pointer is moved around its circumference, than in its static counterpart [16]. In the context of physical touch, our distinction between static and dynamic stimuli is analogous to pressing a cookie cutter against the palm vs. drawing its shape on the palm with a finger or pointy object. Motivated by this analogy, we were specifically interested in studying mid-air touch to test our hypotheses derived from the primary research question: *How accurately and confidently can people identify 2D shapes in mid-air when displayed with a dynamic tactile pointer (DTP), instead of the outline of a stationary shape?*

Two experiments were conducted with $n_1 = 34$ and $n_2 = 25$ participants in which people were asked to identify the shapes they felt, and rate their confidence in their answer. A circle, square, and an equilateral triangle were displayed using the two rendering approaches (static and dynamic). Additionally, we explored both passive and active exploration, where participants were either prohibited or allowed to move their hand freely during the mid-air tactile interaction. Our analysis showed that participants were significantly more accurate and confident in identifying shapes, when presented with the dynamic modality. Furthermore, we also measured that a $300\,\text{ms}$ and $467\,\text{ms}$ pause of the DTP at the corners of the square and triangle respectively, improved people's ability to correctly recognise the displayed shapes by over 30%.

This paper contributes both novel scientific insights about the tactile perception of 2D shapes, and also provides design guidelines for improved mid-air haptic interfaces and haptic visualisations. Both of these contributions are discussed within the context of two application areas (automotive and education) from a haptics and HCI perspective. Specifically, we provide parameter recommendations for optimal shape recognition renderings that could be used for novel assistive technologies that enhance teaching of geometry and mathematics for visually impaired students, or for the rendering of haptic icons and controls in novel gesture controlled car user interfaces [12]. In both cases, a more accurate and confident identification of the communicated haptic shapes can significantly improve their effectiveness and thus improve adoption rates of mid-air haptic interfaces in the future.

II. RELATED WORK

We present a literature review on displaying haptic shapes, the implications of stationary shapes and dynamic tactile stimuli, as well as the role of active and passive touch in recognising geometric features.

A. Static and Dynamic Tactile Stimuli

In tactile graphics design, it is a frequent recommendation to use discontinuous tactile features, for example, to use open arrow heads instead of solid ones [17]. Such design guidelines support the notion that human tactile perception performs better at detecting a change in stimuli, rather than a continuous stimulus. This effect is researched through the comparison of oscillatory and static tactile stimulation. Oyarzabal et al. [18] has shown that indented geometric patterns are more likely to be correctly discriminated when a low frequency vibration is applied to tactile pixels on a tangible shape display. In contrast, Pietrzak et al. [19] studied participants' recognition performance of directional clues. They found that static patterns are better recognised than dynamic ones, when discriminating between eight tactile icons depicting various line gradients. This was associated with the fact that in the static icon condition, participants could explore the pattern in more detail, i.e. an advantage due to active exploration.

B. Active and Passive Touch

In 1962, Gibson not only defined active and passive touch [20], but also performed an experiment on rotating stimuli. Gibson considered passive touch, and asked participants to identify shapes when these were pressed against the hand statically, and when these were rotated. Results showed a 72% accuracy in the rotation condition, opposed to a 49% accuracy in the static condition. Further to the passive (rotation) and passive (static) stimuli, he also found active exploration of the shapes to be superior. He also reports on strategies named by subjects, such as counting corners or points when trying to identify geometric forms.

Schwartz et al. [21] replicated Gibson's experiment, and found controversial results. Active and passive touch recognition of shapes did not differ significantly; however, in the passive (static) condition, an accuracy of only 38.5% was obtained, which was significantly lower than the accuracy obtained in the passive (sequential) condition (92.5%). In Heller's work, the influence of exploration time was discussed in context of form recognition [22]. Heller's study showed that active exploration outperformed both the passive (static) and passive (sequential) stimuli, with 5 seconds of active exploration yielding a similar accuracy to $30\,\text{s}$ of passive touch.

According to Holmes et al. [23] kinaesthetic information plays a key role when we need to discriminate 2D shapes larger than the fingertip. Pasquero and Hayward [24] also remind us how a tactile display should allow freedom of active exploration. Such integration of cutaneous and kinaesthetic perception has been studied in context of mid-air haptics too. Inoue et al. [25] investigated Just-Noticeable-Difference (JND) values of position and angle perception, while allowing active, free-hand exploration for participants to inspect volumetric haptic objects in mid-air. HaptoMime [26], and HaptoClone [27] further discuss active exploration specific applications of volumetric mid-air haptic sensations.

C. Haptic Shape Recognition

Form perception has been studied through multiple tactile interfaces, and multiple body parts. Kaczmarek et al. [28] compared shape recognition via the fingertips on a 49 point electro-tactile array, with a raised dot pattern alternative. Participants discriminated four differently sized circles, squares

and equilateral triangles to an accuracy of 78.5% in the electro-tactile array condition, and 97.2% in the raised dot condition. Bach-Y-Rita et al. [29] replicated the study on the tongue, yielding similar results. Dynamic ways of rendering haptic shapes were also studied by Ion et al. [30]. Error rates of recognising 12 shapes was significantly lower using a skin drag interface than a vibro-tactile system. Participants also classified the stimuli created by the skin drag display, through the movement of a physical probe across the skin, as "clearer" and the vibrating stimuli as "blurry".

Theurel et al. [31] studied the role of visual experience on the prototype effect in the haptic modality of shape recognition. Comparing squares, rectangles, and triangles in their canonical and non-canonical representations, the study with congenitally blind and blindfolded sighted adolescents showed that visual exposure to prototypical representations of shapes, allowed blindfolded participants to achieve faster recognition time. Hence, the prototype effect is not intrinsic to the haptic modality, since the congenitally blind participants were significantly slower, even though they performed $\sim 20\%$ more accurately in recognising shapes. Since our study involved sighted participants and invisible stimuli, we decided to display shapes in their prototypical orientation, eliminating potential confounding variables.

Shape recognition was also studied in mid-air haptics. Korres & Eid [32] studied 2D patterns and measured identification accuracy to be 59.4% with mean recognition time being $13.9\,\mathrm{s}$. Rutten et al [16] tested 2D sensations, where, line based patterns were better recognised than circular ones. It was also noted that a dial like sensation was more accurately recognised than a static shape. Howard et al. [5], studied the ability of people to discriminate line orientation using mid-air haptics. 83% of participants did not express a preference of line orientation in their subjective reports, and this finding was reflected in the indifferent identification scores too. Replicating or contradicting these findings on perception of horizontal, vertical or diagonal lines might be valuable in design processes, such as a decision on using a square shape vs. a triangle. Long et al. [4] also showed that volumetric haptic shapes in mid-air can be perceived at 80% accuracy, but it did not evaluate users' performance on 2D geometry, a challenge that we address, and expand on in the present work.

III. Experimental design

To investigate the main research question on how accurately and confidently people can identify 2D shapes in mid-air, when rendered with DTP instead of a static outline, we defined the following two hypotheses:

H.1   Shapes will be correctly recognised on significantly more occasions when rendered as dynamic stimuli than as static stimuli.

