

Extended Mid-air Ultrasound Haptics for Virtual Reality

STEEVEN VILLA, LMU Munich, DE

SVEN MAYER, LMU Munich, DE

JESS HARTCHER-O'BRIEN, Facebook Reality Labs, USA and Delft University of Technology, NL

ALBRECHT SCHMIDT, LMU Munich, DE

TONJA-KATRIN MACHULLA, TU Chemnitz, DE

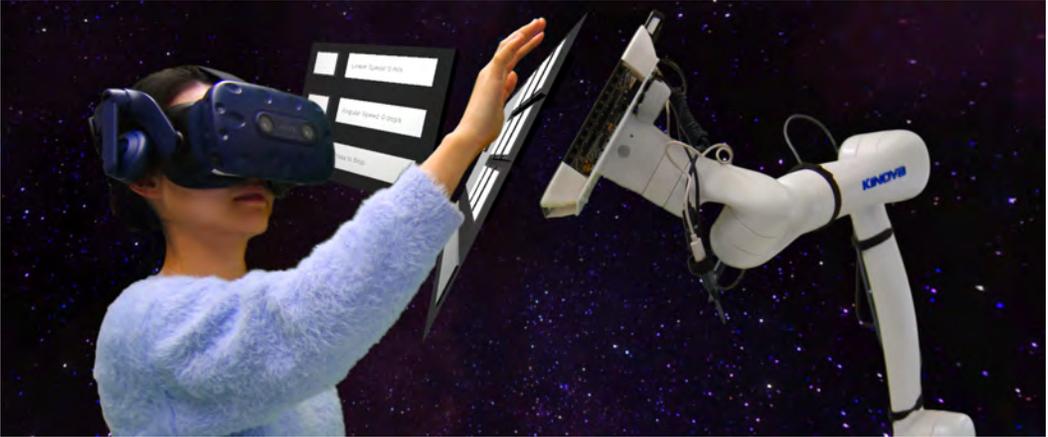


Fig. 1. We present a system to render large-scale mid-air haptic sensations by driving the ultrasound array in space using a 6-degree-of-freedom robot. In the picture: A user interacting with a VR User interface while our system provides haptic feedback.

Mid-air haptics allow bare-hand tactile stimulation; however, it has a constrained workspace, making it unsuitable for room-scale haptics. We present a novel approach to rendering mid-air haptic sensations in a large rendering volume by turning a static array into a dynamic array following the user's hand. We used a 6DOF robot to drive a haptic ultrasound array over a large 3D space. Our system enables rendering room-scale mid-air experiences while preserving bare-hand interaction, thus, providing tangibility for virtual environments. To evaluate our approach, we performed three evaluations. First, we performed a technical system evaluation, showcasing the feasibility of such a system. Next, we conducted three psychophysical experiments, showing that the motion does not affect the user's perception with high likelihood. Lastly, we explored seven use cases that showcase our system's potential using a user study. We discuss challenges and opportunities in how large-scale mid-air haptics can contribute toward room-scale haptic feedback. Thus, with our system, we contribute to general haptic mid-air feedback on a large scale.

Authors' addresses: [Steeven Villa](mailto:Steeven.Villa@ifi.lmu.de), LMU Munich, Munich, 80337, DE, Steeven.Villa@ifi.lmu.de; [Sven Mayer](mailto:info@sven-mayer.com), LMU Munich, Munich, 80337, DE, info@sven-mayer.com; [Jess Hartcher-O'Brien](mailto:Jess.Hartcher-O'Brien@fb.com), Facebook Reality Labs, Redmond, USA and Delft University of Technology, Delft, NL, ; [Albrecht Schmidt](mailto:Albrecht.Schmidt@ifi.lmu.de), LMU Munich, Munich, 80337, DE, albrecht.schmidt@ifi.lmu.de; [Tonja-Katrin Machulla](mailto:tonja.katrin.machulla@tu-chemnitz.de), TU Chemnitz, Chemnitz, DE, .

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2022/12-ART578 \$15.00

<https://doi.org/10.1145/3567731>

CCS Concepts: • **Human-centered computing** → *Virtual reality*; • **Computer systems organization** → *Robotics*.

Additional Key Words and Phrases: Virtual reality, mid-air haptics, ultrasound

ACM Reference Format:

Steeven Villa, Sven Mayer, Jess Hartcher-O'Brien, Albrecht Schmidt, and Tonja-Katrin Machulla. 2022. Extended Mid-air Ultrasound Haptics for Virtual Reality. *Proc. ACM Hum.-Comput. Interact.* 6, ISS, Article 578 (December 2022), 25 pages. <https://doi.org/10.1145/3567731>

1 INTRODUCTION

Haptic feedback improves realism and immersion in virtual reality (VR) environments. Yet, rendering realistic and versatile haptic feedback has been a long-standing challenge for researchers and designers. Over the last year, we have seen a wide range of devices to address this challenge. A large number of haptic feedback systems are body-worn, e.g., [12, 19, 21]. Lately, we have seen the utilization of drones to deliver feedback, e.g., [4, 20]. Beyond body-worn and flown feedback systems, ultrasound mid-air haptics gained popularity due to their potential to render diverse shapes and the possibility of bare-hand interaction. It uses an array of ultrasonic transducers to create focal points (wave collisions) in the air to stimulate the receptors in the user's hand. As such, ultrasound mid-air haptics has great potential due to their capabilities to provide feedback for many use cases.

Especially, the device "Ultrahaptics STRATOS Explore" from Ultraleap sparked new applications to deliver mid-air haptic feedback, cf. Carter et al. [7]. However, the major downside of this approach is the limited rendering volume, cf. Rakkolainen et al. [36]. Indeed, a recent survey on haptic interfaces for virtual reality found that the major challenge of Mid-air Ultrasound haptics is overcoming the workspace limitations [46]. Howard et al. [22] recently made the first attempt to move the haptic array using a pan-tilt platform to enlarge the interaction space, showcasing that increasing the mid-air haptics workspace opens a new set of experience designs for practitioners in VR and haptics. Combining the versatility of ultrasonic rendering with large-scale haptics allows moving from local shape to large-volume renderings, such as furniture elements in games (tables, doors, walls) or environmental elements in a room-scale volume.

With this paper, we go beyond small pan-tilt motion to enrich the mid-air haptics space by combining the great feedback capabilities of an ultrasound array with a robotic arm. Thus, this paper explores a dynamic and user-responsive system that drives an ultrasound array in three dimensions to increase the ultrasonic rendering space using a 6DOF robot. We drive our system to reach the areas where the user interacts in 3D, delivering in situ haptic mid-air feedback wherever needed. Furthermore, we use this strategy to keep the haptic display in a ready position for optimal rendering. In the following, we provide detailed explanations of our systems implementation and tracking strategies, a technical evaluation, the quantization of the perceptual characteristics, and implications for future uses based on a set of showcases. Thus, we performed a technical evaluation of our system by measuring step responses in re-positioning and reaching different rotation poses. Then we perform a psychometric experiment to determine the perceptual characteristics of the robotic-driven system in contrast to a static system. Finally, we developed seven VR showcases to highlight our system's rich versatility to render haptic feedback in an ample space in a second step. Our showcases include feedback for flat, round, and complex surfaces and moving objects, environmental haptics, user interface, and tangibles.

Overall, we showed that our system is technically feasible to control, the perceptual characteristics are not changed, and users see great potential in such a system. In detail, our technical evaluation found that the proposed system increases the original ultrasound array workspace to approximately

19.98 m^3 (With the robot mounted in a pillar), which is 363 times larger than a static ultrasonic transducer array of the same size. Our showcase user study (n=12) also showed that rendering mid-air haptics in motion could satisfactorily increase the feeling of presence and enhance VR experiences' realism and enjoyability. The contributions of this work can be summarized as follows:

- We explore a robotic-driven strategy for increasing the mid-air haptics workspace.
- We performed a technical evaluation of the system.
- We characterized our system and report the first psychophysical study evaluating the impact of a moving ultrasound emitter on perceptual thresholds.
- We report a user study involving seven scenarios to evaluate the user experience of the system.

2 RELATED WORK

2.1 Haptics In Virtual Reality

Today's virtual reality experiences feature high-quality visuals and auditory stimulation but lack haptic feedback. State-of-the-art headsets can render high-resolution images at high refresh rates [23, 32, 44] while constantly becoming more affordable for the average user. The recent launch of the metaverse [28] concept by Meta will presumably increase the interest in VR devices. Furthermore, commercial VR devices partially stimulate the haptic sense with vibration in their controllers. Nevertheless, the haptic sense is especially interesting given its complexity (mechanoreceptors are sensitive to specific mechanical behaviors) and distributed nature (haptic perception is not located in a single part of the body but across it) [17]. These two features make tactile sensations rich but also make them more complex. Several approaches have been explored to provide haptic feedback in VR; handheld devices, for example, have been provided with servos and mechanical actuators to deliver tactile sensations such as pressure or gravity [9, 27, 40, 48]. In the wearable domain, several prototypes have been proposed to render force-feedback, and stiffness [2, 12, 33]. These systems are generally built using servomotors or pneumatic valves. Specialized haptic interfaces have also been integrated into VR to improve immersion [1, 8]. However, the sensations rendered by these devices are very limited and require the user to be in contact with the haptic interface during the experience. In the last decade, mid-air ultrasound haptics has gained interest because of the variety of shapes and sensations generated using acoustic radiation [7].

2.2 Ultrasound Mid-air Haptics

Mid-air haptics refers to a set of technologies that allow mid-air haptic stimulation, mainly by using ultrasound transducers to generate pressure on the palm. Wavelengths above 20kHz are considered ultrasound and are not perceivable by human hearing or touch. Therefore, the ultrasound is sampled to frequencies below $\sim 600\text{Hz}$ [29] to stimulate human touch receptors effectively. A typical ultrasound phased array can render sensations within the 100 to 300 Hz range. The popularity of this technology has increased in the last decade within the field of human-computer interaction, likely due to two factors: the development of new rendering algorithms that facilitate the generation of convincing mid-air sensations [14, 15, 31], and the emerging market interest in such technologies [7]. Ultrasound haptics have been used to create a wide range of applications, including automotive interfaces [18], aviation [16], virtual reality [24], fluid rendering [5], emotion encoding [34], and stiffness rendering [30]. In addition to allowing bare-hand interactions, the main advantage of ultrasound-based mid-air haptics is its ability to render a complex topology flexibly. However, a major limitation of mid-air haptics is its reduced interaction volume; an array of transducers can render a limited volume and the farther from the array the lower the haptic rendering quality.

