



Ultrasound Mid-Air Tactile Feedback for Immersive Virtual Reality Interaction

Thomas Howard, Claudio Pacchierotti, Maud Marchal

► To cite this version:

Thomas Howard, Claudio Pacchierotti, Maud Marchal. Ultrasound Mid-Air Tactile Feedback for Immersive Virtual Reality Interaction. Ultrasound Mid-Air Haptics for Touchless Interfaces, Springer International Publishing, pp.147-183, 2022, Human–Computer Interaction Series, 10.1007/978-3-031-04043-6_6 . hal-03797567

HAL Id: hal-03797567

<https://inria.hal.science/hal-03797567>

Submitted on 4 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Ultrasound Mid-Air Tactile Feedback for Immersive Virtual Reality Interaction ^{*}

Thomas Howard¹, Maud Marchal², and Claudio Pacchierotti¹

¹ CNRS, Univ Rennes, Inria, IRISA – France

² Univ Rennes, INSA, IRISA, Inria, CNRS – France and IUF – France
`{name.surname}@irisa.fr`

Abstract. Ultrasound mid-air haptic (UMH) devices are promising for tactile feedback in virtual reality (VR), as they do not require users to be tethered to, hold, or wear any device. This approach is less cumbersome, easy to set up, can simplify tracking, and leaves the hands free for concurrent interactions. This chapter explores work conducted at CNRS-IRISA dealing with the challenges arising from the integration of UMH interfaces in immersive VR through three main axes. These are discussed in the wider context of the state-of-the-art on UMH for augmented and virtual reality, and illustrated through several VR use-cases.

A first axis deals with device integration into the VR ecosystem. Interaction in immersive VR is based on the synergy between complex input devices allowing real-time tracking of the user and multimodal feedback devices delivering a coherent visual, auditory and haptic picture of a simulated virtual environment (VE). Using UMH in immersive VR therefore hinges on integrating UMH devices such that their operation does not interfere with other input and feedback devices. It is also critical to ensure that UMH feedback is adequately synchronised and co-located with respect to other stimuli, and delivered within a workspace that is compatible with that of VR interaction. Regarding this final point, we propose PUMAH, a robotic solution for increasing the usable workspace of UMH devices.

The second and third axes, respectively, focus on stimulus perception and rendering of VE properties. Virtual object properties can be rendered in a variety of ways, through e.g. amplitude modulation (AM) or spatio-temporal modulation (STM), with many parameters (modulation frequency, spatial sampling, etc.) coming into play, raising questions about the limits of the design space.

To tackle this challenge, we begin by conducting psychophysical experimentation to understand the usable ranges for stimulus parameters and understand the perceptual implications of stimulus design choices. We propose an open-source software framework intended to facilitate UMH stimulus design and perceptual evaluation. These results in turn serve as the basis for the design and evaluation of rendering schemes for VR. Using amplitude variations along a focal point path in STM, we investigate the possibility of rendering geometric details and in a second step, sensations of stiffness in VR.

^{*} This research has received funding from the EU’s H2020 research and innovation programme (grant agreement No 801413, project “H-Reality”).

1 Introduction

The preceding chapters offered a broad view of ultrasound mid-air haptics (UMH), highlighting challenges and perspectives in design and interaction across multiple application domains. In contrast, the present chapter focuses specifically on applications to immersive virtual reality (VR). As such, it serves as somewhat of a bridge between the broader application-independent considerations tackled previously and lower-level UMH rendering and device integration considerations which are dealt with in depth in subsequent chapters. Our work echoes previous considerations relating to UMH’s role in supporting multimodal interaction, user experience and agency in interactions, in the context of using UMH to support VR interaction where inputs to different sensory modalities are fully independent from one another.

The field of VR has developed dramatically over the past decade, thanks in large part to the wide availability of cheap, high-quality head-mounted displays (HMDs) and high-performance graphics processors.

Echoing the early evolution of human-computer interaction (HCI), haptics has been absent or at least slow to be integrated into the VR interaction loop. As was previously noted in relation to touchless user interfaces (Chapter “Augmenting Automotive Gesture Infotainment Interfaces Through Mid-Air Haptic Icon Design”), interaction with digital content inherently lacks natural touch, and VR is no exception to this. Yet, in VR interactions dominated by vision and audio, the addition of the sense of touch has been shown to enhance interaction (e.g. [54]), increase realism (e.g. [55,80]) and improve immersion (e.g. [67,70]).

Immersive VR relies on the use of a HMD worn by the user, blocking out vision and hearing of the real environment while rendering a full 3D visual scene with immersive audio and, depending on the application, some forms of haptic feedback. In its current form, immersive VR however does not occlude the user’s real world perception of smell, taste and touch, both through user’s proprioception and contact with the real environment.