H.2   Shapes will be correctly recognised with significantly more confidence when rendered as dynamic stimuli than as static stimuli.

Evaluating our hypotheses, we performed two controlled experiments and two pilot studies. Both experiment 1 and experiment 2 investigated the primary hypotheses (H.1 and H.2),

as described in sections IV and VI. However, in experiment 2, we modified the dynamic stimuli to also evaluate a new hypothesis (H.3, see section V) conceived after the analysis of experiment 1. Namely, in experiment 2, the dynamic stimuli were changed from a continuous loop to an interrupted loop, which means that the tactile pointer paused its movement for $300\,\mathrm{ms}$ and $467\,\mathrm{ms}$ at the corners of the square and triangle respectively. To find the optimal pause times in the movement of the tactile pointer for the different shapes, we ran two pilot studies, as described in section V. An overview of the experimental design is shown in Figure 2. Research ethics approval was obtained before recruiting participants.


Fig. 2. Summary of the two main experiments including two in-between pilot studies to determine the optimal parameters for experiment 2.

IV. Experiment 1: Single-stroke shapes

In experiment 1, we tested hypotheses H.1 and H.2. Importantly, the tactile pointer was moved around the displayed shape giving no emphasis to any corners, as if drawn using a single continuous (brush) stroke.

A. Method

1) Participants: Participants were selected from the general public and aged 18 to 50 years. We set an upper age limit to account for the potential decline of tactile acuity with age [16]. We recruited 34 participants (f=20, m=14), with a mean age of 27.21 ± 5.79 years. 30 participants were right handed, two left handed, and two reported not having a dominant hand. On a scale from 1 to 7, where 1 meant "no experience at all", and 7 meant "regular user for at least one year", participants' experience with the haptic interface was a mean of 2.00 ± 1.42. Participants declared on the consent form that they did not have any sensory impairment related to their sense of touch.

2) Materials:

a) Stimuli: Originally, we considered eight shapes to test our hypothesis on. These were a circle, square, right-angle triangle, plus-cross, ellipse, rectangle, equilateral triangle and x-cross (see Figure 3). However, for simplification, we decided to limit the study to only three shapes: a circle, square and an upright equilateral triangle, as often seen in literature (e.g. [31], [28]). Using only three prototypical geometric patterns [31], we wanted to eliminate any potential confounding variables due to similarities of shape geometry.

The method of rendering static and dynamic haptic shapes differ both perceptually and in the way that they are generated. The static stimuli employed spatio-temporal modulation (STM) [33], where a single focus of constant amplitude

Fig. 3. Overview on the original set of shapes considered in the study design phase. The final selection of three shapes used in our experiments are highlighted in green.

(intensity $= 1$) is rapidly moved round the shape perimeter. The rotation frequency causes the human skin to vibrate at the same frequency (and its harmonics [34]) along the entire path trajectory, resulting in the perception of a static tactile sensation, analogous to pressing a cookie cutter against the palm. The dynamic stimuli employed amplitude modulation (AM) [35], [3], where a single focus of oscillating amplitude intensity between 0 and 1, is slowly moved around the shape perimeter. The oscillating frequency causes the human skin to vibrate at the same frequency (and its harmonics [34]) but only at the focus, resulting in the perception of a dynamic tactile sensation, analogous to a pointy object or brush drawing shapes on the palm.

To study whether the method of rendering (static vs dynamic) had an effect on identification accuracy, we created a static and a dynamic version of the three chosen shapes, totalling six different stimuli. The parameters were kept constant across all six stimuli. We chose the size of the shapes ($6 \, \mathrm{cm}$ diameter/side length) to fit an average adult palm (anthropometric mean of palm length: $10.56 \, \mathrm{cm} \pm 0.46 \, \mathrm{cm}$) [36]. We chose $70 \, \mathrm{Hz}$ for the STM rotational frequency, as it is near the optimal $5 \, \mathrm{m \, s^{-1}}$ to $10 \, \mathrm{m \, s^{-1}}$ draw speed, for path lengths given by the static shape outlines [6]. For consistency, we chose $70 \, \mathrm{Hz}$ as the AM oscillation frequency, even though the optimal value for a point like stimulus is near $200 \, \mathrm{Hz}$. We used anti-clockwise pointer movements which is the default setting in the experimental device. The rate of drawing shapes using the dynamic stimulus type was chosen to be $0.5 \, \mathrm{Hz}$ ($2 \, \mathrm{s}$ per complete shape), such that the movement feels natural, i.e., as if a finger drew on the palm. The pointer had a diameter of $0.8 \, \mathrm{cm}$, corresponding to the wavelength of the ultrasonic carrier, and simulating the size of a fingertip. The centre of the shapes coincided with the origin of the haptic interface's coordinate system, but vertically translated by $15 \, \mathrm{cm}$ above the surface of the device (see Figure 1).

b) Device: We used a mid-air haptic device manufactured by Ultraleap Ltd, which generates the tactile sensation using 256 ultrasound transducers. In order to fix participants hand at the same height and area where the stimuli are displayed, we placed the device within a hand-support cavity. Participants were instructed to rest their hand on top of the support, over an $\sim 10 \times 10 \, \mathrm{cm}$ opening, as shown in Figure 1. To create the stimuli, we used the Ultrahaptics Sensation Core Library (SCL). The SCL includes a Python scripting interface, which allows developers to design sensations by constructing a graph of inter-connected operations, such as path geometry,

transforms, or animations. The sensations were prepared in advance, such that a Python script can call and display the stimuli on the haptic interface. The script was responsible for logging data, and randomising the order of stimuli.

c) Task: The experimental task was simple: "Tell the researcher the shape you felt, and how confident you are in your answer". We evaluated our hypotheses in two conditions: (1) passive, and (2) active touch as part of the same experiment. In the active condition, participants were allowed to move their hand to explore the stimuli. In passive touch, participants were instructed to keep their hand still. The dynamic and static stimuli were displayed in both active and passive conditions.

Prior to displaying the sequence of shapes, participants were given a chance to familiarise themselves with the experimental setup and the tactile sensations. A matrix of 3×3 focal points were projected on the palm sequentially, from top left to bottom right, with the central point coinciding with the centre of the shapes. Following this, we displayed the six stimuli for $6 \, \mathrm{s}$ respectively, but without disclosing the order of shapes. Although we did not set a maximum number of times the familiarisation could be repeated, none of the participants did the familiarisation session more than twice.

After the familiarisation stage, participants were shown the first stimulus for an indefinite duration and asked to announce what shape they felt. At the moment of announcement the stimulus was terminated. Participants were told that their options were limited to "circle", "square" or "triangle". In experiment 1, we also emphasised, that a "I don't know" response is also allowed. Before moving to the next stimulus, the confidence rating was asked and recorded. This task was repeated 24 times in a randomised order, with each of the three dynamic, and three static stimuli repeated four times, in both of the active and passive conditions. We measured two dependent variables: the *accuracy* of the named shape, and participants' *confidence* in the perceived shape. Accuracy (a dichotomous variable) simply indicated whether the shape was correctly perceived or not. The confidence rating was a self-report scale, from 1 to 7, where 1 meant "not sure at all" and 7 meant "most certain". We also recorded the time between the start and termination of stimuli; however, we did not intend to use this data to test our hypotheses in this study.