2.3 Enlarging Mid-Air Haptics Workspace

Mid-air haptics has mechanical limitations [36] (Wave propagation, transducer power) that constrain the rendering workspace. Some notable attempts have been made to overcome these limitations. For example, Suzuki et al. [41] extended the reachable area of ultrasound waves by increasing the number of ultrasonic transducers from 249 (Commercial ultrasonic arrays have around 256 transducers) to 2241 by creating two 9×9 matrixes of ultrasonic arrays. Such an approach effectively increases the workspace to a 2m cube. However, it brings computational challenges for rendering and fidelity of stimuli, and it requires a setup that can be difficult to integrate into real-world scenarios. In the same year, Brice et al. [6] proposed shifting the ultrasonic array's spatial position to achieve a larger workspace by using a robot arm. However, the array remained static during the interaction, requiring them to pre-set the haptic display positions beforehand. Specifically, they set five predefined positions for the haptic array and tested it. Furthermore, this system needed users to wait for the robot to reposition itself before interacting again, making the system stationary in strict terms. More recently, Howard et al. [22] reported a pan-tilt system that moves the array in a semi-spherical way (2 degrees of freedom). The apparatus effectively increases the workspace and renders ultrasonic sensations in motion. However, the render quality at positions far from the array (similarly to static arrays) decreases. More recently, Ariga et al. [3] reported a setup able to render ultrasonic sensations at long distances; the authors report a setup that features a haptic array and a curved reflector. Unfortunately, while the setup can move the workspace away from the array, it does not increase the workspace significantly compared to the state-of-the-art, e.g., [6, 22, 41]. We present a system that addresses the rendering volume constrain issue and enables a new rendering paradigm by actively driving the ultrasonic array to an optimal rendering position; in contrast with Brice et al. [6] concept, we optimize the array position online by following the user's hand. Furthermore, unlikely Howard et al. [22] we not only rotate but relocate the array, achieving a smaller distance between the array and the user's hand. We do that by using a serial robot arm to drive a single 256 transducer phased array in the space, therefore not requiring an extensive transducer array set like Suzuki et al. [41], thus making the setup more flexible. To the best of our knowledge, this is the setup with higher reachability as of today, as showcased in Section 4.3.

3 CONCEPT: MID-AIR HAPTICS IN MOTION

Real-world interactions are bare-hand and involve large volumes of interaction; we do not restrict ourselves to touching objects within a bounding box in our daily lives. Unfortunately, state-of-the-art haptic experiences tend to break these conditions in one way or another. For example, wearable haptics requires users to attach devices to their hands or bodies to exert a force or induce a vibrotactile stimulus, cf. Fang et al. [12]; force feedback devices require users to be in contact with the end-effector of a robot. In contrast, mid-air haptics offers bare-hand interaction but is constrained in workspace. We propose to render ultrasound mid-air haptics in motion. The proposed setup's dynamic nature allows stimulation from different angles, allowing haptic experience designers to explore a broader range of possibilities.

3.1 Setup

Figure 2 illustrates the system setup; we used a six-degree-of-freedom serial robot manufactured by Kinova robotics (Boisbriand, Canada) for driving a 256 ultrasonic transducer array (Stratos Explore) by Ultraleap (Bristol, United Kingdom) that features a LeapMotion tracker for hand tracking. To attach the haptic array to the robot, we 3D printed a lightweight interface that fits both the casing of the ultrasonic display and the robot's end-effector interface. The robot is controlled as described

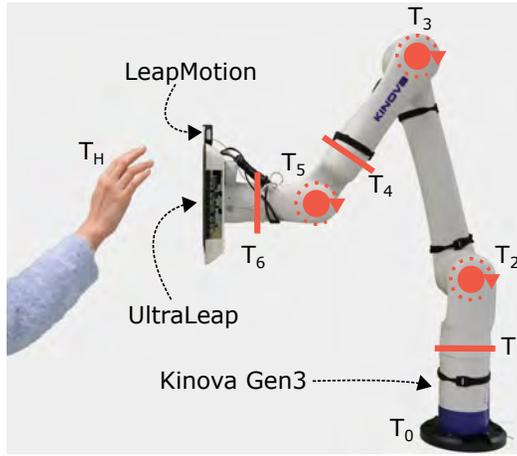


Fig. 2. The setup comprises an ultrasound array (UltraLeap), a six-degree-of-freedom Kinova [®]cobot, and a VR laptop (not visible in the picture). The final position of the haptic array is given by the transformations (T) in every joint of the robot. The transform (3d Position information) of the hand T_H depends on the cameras of the LeapMotion attached to the UltraLeap.

in Section 3.2 and integrated with the graphics engine Unity3D using the toolbox proposed by Villa and Mayer [45].

3.2 Control System

A dynamically driving haptic array setup requires online control to follow the hand and react to the VR interactions. This poses some challenges regarding computing times and control complexity. For example, the inverse kinematics loop must be calculated at a constant frequency of 1kHz (critical for robot stability). The inverse kinematics loop is the one that converts the cartesian coordinates (x,y,z) to joint rotation values (q_1, q_2, \dots, q_6) of the robot. The equations for calculating the joint values are proprietary knowledge of the robot manufacturer (Kinova). So they have to be either derived independently or used through their API. Iterative general IK solvers tend to be non-optimal for real-time calculations. We opted to use the manufacturer's proprietary calculations by creating a velocity control using the C++ Kinova Kortex API. In this way, it is not possible to set a 3D coordinate. However, directional Cartesian speed can be controlled; we measure the end-effector position and set the velocity according to the target position using a P-only control [11, 45]. Due to the different response requirements, we implemented other controllers for rotation and position (see Equation 1 and Equation 2).

$$P_{x,y,z} = P_{x,y,z,bias} + K_p * e(t) \quad (1)$$

$$R_{\theta,\psi,\phi} = R_{\theta,\psi,\phi,bias} + K_{pR} * e(t) \quad (2)$$

Where $P_{x,y,z}$ and $R_{\theta,\psi,\phi}$ represent the pose of the end-effector (Cartesian position and rotations), $P_{x,y,z,bias}$ and $R_{\theta,\psi,\phi,bias}$ are the standard offsets in position and rotation, K_p and K_{pR} are the controller components (PD) and finally $e(t)$ represent the error between target position/rotation and current position/rotation.

Using this approach, we preserve the 1kHz stability on the robot base. Figure 3 illustrates the loops involved in the systems; the Unity plugin we developed is the interface that harmonizes the 1kHz loop with the 100Hz loop of the graphic engine (Physics fixed update of Unity3D). The

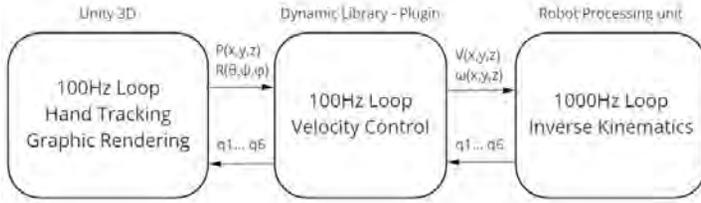


Fig. 3. The setup involves three layers of control, first: low-level inverse kinematic control, provided by the manufacturer, second, our implementation of proportional velocity control, and third, a closed-loop control based on the target tracking.

plugin sends cartesian (\vec{v}) and angular velocities ($\vec{\omega}$) and reads the robot's joint rotation positions to calculate the forward kinematics (measured cartesian position). The joint rotations are also communicated to the graphic engine to reconstruct the robot's real-world position in the VR environment and provide feedback to the user.

3.3 Tracking

The hand position guides the robot's end-effector. We divided the tracking into two levels. 1. short-range: guided by LeapMotion and 2. long-range, guided by a vibe tracker (but could be replaced by any tracking device).

For the short-range tracking, we converted the hand position from the LeapMotion Coordinate System (LCS) to the World Coordinate System (WCS) using the forward kinematics of the robot, which can be calculated using the transformations for every joint of the robot. Equation 3 describes the expression for moving from LCS to WCS, where the last transformation corresponds to the rotations and translations detected by LeapMotion and T_0 to T_6 can be found in Kinova's gen3 manual¹. The encoders in the robot can measure the values of such transformations. In Figure 2 we illustrate the position of every transformation (T_N) in the robot.

$$T_{0,H} = T_{0,1}T_{1,2}T_{2,3}T_{3,4}T_{4,5}T_{5,6}T_{6,H} \quad (3)$$

For the long-range tracking, we used a VIVE tracker attached to the wrist of the user. This allows us to prevent collisions and guide the robot whenever the short-range tracking fails to capture and interpret the user's hand. In addition, this allows us to overcome the limited tracking space and quality of the LeapMotion. The VIVE tracker is only a fallback option when the primary tracking fails, so it can be replaced by alternatives that do not require the user to wear any device (following the rationale of mid-air haptic interaction), for example, Realsense Cameras or AI-based IK solvers that can provide information about the user's arm location with relation to the robot.

3.4 Guidance Strategy

The hand is the target point of the end-effector. We first calculate the end-effector's target position (Ultraleap device) and then the rotations. We trace a vector from the robot's base to the hand to calculate the robot's target position. Then we determine an offset H_{off} , which is especially important since it can determine the ultrasonic rendering's intensity and quality. The β angle offset can be used to define the array's rendering direction (facing downwards or upwards). Finally, the value V_{off} is the height correction necessary to keep the hand in the center of the ultrasonic

¹Kinova Gen3 User Manual

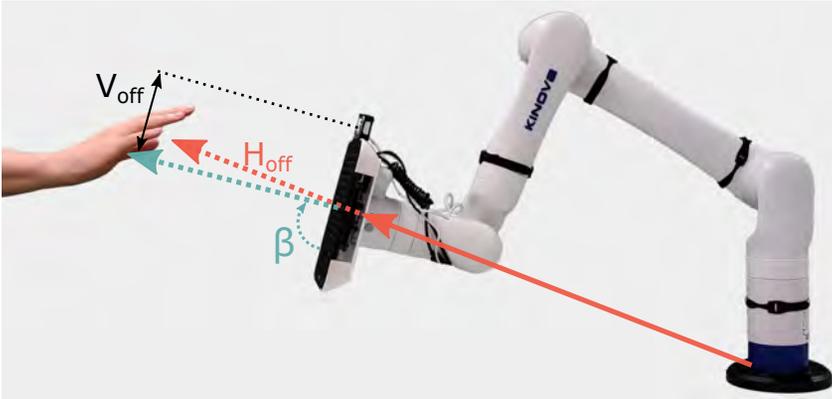


Fig. 4. We guided the haptic array to follow the user’s hand closely. The tracking strategy involves tracing a vector from the robot base to the palm and calculating an offset H_{off} . We also added a vertical correction factor V_{off} to center the hand in front of the haptic array instead of the leap motion coordinate system.

array instead of the LeapMotion sensor. This set of parameters must be tuned depending on the VR scenario to keep the hand within the tracking volume. These values equally impact the short-range and long-range tracking of the hand. It is important to highlight that the rotations of the end-effector are faster than the translation movements. For this reason, we calculate the rotations of the end-effector based on the vector from the end-effector to the hand instead of the vector from the base to the hand. This allows the system to act quickly to adjust the haptic feedback within the space in front of the ultrasonic array, allowing the user to explore an object quickly.