3) Procedure: Upon arrival to the experimental space, participants were introduced to the experimental procedure, and informed consents were obtained. We started collecting demographic data, then participants were instructed to place their right hand above the haptic interface. We carried out a within group experiment, where the active vs. passive conditions were counterbalanced and the stimuli were presented in a random order.

We strived to keep the experimental setup as controlled as possible by keeping the room temperature comfortably warm ($\sim 21°$), to prevent participants from having cold hands and reduced skin sensitivity. Ambient white noise was setup to prevent any audible clues from the haptic device. In the active touch condition, participants were asked to fix their sight on the wall in front of them to avoid speculative guesses of the felt shape, based on the visual inspection of their moving hand. Between the active and passive touch conditions, a $30 \, \mathrm{s}$ break

was allowed. Participants were given a sponge ball to fidget with, and refresh their hand muscles, skin and joints.

At the end of the experimentation, we asked participants two qualitative questions: (1) "*Q1: Which type of stimuli did you find easier to identify?*"; and (2) "*Q3: What strategies did you use, if any, to try to understand the shape?*". We kept written notes on the responses, but did not collect qualitative data systematically in experiment 1. The entire procedure took 30 minutes per participant, who received a £5 Amazon voucher for their time.

B. Results

For the analyses we use *R* (v3.5.2) statistical software. For ease of reading, we grouped the report according to passive and active touch conditions.

1) Passive touch – accuracy metrics: A McNemar's test showed a statistically significant difference ($p < 0.001$) in accuracy across the static and dynamic stimuli. We also analysed data with respect to individual classes (i.e. circle, triangle and square). Figure 4 shows the confusion matrices for both static and dynamic stimuli, but excluding the "I don't know" answers. The overall accuracy for static stimuli was 50.6% and for dynamic stimuli was 56.7%. This supports hypothesis H.1. In both conditions, the matrices show a high level of confusion in participants' answers. In particular, the circle and the square shapes are the most confused. For example, excluding "I don't know" answers, 38% answers of square when the stimulus was a circle, or 33% answers of circle when the stimulus was a square in the static stimulus type, with occasional mistakes in recognising the triangle. This is also supported by the subjective reports of users: *P9: "You could not feel whether it was supposed to be a circle or a square because the shape filled up all of the space, and because you couldn't feel the edges."*.


Fig. 4. Confusion matrix for the passive static (left) and passive dynamic (right) stimuli, expressed as percentage.

2) Passive touch – confidence levels: Figure 5 illustrates the box plot of confidence level for both static and dynamic stimuli. The sample deviates from a normal distribution as assessed by the Shapiro-Wilk's test ($p < 0.05$). Therefore, we ran a Wilcoxon signed-rank analysis to test differences between the confidence levels in static and dynamic stimuli. The test resulted statistically significant ($V = 4794, p < .001$). Participants are more confident in their choices when feeling shapes dynamically drawn (median = 5), than feeling static stimuli (median = 3). This supports hypothesis H.2. The

recorded time measurements were 10.2 ± 8.6 seconds for static stimuli, and 11.2 ± 8.3 seconds for dynamic stimuli.
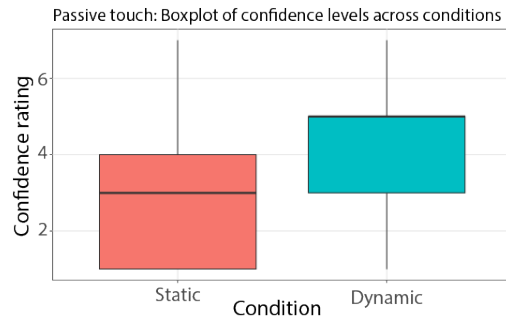

Fig. 5. Box plot of confidence levels across the passive static (red), and passive dynamic (green) stimuli.

3) Active touch – accuracy metrics: McNemar's test did not find significant differences between static and dynamic stimuli in the active condition ($p = 0.22$). This falsifies hypothesis H.1. We again analysed data with respect to individual shapes and created confusion matrices (see Figure 6). The overall accuracy for static stimuli was 57.3%, and for dynamic stimuli was 52.7%. Both types of stimuli brought participants to a high level of confusion in the active condition.


Fig. 6. Confusion matrix for the active static (left) and active dynamic (right) stimuli, expressed as percentage.

4) Active touch – confidence levels: From the box plot shown in Figure 7, it appears that reported confidence levels are higher for dynamic stimuli. This is confirmed by a Wilcoxon signed-rank analysis ($V = 10591, p < .001$). The median scores are 3 and 4, for static and dynamic stimuli respectively, supporting hypothesis H.2. The recorded time measurements were 15.4 ± 10.7 seconds for static stimuli, and 14.8 ± 11.3 seconds for dynamic stimuli.

5) Qualitative results: In the passive condition, every participant said that identifying shapes, as dynamic stimuli was easier. Some only expressed a milder difference: *P15: "It's easier because it feels clearer, whereas the 'cookie cutter' case is more blurry."*. Others expressed a stronger disliking of static stimuli: *P7: "Oh, not again the muddy."*, or *P33: "It's very difficult to grasp when it's a full blast. It just feels like air."*. Multiple participants described the static shapes as too "muddy", "blurry", or "fuzzy" to tell what shape it is. For dynamic stimuli, two different strategies were mentioned. One, focusing on curvature characteristics: *P27: "The circle felt like a smooth curve, whereas with triangle and square you*

Fig. 7. Box plot of confidence levels across the active static (red), and active dynamic (green) stimuli.

Fig. 8. An illustration of rendering squares with DTP, either as a single-stroke (SSDTP) stimulus or as a multi-stroke (MSDTP) stimulus.

could feel the corners.". Two, observing the dynamics of the moving point" *P26: "It slows down around the corners.".*

In the active condition, coherency of reports broke down and depended on the strategies people followed. Participants found dynamic stimuli easier, if they tracked the tactile pointer: *P32: "The moving point was even easier, as you could almost place your hand on it and follow".* However, the majority of people reported static stimuli to be slightly easier to recognise, if they adapted the strategy of tilting their hand, or focusing on points of stimulation on their palm.

C. Summary

Our results show that participants are significantly more accurate in recognising shapes, when these are displayed as dynamic stimuli (56.7%) vs. a static representation (50.6%), but only when their hand is fixed in space. Hence, for passive touch we can verify H.1, even though the effect size is small. For active touch, H.1 is false. Reported confidence levels are also significantly higher for dynamic stimuli, in both passive and active touch, making H.2 true for both conditions. The qualitative data revealed commonly used descriptors referring to the clarity of sensations, which we explore further in experiment 2. Although our time measurements are comparable to the mean recognition time (13.9 s) found by Korres and Eid [32], this finding is only indicative and not conclusive. We did not control how long participants were allowed to think before giving an answer. The high standard deviations also suggest that for some participants identification and announcement might not have happened simultaneously.