4 TECHNICAL EVALUATION

The proposed system is designed to interact closely with humans, sharing the same space, executing similar movements, and following the user’s hand, but most importantly, all of this will happen while the user has reduced awareness of the physical configuration of the robot. Therefore, the control strategy has two primary design criteria. 1. safety of use in shared environments with



Fig. 5. Rendering reachable volumes: nominal space of the robot $3.29m^3$ (Blue), space within ideal rendering quality $7.42 m^3$ (Green) and maximum rendering space $8.27m^3$ (lower quality, in Red). This is $19.98 m^3$ Volume in total.

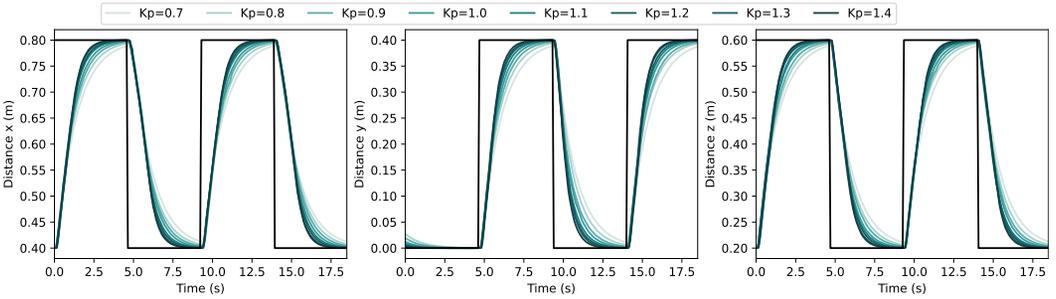


Fig. 6. Translational step response in the x, y, and z-axis: We conducted a series of automated technical tests to evaluate the speed of response of the system. The results show an over-damped system for translations.

humans, and 2. speed performance to translate the haptic array in the space following the hand trajectory. We implemented a P-only control [11] in the translation and rotation speeds to reach the hand's real-time position. As a design consideration, we used the base workspace of the haptic array as the main design driver, given that the base rendering volume allows a range of flexibility when moving the robot since it does not have to strictly maintain a pose relative to the hand. In other words, it means that the robot's end-effector does not have to mirror the hand's rotation and position, but instead, it has to keep the hand within the Ultraleap's rendering boundaries. Therefore, the system's response in translation and rotation becomes important to analyze.

4.1 Step Response for the Translation Axis

The speed at which the array follows the hand is essential to keep the hand inside the rendering volume. To model the translational behavior of the system, we executed an automated test of the robot's speed on every translational axis. The robot was programmed to move from position A to B ($A - B = 40\text{cm}$), and then we recorded the time it needed to reach the final position; we repeated this procedure for several proportional (K_p) values ranging from 0.7 to 1.4. Figure 6 shows the step response measured using these values.

4.2 Step Response for Rotation Axis

Similarly, we evaluated the system's rotation axis; in this case, we focused on pitch and yaw since roll has negligible contributions to the rendering (the roll rotations of the mid-air stimuli can be achieved by the transducer modulation). We rotated the end-effector from an angle α_a to α_b ($\alpha_a - \alpha_b = 90^\circ$). The procedure was repeated using five different K_{pR} values. Figure 7 illustrates the step response of system rotations. As is visible, in this case, the system tends to be under-damped (similar to Howard et al. [22]). However, this does not compromise the user's safety since rotations do not alter the distance from the hand to the array, and the joints q6 and q5 (actuators for T_6 and T_5 , see Figure 2) will always be in the opposite direction to the hand. Fast rotation movements are more relevant than translation movements to keep the hand inside the rendering volume.

4.3 Improvement in Workspace Dimensions

The proposed system can drive the array in a spherical radius of 93.4 cm (length from T_2 to the Ultraleap end-effector Figure 2).

However, the rendering volume is significantly bigger. If we consider a maximum ideal rendering distance of array-to-hand of approximately 45 cm [43] sensations could be rendered up to a 1.384 m radius. Moreover, if we calculate the rendering workspace using the absolute maximum rendering

Table 1. Comparison in workspace of different mid-air haptics systems.

Author	Strategy	Rendering Volume [m^3]	Gain
Ultraleap (measures from [22])	Static Array	0.055	1
Howard et al. [22]	Pan-Tilt	0.95	17.2
Brice et al. [6]	Switch Array Positions	1.5	27.2
Suzuki et al. [41]	Multiple Static Arrays	2.	36.3
Proposed System	Robotic Driven Array	19.98	363.2

distance (Low-quality rendering area), the rendering space would increase to a 1.674 m radius. This maps to $3.29m^3$ of rendering workspace in the nominal volume of the robot (blue volume in Figure 5). In this volume, the ultrasound array can always reach the hand within the ideal rendering distances. Additionally, a volume of $7.42 m^3$ can be rendered with the full robot extension within the ideal ultrasound array to hand rendering distances (green volume in Figure 5). On the boundaries of the ideal rendering distances, an additional volume of $8.27m^3$ can be rendered in a lower quality (Red volume in Figure 5). The total volume renderable by the proposed setup (including lower quality volumes) is $19.98 m^3$. This represents an increase of 21 times over Howard et al. [22] ($0.95 m^3$), 13.32 times over Brice et al. [6] ($1.5 m^3$), and 9.9 times over Suzuki et al. [41] ($2 m^3$), making it to the best of the authors' knowledge, the biggest rendering volume achieved in mid-air haptics. See Table 1 for values relative to a standard static array.

5 PERCEPTUAL CHARACTERIZATION

This study aimed to compare the perceptual impact of rendering ultrasonic haptic feedback using a moving array (referred to as dynamic ULTRASONIC ARRAY) versus a stationary array (referred to as static ULTRASONIC ARRAY) using a robotic manipulator. We executed this comparison to gain insights into the perceptual differences between the proposed system and the state-of-the-art in dynamic rendering setups. In detail, we conducted three experiments that cover the most relevant movement combinations of ULTRASONIC ARRAY, HAND MOVEMENT, and FOCAL POINT.

5.1 Hypotheses

Mid-air ultrasound haptics is rendered using stationary setups. Factors such as hand speed, hand distance from the ULTRASONIC ARRAY and the propagation of the mechanical waves in the air are typically calculated by observing the assumption that the ultrasound actuators are not being displaced from their original position. As we evaluate the same device in a dynamic setting, we

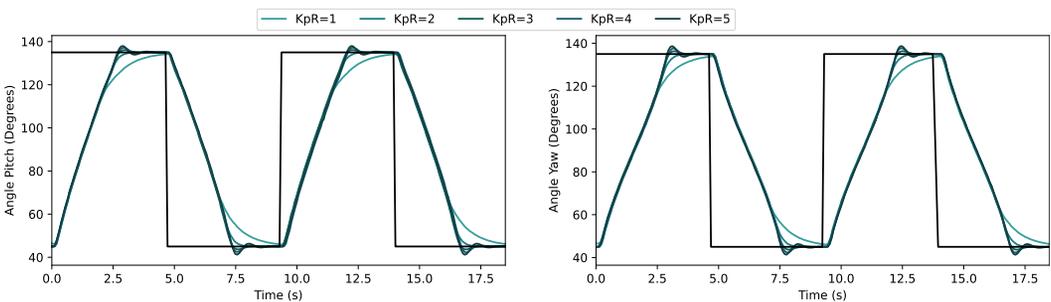


Fig. 7. Rotational step response for pitch and yaw rotations: We conducted a series of automated technical tests to evaluate the speed of the system's response. The results show an under-damped system for rotations.

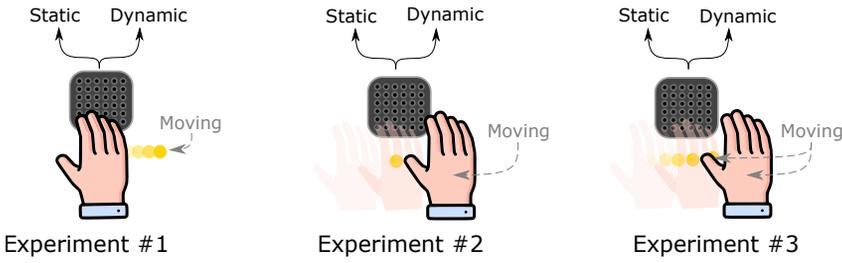


Fig. 8. Experiment description: Experiment one evaluated a moving focal point vs. a static hand, Experiment two evaluated a moving hand vs a static focal point; and Experiment three evaluated a moving hand and moving focal point. All the experiments had static vs. dynamic arrays as conditions

expect the performance to be lower regarding the static ULTRASONIC ARRAY. In particular, we anticipate that the dynamic ULTRASONIC ARRAY will perform worse when rendering stationary elements for active touch since it has to compensate for the robot's movement. As for the moving objects, we expect the dynamic ULTRASONIC ARRAY to closely follow the performance of the static array in both dynamic and static HAND MOVEMENT. However, we expect that the decrease in quality is not significantly high and that the significant gain in the workspace can compensate for such loss.

In Summary:

- **H1:** dynamic ULTRASONIC ARRAY will perform similar to ULTRASONIC ARRAY for passive touch.
- **H2:** dynamic ULTRASONIC ARRAY will perform worse than static ULTRASONIC ARRAY in active touch when the object is stationary in the 3D space.
- **H3:** dynamic ULTRASONIC ARRAY will achieve similar performance to ULTRASONIC ARRAY for active touch with moving objects.

To address these hypotheses, we designed three experiments: **Experiment #1** Evaluates dynamic FOCAL POINT and static HAND MOVEMENT in terms of static and dynamic ULTRASONIC ARRAY. **Experiment #2** Evaluates static FOCAL POINT and dynamic HAND MOVEMENT in terms of static and dynamic ULTRASONIC ARRAY. Finally **Experiment #3** Evaluates dynamic FOCAL POINT and



(a) Stimuli presentation screen: In the screen, guiding hand, visual focal point presentation, user's hand

(b) Forced choice screen: In the screen, forced choice questions, possible answers, and the user's hand

Fig. 9. View of the VR environment used for the perceptual characterization.

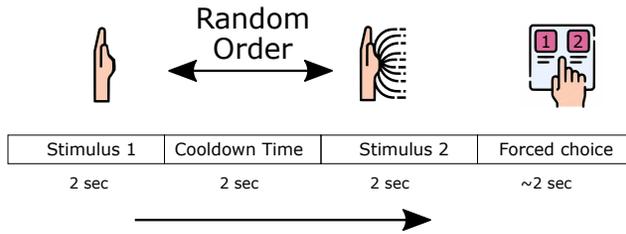


Fig. 10. Stimuli presentation scheme: We present two stimuli sequentially, one of them contains the actual ultrasound feedback, the other is empty, the order of them is random. After the two stimuli were presented, we asked participants in a 2AFC question where the actual stimuli was presented.

dynamic HAND MOVEMENT in terms of static and dynamic ULTRASONIC ARRAY, for graphic details, please refer to [Figure 8](#).

We defined the aims of each experiments as follows:

- **Experiment #1:** Dynamic FOCAL POINT, Static HAND MOVEMENT This experiment was designed to study the passive touch scenario, which is the case where the observer’s hand is not moving but the element that they are touching moves. This is reflected in the current setup by moving the focal point across the bases of the proximal phalanx (Palmar digital creases). The decision to stimulate the hand across the Palmar digital creases is motivated by the density of mechanoreceptors involved in the perception of ultrasonic stimulation (Rapid Adapting Afferents) reported for the hand [10, 25]. The closer to the fingertips, the higher the innervation density; therefore, we selected the largest continuous area of the hand closer to the fingers. Volunteers were instructed to hold their hand still in the position indicated by the virtual guiding hand (more details about this virtual guiding hand in [Section 5.3](#)).
- **Experiment #2:** Static FOCAL POINT, Dynamic HAND MOVEMENT The second experiment aimed to study active touch with a static object setting. This is, the observers actively move their hands but the element they are touching is static. In the setup, this translates into a focal point static in the 3d space and the participant’s hand moving from left to right.
- **Experiment #3:** Dynamic FOCAL POINT, Dynamic HAND MOVEMENT In Experiment #3 This experiment was designed to study the effects of having several moving elements in play. In the setup, this was implemented by moving the focal point from left to right and asking the volunteers to follow the same movement with their hands.

We omitted the combinations of static FOCAL POINT, static HAND MOVEMENT and static ULTRASONIC ARRAY since this case has been widely explored in literature and does not provide additional information regarding the set of hypotheses addressed in this investigation.

5.2 Experimental design:

The three experiments aim to compare the perceptual quality of rendering of dynamic versus static ULTRASONIC ARRAY. Perceptual quality can be expressed as the detection sensitivity and detection threshold. Both detection sensitivity and threshold can be inferred from a psychometric function given a combination of parameters defined for **Experiment #1**, **Experiment #2**, and **Experiment #3**. We opted for an adaptive 2-down-1-up staircase method to sample the data required to calculate the psychometric functions. All the staircases started at 100% rendering intensity and adapted at a step of 20% during the first two reversals; then, the step size decreased to 5%. The order of experiments and conditions were randomized using Latin Square. Volunteers executed two

conditions per experiment (static and dynamic ULTRASONIC ARRAY). The condition ended when volunteers achieved nine reversals, or 35 trials.

Task: Figure 10 illustrates the structure of the trials: volunteers were presented with two stimuli, one of them was an ultrasonic vibration, and the remaining one was empty. That is, it did not contain any vibration. We added a 2 second interval between stimulations to allow the robot to return to the start position before the next stimulus. The stimuli were presented in random order. In a two-alternative forced-choice decision (2AFC), volunteers had to identify which stimuli contained the vibration. The trial ended after the volunteers had contact with both stimuli. We implemented contact and movement checks to verify that the task was executed properly in every experiment. *Experiment Parameters:* The hand was placed at a distance of 25 cm from the haptic array, the speed of movement of the focal point, hand, and array was 0.3 m/s and the total distance rendered was 30 cm.

5.3 Setup

The experimental setup was composed of an UltraLeap for hand tracking, a Kinova Gen3 6 DoF for driving the array, A Valve Index VR headset, and a VR-Ready laptop (Acer Predator Helios 300 with a NVidia RTX 2070). We implemented a VR environment for conducting all three experiments. The ultrasonic array was positioned vertically in front of the volunteers. Volunteers placed their hands on an armrest so they did not get tired of holding their hands in a parallel array position. Finally, we rendered a virtual (guiding) hand in VR that guides the volunteers through all the experiments. (i.e., hold the hand static or Actively explore the stimuli).

5.4 Participants and Procedure

Eight people participated in the study, one volunteer was female, and one participant was left-handed. Ages ranged from twenty-three to thirty-three years old, with an average age of 28.8 ($SD = 3.3$). None of the volunteers reported any skin conditions that impaired their capability to discriminate tactile sensations. One participant was left-handed. Volunteers were compensated with 10€/h for their participation in the study. *Exclusion Criteria:* We targeted young adults (less than forty years old) according to Peters et al. [35] to guarantee skin sensitivity. Additionally, people with reported skin conditions were excluded from the experiment.

Volunteers were instructed to read through the informed consent and fill up a demographics survey. Next, the experimenter explained the tasks, the elements of the study, and the stages of it. When the participant felt comfortable with the setup, we asked them to put on the headset and follow a tutorial about the experiment. The tutorial included tasks like matching the virtual and real hands, moving their hand from left to right, and selecting one stimulus in the same way they would in the experiment (2AFC). After the tutorial, the first assigned experiment started. Once the participant went through all the experiments and conditions, we debriefed them on the experiment's details. Volunteers were compensated with 10€/h. The study's average duration was one hour.

6 PERCEPTUAL CHARACTERIZATION RESULTS

We calculated the psychometric functions for every experiment and condition using bayesian estimation. Precisely, we followed the approach by Schütt et al. [39]. We modeled all the psychometric functions using normal cumulative sigmoids. The points of subjective equality (PSE) are calculated at the 75% chance of correct responses; this value also reflects the absolute detection threshold (T_A). Similarly, the middle point between the PSE and the higher chance of a correct answer is taken as the detection threshold. Finally, the just noticeable difference (JND), which represents the lower change in intensity that an observer can detect, was calculated as the difference between the PSE

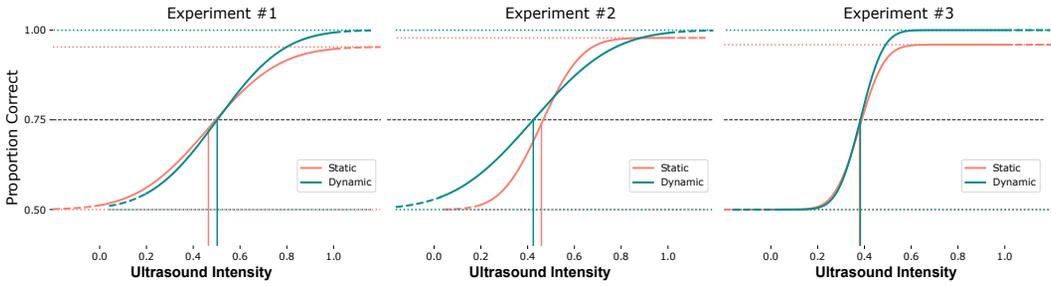


Fig. 11. We tested the system in seven different scenarios exposing its capabilities and the challenges of the paradigm of array-in-motion mid-air rendering.

and detection threshold. All the data from this experiment is available in the supplementary material. For the sake of simplicity, the intensity values are presented as a percentage of the maximum rendering power of the ultrasonic array; peaks on the STRATOS platform range from 0 – 1125Pa.

Results for passive touch (**Experiment #1**) show a slightly lower absolute threshold (T_A) for static ULTRASONIC ARRAY ($T_A = 0.415$, Confidence Interval (CI) = 0.630) than for dynamic ULTRASONIC ARRAY ($T_A = 0.433$, CI = 0.548), however, the detection sensitivity (Sen_{Static}) was higher in dynamic ULTRASONIC ARRAY ($Sen_{Dynamic} = 0.885$) than static ULTRASONIC ARRAY ($Sen_{Static} = 0.745$). Consequently, dynamic ULTRASONIC ARRAY presents a better JND than static ULTRASONIC ARRAY ($JND_{Static} = 0.165$, $JND_{Dynamic} = 0.089$). In the case of active exploration of stationary objects (**Experiment #2**); The absolute threshold was, in fact, lower in the dynamic ULTRASONIC ARRAY condition ($T_A = 0.454$, CI = 0.747) than in the static ULTRASONIC ARRAY condition ($T_A = 0.473$, CI = 0.650). Nonetheless, sensitivity and JND were not better than static ULTRASONIC ARRAY ($Sen_{Dynamic} = 0.592$, $Sen_{Static} = 1.123$, $JND_{Dynamic} = 0.090$, $JND_{Static} = 0.047$). Finally, in **Experiment #3** static and dynamic ULTRASONIC ARRAY showed an identical behavior in all metrics, however, dynamic ULTRASONIC ARRAY performed marginally better sensitivity and JND ($Sen_{Dynamic} = 2.573$, $Sen_{Static} = 2.159$, $JND_{Dynamic} = 0.045$, $JND_{Static} = 0.056$) but also marginally worse in Absolute threshold; namely static ULTRASONIC ARRAY ($T_A = 0.361$, CI=0.420) and dynamic ULTRASONIC ARRAY ($T_A = 0.366$, CI = 0.404).

In all, an overarching question is whether rendering ultrasound mid-air haptics in movement is notably different from the static case, therefore, impacting the rendering quality. We assessed this question by analyzing the likelihood of the two models to be similar (Bayes Factor in favor of the null hypothesis) for the general model across all the experiments using a Bayesian Paired T-test. We found that the Bayes Factor for T_A is 3.8 ($BF_{01} = 3.864$, $Error = 0.015\%$), which means that the distribution of perceptual thresholds obtained from our characterization is 3.8 times more likely to occur under the null hypothesis than under the alternative hypothesis. This likelihood is closely followed by the JND distribution ($BF_{01} = 3.886$, $Error = 0.015\%$). In the case of the sensitivity (Sen) the distribution of data collected is 1.67 ($BF_{01} = 1.672$, $Error = 0.022\%$) times more likely to happen under the case where the systems are similar. This provides support to the fact that even if in specific scenarios, dynamic ULTRASONIC ARRAY can have mechanical advantages and disadvantages, the two systems are likely to be equivalent in terms of absolute detection threshold and just noticeable difference.

6.1 Improvement in quality

The psychometric modeling of the dynamic setup evidenced similar performance in relation to a static ULTRASONIC ARRAY. In specific cases, it did not reduce the rendering quality but increased it.