V. PILOT STUDIES: INCREASING RECOGNITION

The results of experiment 1, backed up with qualitative reports, suggested that participants could not discriminate well between shapes, even if these were dynamically rendered. In particular, people were repeatedly confusing circles and squares. In order to address this, we devised a second experiment that would test an additional hypothesis:

H.3   For dynamic stimuli, displaying shapes as a collection of discrete haptic strokes in form of an interrupted loop, instead of a continuous loop, will further improve the accuracy of shape recognition.

A. Parametrisation and chunking of haptic output

We motivated this hypothesis based on the literature discussing unistroke I/O and cognitive chunking. Considering visual chunking representations, such as a study performed by Zhang et al. [37], it is known that a single continuous line may form a chunk, which represents a straight line, a curve, or a circle. For polygons, it is expected that the number of edges, and vertices are perceived independently as single strokes, but grouped into the appropriate chunk. For example, a group of three strokes form a chunk representing a triangle. Chunking in HCI was discussed by Buxton [38] through multiple scenarios, in search for methods of accelerating the transition between novice and expert users of a computer interface. Buxton concludes that *"The key is gesture-based phrasing to chunk the dialogue into units meaningful to the application. – This desired one-to-one correspondence between concept and gesture leads towards interfaces which are more compatible with the user's model."* [38]. He suggests that this principle is desirable for any application, from terminal commands to input-output interfaces, hence it is worth investigating in cases of novel haptic output devices. Goldberg & Richardson [39] designed a unistroke alphabet to find equivalents of touch typing with the use of a stylus. Such touch input system enables the transition from novice to expert user by means of increased input speed, while also enables higher accuracy interpretation for the recognition system. Robust tools, such as the $1 Recognizer [40] enabled non-experts to incorporate gesture recognition in their UI. However, it also opened up new research topics, such as how gesture articulation speeds affected recognition accuracy. In other words, what parameters of the input contribute to successful recognition by the system. With the evolution of haptic output devices, researching unistroke related parameters, in context of human recognition abilities becomes an interesting research topic. For instance, Hoshi [41] used ultrasonic mid-air haptics to transmit gesture input into unistroke like haptic output, rendered on the palm. An accuracy of 44% recognition was demonstrated, but no rendering parameters were discussed or evaluated.

To test hypothesis H.3, we altered the dynamic stimuli to be composed of a collection of discrete haptic strokes. In experiment 2, the tactile pointer paused its movement when it reached a corner, while in experiment 1, the tactile pointer moved without interruption around the perimeter of the shapes (see Figure 8). Thus, we distinguished between two types of DTP rendering, the single-stroke (SSDTP) and the multi-stroke (MSDTP) mode. However, the duration of

interruption (referred to as "pause") remained a question. To determine the optimal duration of the pause, making the largest impact on recognition, we ran two pilot studies as described below. In the first pilot, we wanted to find out the answer to the question: "*Does recognition of the shape increase with the increase in duration of pauses at the corners?*". The second pilot was responsible for optimising the duration parameter, by determining the model for correlating duration and recognition, such as a linear or quadratic fitting model.

B. Pilot study 1

1) Method:

a) Participants: We recruited nine participants (f=4, m=5, mean age 29.6 ± 4.8 years). All of the qualifying criteria reported in experiment 1 were applicable in this pilot study.

b) Materials: Participants were given two tasks, in the same setup as experiment 1. In task 1, we displayed four repetitions of nine different versions of squares, drawn over $2\,s$, with increasingly long pauses of $0\,ms$ to $400\,ms$, in steps of $50\,ms$, at the corners. We asked participants to rate *"How much does the shape you felt resemble a square, on a scale from 1 (not at all) to 7 (very much)?"*. In task 2, the same task was completed for the triangle.

c) Procedure: The 36 stimuli were presented in a randomised order. Participants were told what the shape was on the display, and they were given standardised instructions of the task in print, since it was crucial they report how much the sensation resembles a shape, and not their ability to recognise it. We measured performance in only the passive touch condition. The pilot took 20 minutes, and a short break was allowed between the two tasks. Task 1 and task 2 were counterbalanced. No compensation was paid.


Fig. 9. Scatter plot of recognition: The mean scores of participants' rating (1-7) is plotted against the nine pause durations tested (ms) for the triangle (left) and square (right) in pilot study 1. A best fit curve is shown in blue.

2) Results: Figure 9 plots the mean scores of participants' ratings of recognition for the different pause durations at the corners of the triangle (left) and square (right). The graphs show that increasing the pause increases participants' perception of feeling a well defined shape. We ran Wilcoxon tests to investigate differences across the various durations. From these analyses, we isolated three groups: 1) [0, 50, 100] ms; 2) [150, 200] ms; 3) [250, 300, 350, 400] ms, for both shapes. Although, the difference between instances of each group were not statistically significant ($p > 0.05$), the scores for the three groups are statistically significantly different.

The results confirm that there is a direct relation between the time spent at the corners, as a kind of emphasis, and the participants' perceived sensation of a shape. However, from the graphs in Figure 9, it is not clear if the trend would descend for longer pauses or continue increasing in a linear fashion. For a clearer representation of the best-fit-curve's trend, we omitted error bars on the scatter plots and zoomed in on the area of interest. To investigate the trend, we ran pilot study 2.

C. Pilot study 2

1) Methods:

a) Participants: The pool of participants was identical to the group of participants taking part in the first pilot study.

b) Materials: We reduced the variation of stimuli by decreasing the tested conditions of the pause duration. However, we increased the repetitions from four to ten, to obtain a cleaner dataset. In task 1, we chose to test values of 0, 150, 300, and $500\,ms$ for squares. Another factor we accounted for in pilot study 2, is the difference between the draw speed of sides in triangles and squares. Since the overall rate of drawing and duration of pauses at corners were identical for both shapes, the speed at which sides are drawn differed. However, since pilot study 1 showed that there are intervals of pause durations at corners, at which no significant differences are observed, we chose to keep the draw speed of sides constant by varying the pause duration. Based on this speed, and the overall rate, we computed the equivalent duration of pauses in the triangle to be 167, 317, 467, and $667\,ms$ respectively. For completeness, we also added the $0\,ms$ baseline condition.

c) Procedure: The procedure was identical to that used in pilot study 1, except the number of trials. Task 1 involved 10 repetitions of four variations on the square, and task 2 involved 10 repetitions of five variations on the triangle.


Fig. 10. Scatter plot of recognition: The mean scores of participants' rating (1-7) is plotted against the five/four pause durations tested (ms) for the triangle (left) and square (right) respectively, in pilot study 2. A best fit curve is shown in blue.