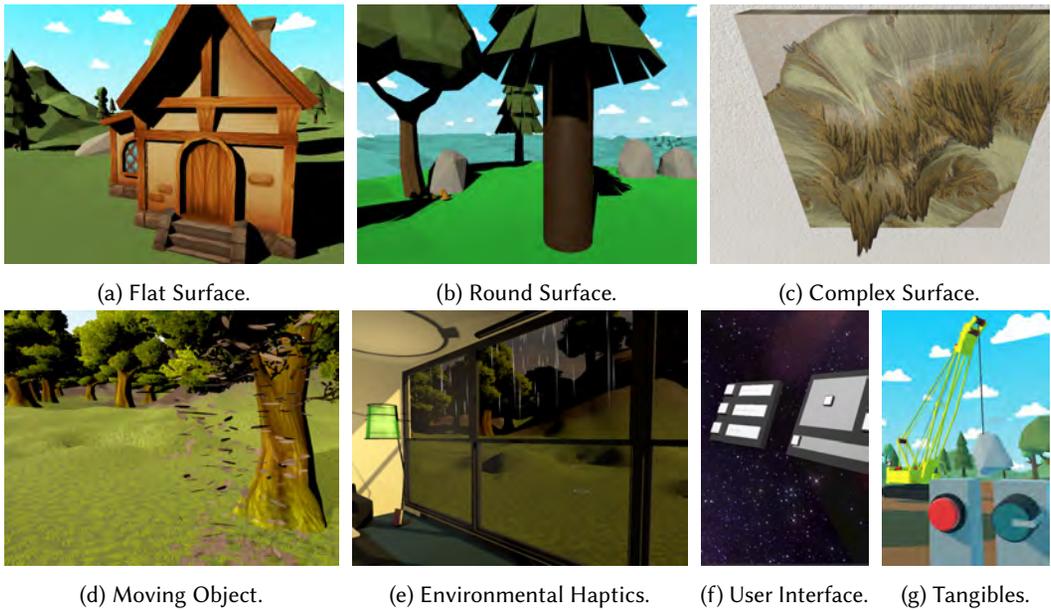


Fig. 12. We tested the system in seven different scenarios, exposing its capabilities and the challenges of the paradigm of array-in-motion mid-air rendering.

We discovered that the dynamic rendering mode introduces a tradeoff between rendering quality (Sensitivity & JND) and perceived intensity (Absolute threshold). For dynamic objects (**H1** and **H3**), it did not impact the rendering quality; in fact, in **Experiment #1**, sensitivity and JND values were better for dynamic ULTRASONIC ARRAY than static ULTRASONIC ARRAY. These results go in the direction of **H1** and **H2**, but also show the possibility of improving the rendering quality by moving the array at a similar speed to the moving object. However, this introduces mechanical limitations for fast-moving objects; further investigations are thus required in this regard. Finally, the dynamic setup did have an evident impact on the rendering quality for static objects. This was expected given the compensation in the position that the ULTRASONIC ARRAY has to perform in order to hold still the focal point while moving. Therefore, such behavior is encompassed with **H2**.

7 SHOW CASES

We developed seven showcases to present the systems' capabilities for rendering large-scale mid-air haptics. Thus, we designed all showcases to reflect the need for haptics in various scenarios and showcase the potential of the proposed system. We classified haptics in VR into three groups: general object haptics, environmental haptics, and haptics for user interfaces. The simplest haptic feedback is to present a *Flat Surface*, more advanced are *Round Surface*, and finally, in the group of general objects, we have *Complex Surface*. A special case of general object haptics is *Moving Objects*, which heavily relies on stable tracking. Moreover, they also deliver a feeling of energy in an environment, such as falling leaves activated by wind. Such feedback can also be ubiquitous in the case of rain, and thus, we identified the group *Environmental Haptics*. Finally, the last group we identified is haptics for user interfaces. Here, we see *User Interface* with buttons, a text field, and a virtual keyboard. On the other hand, we identified 3D user interfaces that align with *Tangibles*.

Flat Surface: The Door

Flat surfaces are one of the most commonly used elements that help build more complex geometric objects. However, their simplicity does not imply that they are not helpful as they are widely used in VR applications, e.g., walls, floors, doors, and tables. Thus, flat surfaces are essential for VR gaming and collaborative environments.

In detail, we used a flat 2m tall 1m wide door to showcase how the proposed system can deliver haptic feedback on a large surface, see [Figure 12a](#). We had the door be part of a building situated in a low-poly flat-shaded world to support immersion. The world consisted of simple elements such as trees and mountains.

For rendering this flat surface, we used the forcefield rendering mode included in the Unity Ultraleap package that calculates the hand's intersections and renders focal points in the collision coordinates. In the flat surface case, we positioned a plane in front of the door, so every time the participant touches the door's surface, the contact points are stimulated.

In this scenario, we expected participants to use a flat open hand in active exploration of the flat surface.

Round Surface: The Tree

After flat surfaces, round surfaces use the third dimension, allowing for more complex combined objects. While they give designers of VR experiences a greater set to pick from, they are harder to realize from a haptic rendering perspective. Examples of such objects are door handles, advertising pillars, or trees.

We selected a tree trunk. Again, we can show the system's capabilities to render large-scale haptics, [Figure 12b](#). As we only focused on scale, we used a simple large cylinder as a tree trunk with a diameter of 0.5m. The greater scenario of this showcase is a small island with infinite see-around. We placed some other trees and some stones on the island. We again used a low-poly flat-shaded visualization with a cartoon-like skybox.

For rendering round surfaces, we used a similar approach to the first scenario. Still, this time we moved the plane tangentially based on the user's fingers' contact points and the cylinder representing the tree's trunk.

In this showcase, we expected participants to explore the surface using a more relaxed and concave hand position.

Complex Surface: The Topology Map

All objects can be modeled as combined complex structures with flat and round surfaces, sometimes comprising thousands of small surfaces. However, this poses several new challenges; for instance, the need for haptic feedback in different directions, e.g., when grabbing smaller objects or the fast switching between haptic feedback directions while moving over a complex surface. For such a scenario, we used a map exploration scenario. Here, touching the different faces requires haptic feedback to move quickly and precisely.

We pick such a map exploration task in an education scenario for showcase, see [Figure 12c](#). Therefore, they hang the map on a classroom wall as one would hang a classical 2D map. The environment was a full-scale classroom with desks, chairs, blackboards, and shelves. Large open windows reveal the skyline of the surrounding city for an immersive impression.

Using a similar approach to round surfaces, we moved the rendering plane based on the contact points, generating focal points in the intersections between the fingers and the mountains' colliders.

For this showcase, we expected participants to rely mainly on their fingertips for the exploration, given the number of details presented by the topology map.

Moving Objects: Falling Leaves

We move away from a static showcase, with the next showcase toward moving objects. Here, we show that the system can provide static haptics at a new position and render constant haptic feedback while the user's hand follows a moving object.

We designed a tree in autumn where leaves fall to the ground to experience moving objects' constant feedback. This showcase is unique as it shows that the proposed system can provide continuous feedback in a large volume and render feedback to many small targets – the many leaves falling. To provide the users with an atmospheric autumn day, the sun is set low with long shadows, more realistic grass textures, and many trees surrounding the user's distance to deliver the impression of a glade in a forest, see [Figure 12a](#).

We generated circular sensations on the palm every time the hand collided with a leaf to render the leaves. The coordinates of the circular sensations followed the leaf movement instead of the hand translation.

In this showcase, we expected participants to quickly change the configuration (from open hand to grasping) and position of their hands in order to catch the leaves.

Environmental Haptics: Rain

For cinematic experiences and storytelling, it is important to deliver haptic feedback for elements in the environment and render environmental haptics such as wind, rain, and snow; cf. Tatarchuk [42]. This further challenges the system as now the tracking and feedback need to provide constant haptic feedback no matter where the user is and how fast they move.

We designed a living room with a large open window to explore environmental haptics for storytelling and not overload the user with haptic feedback. Outside, the user is presented with a rainy summer night [Figure 12e](#). Thus, whenever the users hold their hands, they will get heavy rain rendered onto their hands. Again, the large-scale haptic rendering system allows users to explore the entire area outside the window widely.

We render rain by generating scattered focal points in the hand every time the participant takes their hand out of the virtual window. The haptic feedback is presented as long as the user's hand is outside of the window. Furthermore, the feedback is rendered from below the user's hand and not from above. We rendered this scenario in this specific way given the limitations of the system to placing the array facing upside down at fast speeds when the users flip their hands.

We wanted to explore free-hand exploration across the rendered volume; specifically, we wanted to see the impact of participants flipping their hand in the opposite direction of the array and its impact on enjoyment and perception of the feedback.

User Interfaces: Digital UI

Graphical User Interfaces with buttons, sliders, and scrollbars are standard in traditional desktop applications and to navigate through virtual 3D interfaces; cf. Zhang et al. [49]. While user interface elements are typically flat surfaces, we present them in their own showcase due to their importance in interface design. To enable volunteers to experience the feedback provided by the system, we used a simple interface with buttons placed into an outer space skybox, [Figure 12f](#). This enabled volunteers to experience the rapid clicking of the same buttons and the fast switching between buttons.

Digital interfaces were previously rendered using circular focal point movements on the palm. We followed the same approach for the buttons and the sliders on the interface.

In this showcase, we evaluate a popular use of mid-air haptics: interface interaction. Specifically, we wanted to check one-finger interactions. Therefore, the interface features small buttons so the participants cannot interact with them using their full hands.

Tangibles: Button & Knob

A special form of user interface elements is tangible UIs; here, in contrast to traditional user interface elements, we see 3D elements such as levers, buttons, and knobs; cf. Fang et al. [12].

We compiled a scene in which a user can steer the rotation of a crane using a button and a knob, see Figure 12g. While the knob allows manipulating the crane's current orientation, the buttons allow switching the direction of the rotation. While this is a toy example, the focus is on the tangible elements and not on the complexity of the task.

The haptic rendering in this environment aimed to simulate coherence between the visual and haptic feedback; for the push button, we simulated pressure on the palm whenever the user touched the buttons. In the case of the knob, we rendered a dial sensation on the palm, rotating at the same speed as the knob.

In this scenario, we wanted to complement the interface interactions by adding elements that must be used using the full hand and not only one-finger interactions, such as a big push button and a hand-sized knob.

8 SHOW CASE EVALUATION

To evaluate the system and its capabilities to render haptic feedback at a large volume, we invited 12 volunteers to interact with the seven showcases above. While the proposed system can render different haptic feedback patterns, our investigation solely focused on providing feedback in a large space, allowing the user to move around an object or area of interest freely. Volunteers experienced all seven SCENES with *Haptic Feedback* and with *No Feedback*. We randomized the order of the scenes within each haptic condition and counter-balanced the independent variable HAPTIC.

While there are haptic feedback systems out there, we opted for the *No Feedback* as this is the only real possible comparison. Haptic feedback systems using vibration do require a controller in the hand of the user. The alternative to controllers is body-worn haptic systems (cf. Fang et al. [12]); however, they augment the user too. On the other hand, our system does not occupy the user's hand, nor is it body-worn. Thus, currently, there is no system that provides haptic feedback on this scale. Thus, the only fair comparison is a *No Feedback* comparison.

8.1 Apparatus

We used the system in a silenced $\sim 30m^2$ large room. The robotic arm was mounted on a desk situated in the center of the room. We provided visual feedback of the robot end-effector inside the virtual environment to ensure participants' safety and avoid breaks in immersion due to unexpected elements in the real world. A semi-transparent red square represents the end-effector. In all the conditions, participants had feedback about the position of their hands.

8.2 Participants and Procedure

In total, we invited 12 volunteers (6 female, 6 male), with an average age of 26.2 ($SD = 2.2$). Two volunteers own VR headsets, and two stated they had no prior experience with VR. In addition, all but one participant were right-handed. Finally, three volunteers indicated that they had experienced haptic devices beyond common vibration feedback (e.g., smartphones) beforehand.

After welcoming volunteers, we explained the study and answered any questions before starting the study. One experiment guided the volunteers through the survey during the study, which took around 40 minutes. Volunteers started either with or without haptic feedback. Here, volunteers were

asked only to use their right hand to explore the environment. The exploration tasks did not include any dexterous manipulation, and task performance was not measured. Therefore, the handedness of the participant would not impact the execution of the task. The volunteers experienced each scene for around 1 minute, in which they were free to move and explore the target objects. On a 7-point scale, we asked volunteers to rate the following four questions: "The experience was enjoyable?" "I felt a strong feedback," "The virtual object felt real," and "I did not feel anything when my hand touched the object." The first one was inspired by the User Experience Questionnaire [38], and the other three by other haptic investigations, e.g., [12, 22, 26]. Additionally, after the volunteers finished one haptic condition, we asked them to fill out a Igroup Presence Questionnaire (IPQ) [37]. Finally, we noted down relevant comments and reactions the volunteers had during their participation. While the robot is designed for humans as an assertive robot and, thus, should not harm a human in a way, the experimenter always had an emergency button within arms reach. Moreover, to protect volunteers from the ultrasound, we asked them to wear earplugs, which we provided.

8.3 Results

Overall, all volunteers reported an excellent level of positive feedback on the use of the system. In particular, they were positively surprised by how well they could experience the complex surface – the mountain; for instance, P5 stated, "[the mountain] was the coolest." Other comments include P6 "I felt the leaf falling through my hand." Moreover, P12 stated that the environment feels not responsive; here, P12 described being disconnected from the virtual world.

With the designs, e.g., using low-poly flat-shaded scenes, we aimed to not focus on the feedback itself but on the system's ability to render feedback in a large space. However, volunteers would have liked to feel the tree's bark and stated that this was a missed opportunity not to render the detailed texture (P11); cf. Freeman et al. [13]. Moreover, P3 stated they felt a mismatch between the real world and VR as the feedback was mid-air and not force feedback. Finally, P1 and P2 reported tracking issues that we could trace back to the LeapMotion, as these two volunteers had particularly small hands.

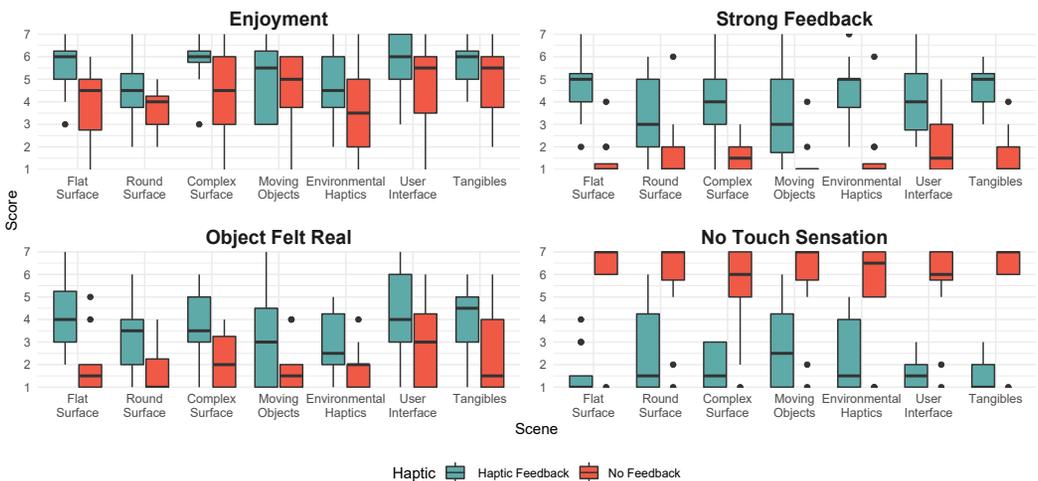


Fig. 13. The score of the four questions independent for HAPTIC \times SCENE.

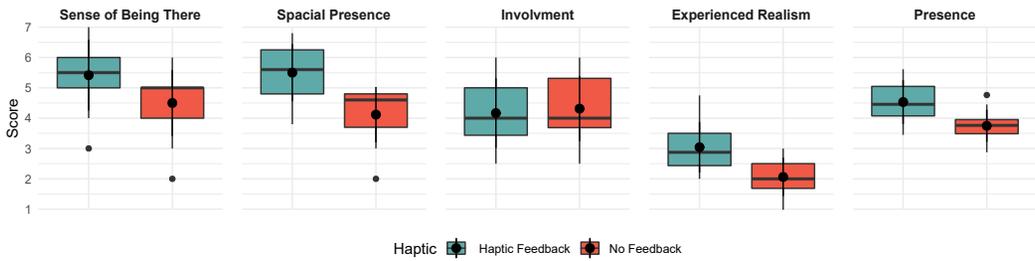


Fig. 14. The rating of the four sub-scales of the IPQ questionnaire [37] and the combined presents the score for the two conditions *Haptic Feedback* and *No Feedback*.

8.4 Quantitative Feedback

As all of the questionnaire results for enjoyment, strong feedback, object felt real, and no touch sensation were not normally distributed ($W = 0.920, p < .001$; $W = 0.854, p < .001$; $W = 0.890, p < .001$; $W = 0.800, p < .001$; respectively) we aliened and ranked the data first using ARTool [47].

First, we investigated whether HAPTIC or SCENE significantly influenced **enjoyment**. Therefore, we conducted a two-way ANOVA. The analysis revealed a significant effect for HAPTIC and SCENE ($F_{1,143} = 32.196, p < .001$; $F_{6,143} = 3.410, p < .004$; respectively). However, we found no significant interaction effect ($F_{6,143} = .689, p = .659$). Post hoc t-tests revealed that only *Round Surface* vs. *Tangibles* is significantly different ($t(143) = -3.305, p = .025$), all others are ($p > .05$).

Next, we were interested if the **strong feedback** was significantly influenced by HAPTIC or SCENE. The analysis revealed a significant effect for HAPTIC and SCENE, and a significant interaction effect ($F_{1,143} = 267.386, p < .001$; $F_{6,143} = 2.547, p = .023$; respectively). Due to the significant interaction effects, we only looked at the six important post hoc comparisons comparing only within each SCENE. T-test reveal that all differences are $p < .001$, see Figure 13.

Next, we investigated whether HAPTIC or SCENE significantly influenced if the **object felt real**. We again conducted a two-way ANOVA. The analysis revealed a significant effect for HAPTIC and SCENE ($F_{1,143} = 67.197, p < .001$; $F_{6,143} = 2.684, p < .017$; respectively). However, we found no significant interaction effect ($F_{6,143} = .757, p = .605$). Post hoc t-tests reviled that only *Environmental Haptics* vs. *Tangibles* is significantly different ($t(143) = -3.181, p = .038$), all others are ($p > .05$).

Finally, we were interested if the feeling of **no touch sensation** was significantly influenced by HAPTIC or SCENE. The analysis revealed a significant effect for HAPTIC and a significant interaction effect ($F_{1,143} = 248.679, p < .001$; $F_{6,143} = 2.231, p = .044$; respectively). However, no significant effect for SCENE ($F_{6,143} = 1.611, p = .148$). Due to the significant interaction effects, we again looked at the six meaningful post hoc comparisons, comparing only within each SCENE. The T-test reveals that all differences are $p < .001$, see Figure 13.

8.5 Presence

First, we analyzed the IPQ questionnaire [37] (scale: from 1 to 7) to determine whether HAPTIC significantly influenced the presence in VR, see Figure 14. As the sub-scale "Sense of Being There" was not normally distributed ($W = 0.908, p = 0.032$), we decided to run the Wilcoxon signed-rank test, which does not assume the normality of the data.

The tests revealed that the sub-scales "Sense of Being There," "Spacial Presence," and "Experienced Realism" are all significantly different between the condition *Haptic Feedback* vs. *No Feedback* ($V = 2, p = 0048$; $V = 1, p = .005$; $V = 0, p = .004$; respectively) where *Haptic Feedback* outperformed *No*

Feedback. However, there was no significant difference for the subscale "Involvement" ($V = 27$, $p = .635$). Finally, the overall feeling of presence scored by the IPQ was higher in the condition with haptic feedback ($V = 5$, $p = .005$).

9 DISCUSSION

We encompassed our efforts towards outperforming the current state-of-the-art in workspace size and render quality for mid-air ultrasound haptics for Virtual Reality. With the proposed system, we achieved a rendering volume almost ten times bigger than the highest rendering volume reported in the literature (Suzuki et al. [41]). The synergy between an ultrasound emitter and a serial robot increases the haptic array's workspace by optimizing the array rotation and location within the nominal actuation range of the robot. However, physically moving parts introduce an additional set of safety considerations to the design process. We tackled this by maintaining the robot's safety features (such as collision detection and maximum speed) and implementing an over-damped PD control that ensures soft movements of the haptic array. On the other hand, observing such safety measures reduces the system's versatility in speed or acceleration. It is technically possible to increase the robot's reaction speed to match the hand's movement speed (for example, with a full PID controller setup and overriding the speed constraints). However, additional safety measures should be introduced to protect the user from unwanted collisions. We executed a set of tests to model the end-effector behavior in terms of rotations and translations (see Figure 6 and 7). Additionally, the hand guidance strategy used to enable the robot to follow the hand favors configurations where the haptic array is facing outwards from the base of the robot; While this strategy can exploit the potential of the increased workspace, a more elaborated strategy could allow the end-effector to quickly relocate facing down to render sensations even if the palm of the user is facing up.

9.1 Challenges of extending ultrasound haptics workspace using moving arrays

Besides the technical challenges that can be faced when using a setup like the one presented in this manuscript, a primary design concern is the impact of motion on rendering quality. This is how the movement of the emitter alters the perceived quality of the ultrasound stimuli. We explored this impact using psychometric characterization as this is a central topic in the context of this manuscript. This puts our perceptual characterization as the first study reporting the impact of the emitter's movement on the perceived rendering quality of mid-air ultrasound haptics.

With this exploration, we identified that while a moving array closely matching the speed of the virtual moving object could lead to better perceptual features and, in the opposite case, a static object with a moving array could tend to perform less efficiently. We discovered that the overall performance of both systems is likely to be equivalent (Bayesian T-test). This means that even with the gain on rendering volume, both arrays can perform similarly; with an added gain, the rendering distance of the dynamic array to the hand is smaller, so the hand stays within the optimal rendering distance.