2) Results: For the triangle, we see from Figure 10 that the best fit curve follows a quadratic trend, although it is less sharp than in the case of the square. The central values of $467\,ms$ and $300\,ms$ for the triangle and square respectively were statistically different ($p < 0.05$) from other values tested using Wilcoxon tests. We see that a too long a pause may decrease performance. In case of the square, participants may benefit from feeling the edges. A square rendered in $2\,s$, with a $500\,ms$ pause at the corners, means that there is no time left to render edges. The tactile pointer is repositioned discontinuously from corner to corner.

D. Summary

Two pilot studies were conducted to investigate the effect of pauses at corners on shape recognition. The pauses interrupted the movement of the tactile pointer, rendering a haptic shape. It was shown that different pause durations can have a noticeable impact on recognition, and that the optimal pause durations differ from shape to shape. Although the results we obtained were indicative of the most appropriate duration to use, it was not conclusive whether participants were going to be able to discriminate the shapes, once the stimuli were mixed, as in experiment 1. This was the objective of experiment 2.

VI. EXPERIMENT 2 – MULTI-STROKE SHAPES

This experiment studied all three hypotheses H.1, H.2 and H.3. We measured participants' accuracy and confidence in mid-air haptic shape recognition, for static and dynamic stimuli in passive and active conditions. Importantly, we used the modified dynamic stimuli, where the tactile pointer took short pauses at the corners of the displayed shape, as if drawn using multiple (brush) strokes.

A. Method

1) Participants: We recruited 25 participants (f=14, m=11), with a mean age of 30.24 ± 7.80 years. 22 participants were right handed and 3 were left handed. Their experience with the haptic interface, on a scale from 1 to 7, was 2.08 ± 1.20. No one declared a disorder compromising their tactile acuity. Participants of the pilot studies were excluded from taking part in this experiment.

2) Materials: The stimuli used in the static condition were identical to those used in experiment 1. In the dynamic method of rendering, we exchanged the single-stroke stimuli with multi-stroke sensations. Based on the results of the two pilot studies, we chose $300 \, \text{ms}$ and $467 \, \text{ms}$ long pauses at the corners of the squares and triangles respectively. We expected that this method would help in distinguishing between circles and squares displayed as dynamic stimuli.

3) Procedure: The task and procedure for experiment 2 followed the same protocol as in experiment 1, except in two aspects. First, we did not allow for an "I don't know" answer when identifying the presented shape. We chose to make this change to feed the confusion matrix with more relevant data. The minimum confidence score accounted for the "I don't know" option. Secondly, we wanted to perform a more thorough qualitative analysis, hence, we audio recorded the final five minute interviews, and included a third question, asking participants *"Q2: Using 2-3 adjectives, how would you describe the clarity, or sharpness of the shapes you felt in each of the conditions?"*.

B. Results

1) Passive touch – accuracy metrics: Confusion matrices for the two types of stimuli are shown in Figure 11. The overall accuracy for static stimuli was 51.7%, and for dynamic stimuli was 83.0%. This is a statistically different result (McNemar's test, $p < 0.001$) and a significant improvement compared to

the results in experiment 1, supporting hypothesis H.1. Values for the dynamic stimuli highlight how the shapes are better perceived with the introduction of multi-stroke shapes. Only 14% answers of square were given, where the shape was a circle; and only 9% answers of circle were given, where the shape was a square.


Fig. 11. Confusion matrix for the passive static (left) and passive dynamic (right) stimuli, expressed as percentage.

2) Passive touch – confidence levels: A Wilcoxon signed-rank analysis confirmed a significant difference ($V = 912, p < .001$) between confidence levels in the two stimulus types. Once again, participants were more confident in dynamic stimuli (median = 5), than in static stimuli (median = 3), as shown on the box plot in Figure 12. This supports hypothesis H.2. The recorded time measurements were 7.8 ± 5.6 seconds for static stimuli, and 7.8 ± 5.3 seconds for dynamic stimuli.

Fig. 12. Box plot of confidence levels across the passive static (red), and passive dynamic (green) stimuli, in experiment 2.

3) Active touch – accuracy metrics: Figure 13 shows the confusion matrices for the active condition. The overall accuracy for static stimuli was 57.3%, and for dynamic stimuli was 84.7%. This is a statistically significant difference (McNemar's test, $p < 0.001$) and makes hypothesis H.1 true.

4) Active touch – confidence levels: The reported confidence levels are again higher for dynamic stimuli (Wilcoxon signed-rank test: $V = 2574, p < 0.001$). The median score for the confidence level rating is 4 for static stimuli and 5 for the dynamic types (see Figure 14). This supports hypothesis H.2. The recorded time measurements were 9.3 ± 5.7 seconds for static stimuli, and 8.4 ± 5.5 seconds for dynamic stimuli.

5) Qualitative results: Our aim was to quantify observations on participants' comments from experiment 1, and systematically collect linguistic descriptors of the two types of stimuli. To do this, we transcribed all five minute interviews conducted at the end of the experiment. Relevant snippets

Fig. 13. Confusion matrix for the active static (left) and active dynamic (right) stimuli, expressed as percentage.


Fig. 14. Box plot of confidence levels across the active static (red), and active dynamic (green) stimuli, in experiment 2.

TABLE I
DESCRIPTORS OF PERCEIVED QUALITY OF SENSATIONS, QUALITY OF SHAPES, AND ABILITY TO RECOGNISE SHAPES.

Valance	**Positive**	**Negative**
Theme	**Perceived quality of sensation**	
Total count (static)	4	13
Total count (dynamic)	12	3
Frequent descriptors (static)	–	blow, wall of air (5) block (3)
Frequent descriptors (dynamic)	pencil/finger tip (3), smooth (1)	–
Theme	**Perceived quality of shapes**	
Total count (static)	2	32
Total count (dynamic)	28	4
Frequent descriptors (static)	–	fuzzy (7), blurry (3) unclear (5)
Frequent descriptors (dynamic)	clear (8), sharp (5), higher definition (4)	–
Theme	**Perceived ability to recognise shapes**	
Total count (static)	3	20
Total count (dynamic)	17	1
Frequent descriptors (static)	–	hard (10), indistinguishable (5)
Frequent descriptors (dynamic)	easy (10), makes mental image (3)	–

of the transcripts were extracted, and grouped into three categories, coded as: (Q1) Preference, (Q2) Descriptor, and (Q3) Strategy. After the coding process, we further abstracted information relevant to the respective category.

In Q1, we looked for how many people found either of the stimulus types easier based on their subjective reports, and how varied the spectrum of expressed difficulty is (from a little easier to a lot easier). We found that 22 of 25 participants reported that the dynamic condition was "easier". 3 participants said it depended on whether they explored actively or not. In the active touch they felt the static shapes were easier to recognise, though they still preferred the dynamic display mode when their hand was fixed. We also identified 11 positive, and 5 negative signifiers. Positive signifiers included adjectives, such as "definitely" (7 instances), or "much" (2 instances): *P9: "The moving one was definitely a lot easier.".* On the other hand, negative signifiers, such as "I think" (4 instances) or "perhaps" (1 instance) indicated a weaker preference: *P2: "I think the moving one was perhaps better.".*

In Q2, we abstracted a list of 28 adjectives, descriptive phrases associated with the individual conditions. We counted the frequency of these descriptors, and coded them according to three themes. The themes were divided into positive and negative attributes. For the most frequent adjectives and their occurrences in each theme, see Table I.