9.2 Contributions of an extended haptic rendering workspace to presence

Only one IPQ scale did not perform better than the baseline. However, this is not surprising as involvement with the scene and the environment's interaction did not change. Moreover, we found that volunteers enjoyed exploring the different showcases. We also found that we could significantly boost their enjoyment when providing haptic mid-air feedback. Combined with the technical evaluation and its result that the workspace is more than $19 m^3$ large, we believe that with the proposed setup, we make a step towards large-scale mid-air haptic feedback, which many systems can benefit from.

In the STRONG FEEDBACK questionnaire, a small number of participants rated the feedback as strong (above 1) in the NO FEEDBACK condition; the reason for this phenomenon could be the presence of phantom vibrations after being exposed to the vibrations produced by mid-air haptics. However, more evidence is necessary since the mid-air haptics community has not explored this phenomenon so far.

9.3 Extending Ultrasound Haptics Workspace, What is the best approach?

While several authors have reported a significant increase in the workspace of ultrasound-based mid-air haptics, these approaches have been highly heterogeneous, and yet ours significantly differs from those already reported in the literature:

Howard et al. [22] pan-tilt device tackles the problem by mounting the haptic array on a two-degree-of-freedom rotational platform. Such a setup considerably increases the workspace (around 17 times). However, the fundamental problem of the rendering quality remains untouched, given that the array is still spatially anchored to a fixed point. Instead, the proposed system drives the array closer to the user's hand, which impacts the rendering intensity and quality. A significant disadvantage of the proposed system is the control complexity and cost difference (Howard et al. [22] reported that their system cost only 150 Euro). However, the rendering volume and quality difference is considerable; the proposed system has a workspace 21 times larger than Howard et al. [22].

Brice et al. [6] approach instead was to relocate the haptic array in a set of pre-defined locations using a serial manipulator. Their approach increased the volume-rendered around 27 times with five pre-defined locations. While this setup can drive the array from one position to the other, it does not dynamically adapt to the user's hand position. Furthermore, it does not consider rotations of the haptic array and does not render while the array is relocating. Such an approach is inherently simpler from the control perspective and safer in many conditions, given that the number of situations where the robot is moved is considerably lower. Nevertheless, the implementation carries many of the static array limitations given that the approach is fundamentally the same but with added relocation. The proposed system instead continuously optimizes the position and rotation of the array to be close to the user's hand across a bigger workspace.

Finally, perhaps the most divergent approach in this ecosystem is the one proposed by Suzuki et al. [41]; Enlarging the ultrasound array by increasing the number of transducers. The approach by Suzuki et al. [41] advantage is the lack of moving parts in the setup, making it inherently safe concerning collisions within the interaction space. Additionally, the focal point relocation is potentially faster in extreme positions, given that the emitters do not have to switch positions or locations. Suzuki et al. [41] proposed setup enlarged the rendering volume around 36 times compared with the original Ultraleap workspace. However, the setup is technically more complex and energy inefficient, considering the overheating of the haptic arrays. Furthermore, the rendered volume is still considerably smaller than the one achieved with the proposed setup.

A major advantage of the proposed setup over the cited alternatives is the system's scalability. Cobots are increasingly becoming more popular and affordable. They can be used in a broad number of contexts, including haptics. This favors the proposed modular design that allows us to mount and unmount the haptic array from the robot's end-effector to switch from one proposed to the other easily. Moreover, a potential extension of the current work could be the integration of encountered-type haptics and extended ultrasound mid-air haptics to render kinaesthetic and tactile-only sensations with a single setup. Such an advantage is unfortunately not possible with the compared alternatives given the highlighted limitations on their approaches.

Yet, the proposed approach also carries a series of disadvantages and challenges, which we will discuss in the following.

9.4 Limitations and Challenges

Although the proposed setup enables considerably larger mid-air haptics rendering volume and quality, we discovered a set of limitations. Perhaps the most important one is the speed of movement. Although cobots are safe to interact with humans, the end-effector's speed of movement should be limited to avoid potential harm to the users; this constrains the reaction time of the array to be positioned from one point to the other.

Also, the proposed system has its optimal actuation dynamics and rendering when the user's palms are facing the robot's base and moving perpendicularly to it. Although it can drive the array upside down, such movements are more complex and take the robot to the boundaries of its workspace more quickly.

9.4.1 Hand Tracking. We used LeapMotion as the main tracking system and the VIVE tracker as an auxiliary tracking fallback. Even though LeapMotion can track fast-moving hands, the combination of moving hands and moving array seemed to decrease the hand recognition stability. This was especially noticeable in scenarios that required moving objects, such as falling leaves or rain. The latter posed an exciting challenge in tracking since the rain is perceived as falling from the sky. Therefore, the array had to face downwards. Thus, the low quality of the LeapMotion's tracking restricted the possible movement space of the robot. In this regard, we recommend exploring new alternatives for detailed hand movements; perhaps a multi-camera AI-based system could contribute to attenuating this issue.

9.4.2 Safety. A moving array in a shared environment with a human requires maximum safety considerations. The moving elements can hit the user or other individuals in the room. For this reason, redundant safety is required. Measures like maximum speed limits, collision detection, and the robot's visual feedback inside the VR environment are not optional. One possible workaround could be including the robot's feedback as a dietetic interface. In addition to the mechanical risk, rendering ultrasound with a moving array can misguide the hand's focal points to the user's ears. For this reason, ear protection is required during the whole VR experience.

9.4.3 Noise. Both the robot and the haptic display are loud when rendering and moving in space. Even with ear protection, volunteers reported hearing the robot actuating motors and the sound of the focused ultrasound being rendered, especially when moving the hand from one extreme position to the opposite side. This should be considered when designing experiences using the proposed approach, since such noises can break immersion if they are not correctly blended into the experience.

10 CONCLUSIONS

In this paper, we present a system to extend the workspace of mid-air ultrasound haptics to more than $18.98m^3$. We evaluated the proposed system technically, perceptually, and qualitatively. Our technical tests modeled the system's response in translation and rotation to guarantee repeatability and enough speed to follow the hand while preserving the safety of the users. We also studied the implications of rendering ultrasound haptics while moving the emitter and how this impacts the rendering quality. While we found a tendency for the dynamic array to perform better in scenarios with moving objects, the results showed that the perceptual features of both configurations are similar. Such validation indicates that even putting the haptic array under strain (rendering in motion), rendering quality is not worsened, making it suitable for the proposed setup. Furthermore, we found that the system enables users to enjoy mid-air haptics in a larger space than before. We showcased this in seven different environments that featured different haptic explorations. These scenarios illustrate the system's capabilities and depict the challenges that such a paradigm

brings. We presented the first study on the impact of motion of the ultrasound emitter on rendering quality on a perceptual level using psychometric testing. Drawing on that evidence, we exposed the advantages and challenges of ultrasound arrays in motion, compared the proposed approach with the state-of-the-art alternatives, and discussed the advantages and disadvantages of the alternative approaches. Even though the proposed system has some limitations, the considerable gain in workspace, quality, and scalability of the setup in terms of extensibility to encountered type haptics makes us believe that it is worth considering as a reasonable option to deliver haptics for VR.

ACKNOWLEDGMENTS

This work was partly funded by the European Research Council (ERC AMPLIFY, no. 683008), the Federal Ministry of Education and Research of Germany (BMBF Hive, no. 16SV8183) and One Munich Strategy Forum at LMU Munich, funded under the Hightech Agenda Bavaria.

REFERENCES

- [1] Ahmad Abiri, Yen-Yi Juo, Anna Tao, Syed J Askari, Jake Pensa, James W Bisley, Erik P Dutson, and Warren S Grundfest. 2019. Artificial palpation in robotic surgery using haptic feedback. *Surgical Endoscopy* 33, 4 (2019), 1252–1259. <https://doi.org/10.1007/s00464-018-6405-8>
- [2] Nathaniel Agharese, Tyler Cloyd, Laura H Blumenschein, Michael Raitor, Elliot W Hawkes, Heather Culbertson, and Allison M Okamura. 2018. HapWRAP: Soft growing wearable haptic device. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, New York, NY, USA, 5466–5472. <https://doi.org/10.1109/ICRA.2018.8460891>
- [3] Kentaro Ariga, Masahiro Fujiwara, Yasutoshi Makino, and Hiroyuki Shinoda. 2021. Workspace Evaluation of Long-Distance Midair Haptic Display Using Curved Reflector. In *2021 IEEE World Haptics Conference (WHC)*. IEEE, New York, NY, USA, 85–90. <https://doi.org/10.1109/WHC49131.2021.9517193>
- [4] Jonas Auda, Nils Verheyen, Sven Mayer, and Stefan Schneegass. 2021. Flyables: Haptic Input Devices for Virtual Reality Using Quadcopters. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology (Osaka, Japan) (VRST '21)*. Association for Computing Machinery, New York, NY, USA, Article 40, 11 pages. <https://doi.org/10.1145/3489849.3489855>
- [5] Héctor Barreiro, Stephen Sinclair, and Miguel A Otaduy. 2019. Ultrasound rendering of tactile interaction with fluids. In *2019 IEEE World Haptics Conference (WHC '19)*. IEEE, New York, NY, USA, 521–526. <https://doi.org/10.1109/WHC.2019.8816137>
- [6] Daniel Brice, Thomas McRoberts, and Karen Rafferty. 2019. A proof of concept integrated multi-systems approach for large scale tactile feedback in VR. In *International Conference on Augmented Reality, Virtual Reality and Computer Graphics*. Springer International Publishing, Cham, 120–137. https://doi.org/10.1007/978-3-030-25965-5_10
- [7] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. 2013. UltraHaptics: Multi-Point Mid-Air Haptic Feedback for Touch Surfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (St. Andrews, Scotland, United Kingdom) (UIST '13)*. Association for Computing Machinery, New York, NY, USA, 505–514. <https://doi.org/10.1145/2501988.2502018>
- [8] Xiaojun Chen and Junlei Hu. 2018. A review of haptic simulator for oral and maxillofacial surgery based on virtual reality. *Expert Review of Medical Devices* 15, 6 (2018), 435–444. <https://doi.org/10.1080/17434440.2018.1484727>
- [9] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. 2018. CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174228>
- [10] Giulia Corniani and Hannes P. Saal. 2020. Tactile innervation densities across the whole body. *Journal of Neurophysiology* 124, 4 (2020), 1229–1240. <https://doi.org/10.1152/jn.00313.2020>
- [11] Richard C. Dorf and Robert H. Bishop. 2011. *Modern control systems*. Pearson, London, England, UK.
- [12] Cathy Fang, Yang Zhang, Matthew Dworman, and Chris Harrison. 2020. Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3313831.3376470>
- [13] Euan Freeman, Ross Anderson, Julie Williamson, Graham Wilson, and Stephen A. Brewster. 2017. Textured Surfaces for Ultrasound Haptic Displays. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction (Glasgow, UK) (ICMI '17)*. Association for Computing Machinery, New York, NY, USA, 491–492. <https://doi.org/10.1145/3136755.3143020>