In Q3, we abstracted two key strategies. First, people who counted corners or edges in the passive (dynamic) condition, and people who moved their hand with the moving tactile pointer, in the active (dynamic) condition. In the former case, people reported that counting helped them create a mental picture of the shape: *P19: "I could see this almost like tracing something on my skin so I could kind of mentally construct*

the shape". In the latter case, participants relied on whether the movement of tactile stimulus on their hand, matched the self-initiated, kinaesthetic movement.

C. Summary

Comparing the accuracy results obtained for dynamic stimuli in experiment 1 and experiment 2, using a χ^2 test of homogeneity, we see a statistically significant difference in both the passive ($\chi^2 = 87.23, df = 1, p < 0.001$) and active conditions ($\chi^2 = 61.23, df = 1, p < 0.001$). Thus, we can claim H.3 to be true, since the results of experiment 2 show that displaying shapes as a collection of multiple strokes rather than a single stroke, can significantly improve accuracy of shape recognition. In particular, the overall accuracy in the passive touch for dynamic stimuli increased from 56.7% to 83.0%; while the accuracy also increased in the active touch, dynamic stimuli, from 52.7% to 84.7%. These results confirm hypothesis H.1. We see that for the dynamic stimuli in both passive and active touch, the median value of confidence is 5, which is significantly different from that for static stimuli, thus supporting H.2. The qualitative analysis also shows that people find static shapes more blurry or fuzzy, compared to dynamically drawn shapes, which have been named as clear, or having a higher definition. The answers given by participants to the interview questions show that recognising shapes presented as dynamic stimuli is easy, while it is hard for static stimuli.

VII. DISCUSSION

Our study reports on how accurately and confidently people can identify 2D shapes using mid-air haptic stimulation. Here, we discuss how our work contributes to haptics and HCI

research. We also outline possible application scenarios that can benefit from our findings.

A. Mid-Air Haptic Shape Recognition

We learnt three key lessons. First, in experiment 1 we showed that people can recognise more accurately and confidently the tested shapes, when these were rendered with DTP, instead of a stationary outline. Remarkably, while passive touch dynamic stimuli performed 6.1% better on accuracy than static shapes, in active exploration the dynamic stimuli performed 4.6% less accurately. Although the results in active touch are not statistically different, this is in line with previous studies [20]. It is likely that a shape presented as a full outline is better understood while explored actively, than when passively felt. This is apparent from comparing the accuracy results of static stimuli in the passive (50.6%) and the active (57.3%) conditions. In contrast, if both the tactile pointer and the participant's hand is moving, this may conflict the creation of accurate mental representations.

Secondly, experiment 2 showed that breaking down a shape into individual chunks (i.e. using multiple strokes) can increase the accuracy of shape recognition by $\sim 30\%$. Feeling a continuous loop led to higher levels of association with a circle, and feeling well distinguished corners, enabled participants to make a clear link with either triangle, or square: *P18: "Counting the corners, and if I didn't feel a corner and I felt a constant movement, then I thought it was a circle."*.

Thirdly, we obtained comparable results to those cited in the literature. Gibson found a 72% accuracy of shape recognition, in a passive (rotation) touch condition. This is similar to our results of 83.0% accuracy of dynamic stimuli in the passive condition. He also reported participants' recognition strategy to be "counting corners and points" [20], which we also found. Ion et al. [30] also found vibro-tactile interfaces to perform $\sim 20\%$ less accurately on a shape recognition task, compared to a skin drag display. This is in line with the $\sim 30\%$ difference between accuracy of identifying dynamic and static shapes in experiment 2. The qualitative reports of Ion et al. [30] "clearer" skin drag stimulus vs. "blurry" vibro-tactile stimulus are also matching our qualitative findings.

In addition, the two pilot studies provided the optimal pause duration parameters for the specific size and draw speed of the tested shapes. These were experimentally deduced, however we believe that this parameter can be defined precisely for a general geometry, as a function of other parameters, such as perimeter, number of sides, or rate of drawing. Reports of participants also clearly support the numerical findings: *P9: "Having definitive pauses at the vertices, meant that I could definitely feel four points. That must mean it's a square. I can definitely feel three points. That must mean it's a triangle. That helped immensely."*. Although we obtained an optimal pause duration for shape identification, it did not consider any use case restrictions. In some control interfaces, such as automotive, time is of the essence and therefore a trade-off may exist between accuracy and sensation duration. Optimising shape recognition time remains an open question.


Fig. 15. (left) A closeup photo of a finger, drawing a triangle into a palm; (middle) A person drawing a triangle on a tablet computer; and (right) A mid-air haptic kit stimulating a hand, in the pattern of a triangle.

B. Application Opportunities: Two User Scenarios

1) Scenario 1: Haptic Controls in Automotive Systems: Imagine a driver wishes to turn the volume of the radio down, and increase the temperature in the car. It is an important interaction design task of in-car interaction to provide interfaces that do not require the driver to take their eyes off the road [12], [42]. One possibility is to use gesture control interfaces with integrated haptic feedback. Given that people can easily distinguish between simple shapes, such as a circle and triangle, it becomes possible to design a gesture control interface with added haptic feedback. Placing the hand in an interaction space, a haptic icon appears. If it is a circle, a rotating movement in either direction could adjust the radio volume. Swiping movement brings up a new icon, for instance a triangle. Here, rotating movement of the hand in either direction results in changing the temperature. To evaluate the effectiveness and safety of such a system, we foresee an experiment replicating our findings in a car simulator, especially focusing on circumstances where users are subject to high cognitive demand, or potential risk.

2) Scenario 2: Geometry Instruction for Visually Impaired Students: Imagine a visually impaired student learning elementary geometry. Traditionally tactile graphics is embossed on paper, to aid the instruction. In certain scenarios, such as in secluded areas, the student requires remote help revising the concepts. In this case, through a voice call and the haptic interface, the tutor is able to assist, as illustrated in Figure 15. If the tactile paper is acoustically transparent, the mid-air haptics can be used as an auxiliary tool, highlighting areas on the paper. The regions of interest are discussed through guided exploration using the tactile pointer. Providing appropriate input devices for content creation, the immediate tactile feedback is also possible, which is a critical requirement [43]. To evaluate the merit of such a system, we foresee an experiment, which studies tactile shape perception in mid-air vs. tactile graphics in novice users.