- [14] William Frier, Damien Ablart, Jamie Chilles, Benjamin Long, Marcello Giordano, Marianna Obrist, and Sriram Subramanian. 2018. Using Spatiotemporal Modulation to Draw Tactile Patterns in Mid-Air. In *Haptics: Science, Technology, and Applications*. Springer International Publishing, Cham, 270–281. https://doi.org/10.1007/978-3-319-93445-7_24
- [15] William Frier, Dario Pittera, Damien Ablart, Marianna Obrist, and Sriram Subramanian. 2019. Sampling Strategy for Ultrasonic Mid-Air Haptics. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI '19*). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300351>
- [16] Alex Girdler and Orestis Georgiou. 2020. Mid-Air Haptics in Aviation—creating the sensation of touch where there is nothing but thin air. *arXiv preprint arXiv:2001.01445* (2020). <https://doi.org/10.48550/ARXIV.2001.01445>
- [17] Blake Hannaford and Allison M Okamura. 2016. *Haptics*. Springer International Publishing, Cham, 1063–1084. https://doi.org/10.1007/978-3-540-30301-5_31
- [18] Kyle Harrington, David R. Large, Gary Burnett, and Orestis Georgiou. 2018. 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* (Toronto, ON, Canada) (*AutomotiveUI '18*). Association for Computing Machinery, New York, NY, USA, 11–20. <https://doi.org/10.1145/3239060.3239089>
- [19] Seongkook Heo, Christina Chung, Geehyuk Lee, and Daniel Wigdor. 2018. Thor's Hammer: An Ungrounded Force Feedback Device Utilizing Propeller-Induced Propulsive Force. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3173574.3174099>
- [20] Matthias Hoppe, Pascal Knierim, Thomas Kosch, Markus Funk, Lauren Futami, Stefan Schneeggass, Niels Henze, Albrecht Schmidt, and Tonja Machulla. 2018. VRHapticDrones: Providing Haptics in Virtual Reality through Quadcopters. In *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia* (Cairo, Egypt) (*MUM 2018*). Association for Computing Machinery, New York, NY, USA, 7–18. <https://doi.org/10.1145/3282894.3282898>
- [21] Matthias Hoppe, Daria Oskina, Albrecht Schmidt, and Thomas Kosch. 2021. Odin's Helmet: A Head-Worn Haptic Feedback Device to Simulate G-Forces on the Human Body in Virtual Reality. *Proc. ACM Hum.-Comput. Interact.* 5, EICS, Article 212 (may 2021), 15 pages. <https://doi.org/10.1145/3461734>
- [22] Thomas Howard, Maud Marchal, Anatole Lécuyer, and Claudio Pacchierotti. 2019. PUMAH: Pan-tilt Ultrasound Mid-Air Haptics for larger interaction workspace in virtual reality. *IEEE transactions on haptics* 13, 1 (2019), 38–44. <https://doi.org/10.1109/TOH.2019.2963028>
- [23] HTC. 2018. HTC Vive Pro. <https://www.vive.com/eu/product/vive-pro/> Accessed: 2022-08-16.
- [24] Inwook Hwang, Hyungki Son, and Jin Ryong Kim. 2017. AirPiano: Enhancing music playing experience in virtual reality with mid-air haptic feedback. In *2017 IEEE World Haptics Conference (WHC '17)*. IEEE, New York, NY, USA, 213–218. <https://doi.org/10.1109/WHC.2017.7989903>
- [25] Roland S Johansson and Ake B Vallbo. 1979. Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *The Journal of physiology* 286, 1 (1979), 283–300. <https://doi.org/10.1113/jphysiol.1979.sp012619>
- [26] Pascal Knierim, Valentin Schwind, Anna Maria Feit, Florian Nieuwenhuizen, and Niels Henze. 2018. Physical Keyboards in Virtual Reality: Analysis of Typing Performance and Effects of Avatar Hands. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3173574.3173919>
- [27] Robert Kovacs, Eyal Ofek, Mar Gonzalez Franco, Alexa Fay Siu, Sebastian Marwecki, Christian Holz, and Mike Sinclair. 2020. Haptic PIVOT: On-Demand Handhelds in VR. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (*UIST '20*). Association for Computing Machinery, New York, NY, USA, 1046–1059. <https://doi.org/10.1145/3379337.3415854>
- [28] Sascha Kraus, Dominik K. Kanbach, Peter M. Krysta, Maurice M. Steinhoff, and Nino Tomini. 2022. Facebook and the creation of the metaverse: radical business model innovation or incremental transformation? *International Journal of Entrepreneurial Behavior & Research* 28, 9 (01 Jan 2022), 52–77. <https://doi.org/10.1108/IJEBR-12-2021-0984>
- [29] Benjamin Long, Sue Ann Seah, Tom Carter, and Sriram Subramanian. 2014. Rendering Volumetric Haptic Shapes in Mid-Air Using Ultrasound. *ACM Trans. Graph.* 33, 6, Article 181 (Nov. 2014), 10 pages. <https://doi.org/10.1145/2661229.2661257>
- [30] Maud Marchal, Gerard Gallagher, Anatole Lécuyer, and Claudio Pacchierotti. 2020. Can Stiffness Sensations be Rendered in Virtual Reality Using Mid-air Ultrasound Haptic Technologies?. In *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer International Publishing, Cham, 297–306. https://doi.org/10.1007/978-3-030-58147-3_33
- [31] Jonatan Martinez, Adam Harwood, Hannah Limerick, Rory Clark, and Orestis Georgiou. 2019. Mid-Air Haptic Algorithms for Rendering 3D Shapes. In *2019 IEEE International Symposium on Haptic, Audio and Visual Environments*

- and Games (HAVE). IEEE, New York, NY, USA, 1–6. <https://doi.org/10.1109/HAVE.2019.8921211>
- [32] Meta. 2021. Oculus Quest 2. <https://www.oculus.com/quest-2/> Accessed: 2022-08-16.
- [33] Kouta Minamizawa, Souichiro Fukamachi, Hiroyuki Kajimoto, Naoki Kawakami, and Susumu Tachi. 2007. Gravity Grabber: Wearable Haptic Display to Present Virtual Mass Sensation. In *ACM SIGGRAPH 2007 Emerging Technologies*. Association for Computing Machinery, New York, NY, USA, 8–es. <https://doi.org/10.1145/1278280.1278289>
- [34] Marianna Obrist, Sriram Subramanian, Elia Gatti, Benjamin Long, and Thomas Carter. 2015. Emotions Mediated Through Mid-Air Haptics. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 2053–2062. <https://doi.org/10.1145/2702123.2702361>
- [35] Ryan M Peters, Monica D McKeown, Mark G Carpenter, and J Timothy Inglis. 2016. Losing touch: age-related changes in plantar skin sensitivity, lower limb cutaneous reflex strength, and postural stability in older adults. *Journal of Neurophysiology* 116, 4 (2016), 1848–1858. <https://doi.org/10.1152/jn.00339.2016>
- [36] Ismo Rakkolainen, Euan Freeman, Antti Sand, Roope Raisamo, and Stephen Brewster. 2021. A Survey of Mid-Air Ultrasound Haptics and Its Applications. *IEEE Transactions on Haptics* 14, 1 (2021), 2–19. <https://doi.org/10.1109/TOH.2020.3018754>
- [37] Holger T. Regenbrecht, Thomas W. Schubert, and Frank Friedmann. 1998. Measuring the Sense of Presence and its Relations to Fear of Heights in Virtual Environments. *International Journal of Human-Computer Interaction* 10, 3 (1998), 233–249. https://doi.org/10.1207/s15327590ijhc1003_2
- [38] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. 2017. Construction of a Benchmark for the User Experience Questionnaire (UEQ). *International Journal of Interactive Multimedia and Artificial Intelligence* 4, 4 (2017), 40–44. <https://doi.org/10.9781/ijimai.2017.445>
- [39] Heiko H Schütt, Stefan Harmeling, Jakob H Macke, and Felix A Wichmann. 2016. Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision research* 122 (2016), 105–123. <https://doi.org/10.1016/j.visres.2016.02.002>
- [40] Mike Sinclair, Eyal Ofek, Mar Gonzalez-Franco, and Christian Holz. 2019. CapstanCrunch: A Haptic VR Controller with User-Supplied Force Feedback. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 815–829. <https://doi.org/10.1145/3332165.3347891>
- [41] Shun Suzuki, Masahiro Fujiwara, Yasutoshi Makino, and Hiroyuki Shinoda. 2019. Midair Ultrasound Haptic Display with Large Workspace. In *Haptic Interaction*. Springer Singapore, Singapore, 3–5. https://doi.org/10.1007/978-981-13-3194-7_1
- [42] Natalya Tatarchuk. 2006. Artist-Directable Real-Time Rain Rendering in City Environments. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) (SIGGRAPH '06). Association for Computing Machinery, New York, NY, USA, 23–64. <https://doi.org/10.1145/1185657.1185828>
- [43] UltraLeap. 2022. Ultrahaptics Developer Information: Interaction zones and form factors. <https://developer.ultrahaptics.com/knowledgebase/interaction-zones-form-factors/> Accessed: 2022-08-16.
- [44] Valve. 2019. Valve Index. <https://store.steampowered.com/valveindex> Accessed: 2022-08-16.
- [45] Steeven Villa and Sven Mayer. 2022. Cobity: A Plug-And-Play Toolbox to Deliver Haptics in Virtual Reality. In *Proceedings of Mensch Und Computer 2022* (Darmstadt, Germany) (MuC '22). Association for Computing Machinery, New York, NY, USA, 78–84. <https://doi.org/10.1145/3543758.3543775>
- [46] Chyanna Wee, Kian Meng Yap, and Woan Ning Lim. 2021. Haptic Interfaces for Virtual Reality: Challenges and Research Directions. *IEEE Access* 9 (2021), 112145–112162. <https://doi.org/10.1109/ACCESS.2021.3103598>
- [47] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 143–146. <https://doi.org/10.1145/1978942.1978963>
- [48] Andre Zenner and Antonio Krüger. 2017. Shifty: A weight-shifting dynamic passive haptic proxy to enhance object perception in virtual reality. *IEEE transactions on visualization and computer graphics* 23, 4 (2017), 1285–1294. <https://doi.org/10.1109/TVCG.2017.2656978>
- [49] Yang Zhang, Wolf Kienzle, Yanjun Ma, Shiu S. Ng, Hrvoje Benko, and Chris Harrison. 2019. ActiTouch: Robust Touch Detection for On-Skin AR/VR Interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 1151–1159. <https://doi.org/10.1145/3332165.3347869>