C. Limitations and Future Work

One of the drawbacks of our method is the arbitrary choice of shape size. Recent work by Frier et al [44] suggests that the size of stimulus is affecting the perceived intensity of ultrasonic mid-air haptics. A potential solution is to personalise the size of the stimulus. Similarly, the arbitrary choice of rate at which the DTP completed a loop needs to be tested to

identify the optimal parameters. In physical touch it was shown that slower movement creates a sensation of curvature, while faster rates are perceived straighter [45]. This could contribute to confusions between a square and a circle when described with a continuously moving pointer. Further limitation of our study is the number of shapes tested. We have shown that displaying dynamic shapes is better recognised if it is either a circle, square or equilateral triangle; however, we know little about how well people could distinguish between shapes, such as a circle and an oval, or a triangle in different orientations. In future work, we wish to optimise parameters, such as rate, orientation, size, or type of stimulus used as a tactile pointer.

VIII. Conclusion

It is recommended that mid-air haptic devices render two-dimensional geometric shapes through the use of a dynamic tactile pointer, instead of displaying the full outline of the shape. It is also recommended to break down polygons into discrete sides, by interrupting the movement of the pointer at the vertices. The optimal pause duration for a 6 cm square, and equilateral triangle is 300 ms, and 467 ms respectively, when displayed at a rate of 2 s. According to these specifications, the accuracy of passive touch shape recognition is 83.0%, with active touch at 84.7%. These results are comparable to accuracies measured for mid-air haptics displaying 3D shapes, as well as studies using raised pin arrays and vibro-tactile displays. These insights may play a crucial role in a plethora of application areas, such as mid-air haptics control design, in an automotive context, or as assistive technologies for visually impaired distance learning children.

References

[1] T. Iwamoto, M. Tatezono, and H. Shinoda, "Non-contact method for producing tactile sensation using airborne ultrasound," 06 2008, pp. 504–513.

[2] D. A. Mahns, N. M. Perkins, V. Sahai, L. Robinson, and M. J. Rowe, "Vibrotactile Frequency Discrimination in Human Hairy Skin," *Journal of neurophysiology*, vol. 95, no. 3, pp. 1442–1450, 2006.

[3] T. Carter, S. A. Seah, B. Long, B. Drinkwater, and S. Subramanian, "Ultrahaptics: Multi-point mid-air haptic feedback for touch surfaces," in *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '13. ACM, 2013, pp. 505–514.

[4] B. Long, S. A. Seah, T. Carter, and S. Subramanian, "Rendering Volumetric Haptic Shapes in Mid-air Using Ultrasound," *ACM Trans. Graph.*, vol. 33, no. 6, pp. 181:1–181:10, nov 2014.

[5] T. Howard, G. Gallagher, A. Lécuyer, C. Pacchierotti, and M. Marchal, "Investigating the recognition of local shapes using mid-air ultrasound haptics," *IEEE World Haptics Conference*, July 2019.

[6] W. Frier, D. Ablart, J. Chilles, B. Long, M. Giordano, M. Obrist, and S. Subramanian, "Using spatiotemporal modulation to draw tactile patterns in mid-air," in *Haptics: Science, Technology, and Applications*. Springer International Publishing, 2018, pp. 270–281.

[7] A. Matsubayashi, Y. Makino, and H. Shinoda, "Direct finger manipulation of 3d object image with ultrasound haptic feedback," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 2019, p. 87.

[8] C. Vi, D. Ablart, E. Gatti, C. Velasco, and M. Obrist, "Not just seeing, but also feeling art: Mid-air haptic experiences integrated in a multisensory art exhibition," *International Journal of Human Computer Studies*, vol. 108, 2017.

[9] D. Ablart, C. Velasco, and M. Obrist, "Integrating mid-air haptics into movie experiences," in *Proceedings of the 2017 ACM International Conference on Interactive Experiences for TV and Online Video*, ser. TVX '17. ACM, 2017, pp. 77–84.

[10] D. Pittera, E. Gatti, and M. Obrist, "I'm sensing in the rain: Spatial incongruity in visual-tactile mid-air stimulation can elicit ownership in vr users," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. ACM, 2019, pp. 132:1–132:15.

[11] O. Georgiou, C. Jeffrey, Z. Chen, B. X. Tong, S. H. Chan, B. Yang, A. Harwood, and T. Carter, "Touchless haptic feedback for VR rhythm games," in *2018 IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2018, Tuebingen/Reutlingen, Germany, 18-22 March 2018*, 2018, pp. 553–554.

[12] K. Harrington, D. R. Large, G. E. Burnett, and O. Georgiou, "Exploring the use of mid-air ultrasonic feedback to enhance automotive user interfaces," in *Proceedings of the 10th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI 2018, Toronto, ON, Canada, September 23-25, 2018*, 2018, pp. 11–20.

[13] G. Shakeri, J. H. Williamson, and S. Brewster, "May the force be with you: Ultrasound haptic feedback for mid-air gesture interaction in cars," in *Proceedings of the 10th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, 2018, pp. 1–10.

[14] S. J. Lederman and R. L. Klatzky, "Haptic perception: A tutorial," *Attention, Perception, & Psychophysics*, vol. 71, no. 7, pp. 1439–1459, Oct 2009.

[15] E. Freeman, R. Anderson, J. Williamson, G. Wilson, and S. A. Brewster, "Textured surfaces for ultrasound haptic displays," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, ser. ICMI '17. ACM, 2017, pp. 491–492.

[16] I. Rutten, W. Frier, L. Van den Bogaert, and D. Geerts, "Invisible touch: How identifiable are mid-air haptic shapes?" in *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '19. ACM, 2019, pp. LBW0283:1–LBW0283:6.

[17] B. A. of North America and C. B. Authority, *Guidelines and Standards for Tactile Graphics*, 2010.

[18] M. Oyarzabal, M. Nakatani, and R. D. Howe, "Vibration enhances geometry perception with tactile shape displays," in *Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, ser. WHC '07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 44–49.

[19] T. Pietrzak, I. Pecci, and B. Martin, "Static and dynamic tactile directional cues experiments with VTPlayer mouse," in *Proceedings of the Eurohaptics conference (Eurohaptics 2006)*, Paris, France, 2006, pp. 63–68.

[20] J. J. Gibson, "Observations on active touch," *Psychological review*, vol. 69, pp. 477–91, 12 1962.

[21] A. S. Schwartz, A. J. Perey, and A. Azulay, "Further analysis of active and passive touch in pattern discrimination," *Bulletin of the Psychonomic Society*, vol. 6, pp. 7–9, 07 1975.

[22] M. Heller, "Active and passive touch: The influence of exploration time on form recognition," *The Journal of general psychology*, vol. 110, pp. 243–9, 05 1984.

[23] E. Holmes, B. Hughes, and G. Jansson, "Haptic perception of texture gradients," *Perception*, vol. 27, no. 8, pp. 993–1008, 1998.

[24] J. Pasquero and V. Hayward, "Stress: A practical tactile display system with one millimeter spatial resolution and 700 hz refresh rate," 01 2003, pp. 94–110.

[25] S. Inoue, Y. Makino, and H. Shinoda, "Active touch perception produced by airborne ultrasonic haptic hologram," pp. 362–367, 08 2015.

[26] Y. Monnai, K. Hasegawa, M. Fujiwara, K. Yoshino, S. Inoue, and H. Shinoda, "Haptomime: Mid-air haptic interaction with a floating virtual screen," in *Proceedings of the 27th Annual ACM Symposium on*

User Interface Software and Technology, ser. UIST '14. ACM, 2014, pp. 663–667.

[27] Y. Makino, Y. Furuyama, S. Inoue, and H. Shinoda, "Haptoclone (haptic-optical clone) for mutual tele-environment by real-time 3d image transfer with midair force feedback," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ser. CHI '16. ACM, 2016, pp. 1980–1990.

[28] K. Kaczmarek, M. Tyler, and P. Bach-y Rita, "Pattern identification on a fingertip-scanned electrotactile display," vol. 4, 01 1997, pp. 1694 – 1696 vol.4.

[29] P. Bach-y Rita, K. Kaczmarek, M. Tyler, and J. Garcia-Lara, "Form perception with a 49-point electrotactile stimulus array on the tongue: A technical note," *Journal of rehabilitation research and development*, vol. 35, pp. 427–30, 11 1998.

[30] A. Ion, E. J. Wang, and P. Baudisch, "Skin drag displays: Dragging a physical tactor across the user's skin produces a stronger tactile stimulus than vibrotactile," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '15. ACM, 2015, pp. 2501–2504.

[31] A. Theurel, S. Frileux, Y. Hatwell, and E. Gentaz, "The haptic recognition of geometrical shapes in congenitally blind and blindfolded adolescents: Is there a haptic prototype effect?" in *PloS one*, 2012.

[32] G. Korres and M. Eid, "Haptogram: Ultrasonic point-cloud tactile stimulation," *IEEE Access*, vol. PP, pp. 1–1, 09 2016.

[33] B. Kappus and B. Long, "Spatiotemporal modulation for mid-air haptic feedback from an ultrasonic phased array," *The Journal of the Acoustical Society of America*, vol. 143, no. 3, pp. 1836–1836, 2018.

[34] J. Chilles, W. Frier, A. Abdouni, M. Giordano, and O. Georgiou, "Laser doppler vibrometry and fem simulations of ultrasonic mid-air haptics," in *2019 IEEE World Haptics Conference (WHC)*. IEEE, 2019, pp. 259–264.

[35] T. Hoshi, M. Takahashi, T. Iwamoto, and H. Shinoda, "Noncontact tactile display based on radiation pressure of airborne ultrasound," *IEEE Transactions on Haptics*, 2010.

[36] A. Chandra, "Analysis of hand anthropometric dimensions of male industrial workers of haryana state," *International Journal of Engineering*, vol. 5, 08 2011.

[37] D. Zhang, Y. Ding, J. Stegall, and L. Mo, "The effect of visual-chunking-representation accommodation on geometry testing for students with math disabilities," *Learning Disabilities Research & Practice*, vol. 27, 11 2012.

[38] W. A. S. Buxton, "Human-computer interaction," R. M. Baecker, J. Grudin, W. A. S. Buxton, and S. Greenberg, Eds. Morgan Kaufmann Publishers Inc., 1995, ch. Chunking and Phrasing and the Design of Human-computer Dialogues, pp. 494–499.

[39] D. Goldberg and C. Richardson, "Touch-typing with a stylus," in *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, ser. CHI '93. ACM, 1993, pp. 80–87.

[40] J. O. Wobbrock, A. D. Wilson, and Y. Li, "Gestures without libraries, toolkits or training: A $1 recognizer for user interface prototypes," in *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '07. ACM, 2007, pp. 159–168.

[41] T. Hoshi, "Handwriting transmission system using noncontact tactile display," *Haptics Symposium 2012, HAPTICS 2012 - Proceedings*, 03 2012.

[42] G. Shakeri, J. H. Williamson, and S. A. Brewster, "Bimodal feedback for in-car mid-air gesture interaction," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, ser. ICMI 2017. ACM, 2017, pp. 518–519.

[43] J. Bornschein, D. Bornschein, and G. Weber, "Comparing computer-based drawing methods for blind people with real-time tactile feedback," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI '18. ACM, 2018, pp. 115:1–115:13.

[44] W. Frier, D. Pittera, D. Ablart, M. Obrist, and S. Subramanian, "Sampling strategy for ultrasonic mid-air haptics," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. ACM, 2019, pp. 121:1–121:11.

[45] N. Langford and R. J. Hall, "Cutaneous perception of a track produced by a moving point across the skin," *Journal of experimental psychology*, vol. 97, pp. 59–63, 02 1973.

Daniel Hajas received his M.Phys in Theoretical Physics from the University of Sussex (2017). He is a PhD student in the SCHI Lab at the University of Sussex, working on the intersection of mid-air haptics and science communication. Daniel's research is targeting the use of tactile experiences for purposes of provoking personal responses, which are known to be relevant in science communication, such as interest or enjoyment. Daniel was a Junior Research Associate in 2015 at the University of Sussex, researching tangible user interfaces and actuators.

Dario Pittera received his B.S. in Communication and Psychology (2011) and his M.S. in Clinical psychology, developmental psychology and neuropsychology (2014), both from the Universitá di Milano-Bicocca, Italy. He was a visiting researcher at the University of Birmingham 2015-16, where he was interested in psychophysics and haptic feedback technology. After that he started his PhD in the SCHI Lab at the University of Sussex. His research is focused on the psychophysical exploration of mid-air haptics, tactile illusions, and embodiment in VR using mid-air haptic technology. 2016-2017, Dario completed a four-months research internship at Disney Research, PA, USA.

Antony Nasce has a PhD in non-contact audio signal reproduction and an acoustical engineering degree from the University of Southampton. He is currently the Technology Evangelist for software tools at Ultrahaptics, and actively researches mid-air haptic interaction design. Before that, Antony worked at a leading art and image processing software company, The Foundry, where he helped test and develop a number of creative software tools for the visual effects industry.

Orestis Georgiou has a PhD in Applied Mathematics from the University of Bristol and has published over 65 articles in leading journals and conferences of Mathematics, Physics, Computer Science, Engineering and Medicine and holds 4 US patents. He is currently the Director of Research at Ultrahaptics and co-PI of the FET Open projects Levitate and H-Reality. Before that, Orestis was a Senior Researcher at Toshiba TRL, and a Postdoc at the Max Planck Institute PKS in Dresden. He got two Best Paper Awards and the 2019 IEEE Heinrich Hertz Award.

Marianna Obrist is Professor of Multisensory Experiences and head of the Sussex Computer Human Interaction (SCHI) Lab at the School of Engineering and Informatics, University of Sussex. Her research ambition is to establish touch, taste, and smell as interaction modalities in HCI. Before joining Sussex, Marianna was a Marie Curie Fellow at Newcastle University and prior to this an Assistant Professor at the University of Salzburg, Austria. She was selected Young Scientist 2017 and 2018 to attend the World Economic Forum in the People's Republic of China.