Interaction is considered as the reciprocal action of the user on the computer system implementing a VR environment through various input devices, and of said computer system on the user through various feedback devices (see Fig. 1). Haptic feedback enhances interaction by enriching the action of the system on the user, providing additional immediate and informative feedback, in particular during contact with and manipulation of virtual objects (green frame in Fig. 1). We define realism as the quality of the representation of a virtual environment (VE) to the user in a way that is accurate and true to the VE’s internal structure and rules. Haptic feedback can increase realism by providing sensations corresponding to the virtual physical processes occurring in the virtual environment (purple frame in Fig. 1). Immersion can be understood as the result of a human-computer interface’s capability to remove or mask as many real world sensations as possible, and to substitute them with sensations corresponding to the VE [57]. Haptic feedback thus improves immersion by stimulating an otherwise unused sensory modality, which would tend to anchor the user in the perception of the real environment rather than that of the VE (red frame in Fig. 1). In addition,

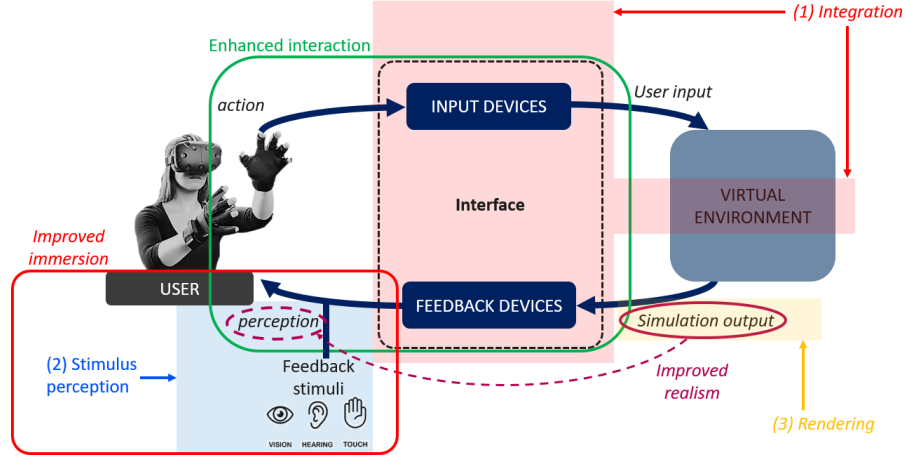


Fig. 1. The HCI interaction loop as applied to immersive VR. Multiple input devices acquire the user’s state and actions, which can be mapped to the user’s avatar, i.e. their virtual representation in the simulated virtual environment (VE), or to other VE properties. The state of said VE and its actions on the avatar in turn generate outputs which are fed to the multiple feedback devices interfacing with the user. In immersive VR, these feedback devices have the specific characteristic of blocking out real-world sensory inputs for vision and hearing thanks to the HMD, replacing them with sensory inputs related to the VE. Regarding touch, the real-world sensory inputs from the user’s proprioception and contact with their physical environment are not masked, but rather overlaid with virtual sensory inputs from various haptic feedback devices. The broad areas where haptic feedback can provide benefits are circled in green, purple and red respectively for enhanced interaction, improved realism and improved user immersion. The shaded boxes highlight the three axes discussed in this Chapter: device integration (red), stimulus perception (blue), and rendering (yellow).

or as a byproduct of the former, added haptics in interaction with virtual environments appears to have positive effects on user presence [46], engagement and affective response [73], as well as factors relating to embodiment – in particular sense of agency [13].

A wide variety of haptic technologies have been applied to achieve these improvements, for the most part relying on direct physical contact between the user and an actuator. These encompass grounded force-feedback (e.g. [11]), exoskeletons (e.g. [48]), wearables (e.g. [76]), tangible objects (e.g. [81]) and encounter-type haptic devices (e.g. [56]), to name a few. In contrast, contactless haptics aim to provide the benefits of haptic feedback discussed above in contexts where direct physical contact between the user and the interface is not desired [1,10]. Ultrasound mid-air haptics (UMH) [28,8] provides an interesting recent addition to this arsenal of technologies, combining the general benefits from contactless haptics with benefits specific to the stimuli delivered by ultrasound interfaces.

UMH and its many applications have been reviewed in detail by Rakkolainen et al. [69].

There are many potential benefits to using contactless haptics when interacting with virtual environments. Removing unwanted physical contact with a controller or actuator may improve interaction by freeing up the hand from constraining devices. This in turn can facilitate better manipulation of virtual objects and ease tracking. Freeing up the hand also allows concurrent manipulation, for example of tangible objects, without impeding real world touch feedback when it is desired. Physical contact may also be undesirable for the sake of improved immersion. Contactless interfaces achieve the best possible interface transparency by preventing any haptic stimuli arising from the contact of the actuator on the user's skin [69]. This improved immersion also translates to an advantage for VE designers as contactless devices do not constrain VR interaction design to match the unavoidable perceived actuator properties (e.g. forcing the user's avatar to wear gloves or a wristband to match haptic devices worn in the real world and thus maintain immersion). A final potential advantage from the absence of contact comes from the perspective of user safety and hygiene. As no physical contact between the user and actuator takes place, there is a reduced risk of user injury from the device or contact with contaminated surfaces. While the latter point may not be as prominent of an issue as for e.g. digital signage [12], it may still be relevant e.g. in medical AR and VR applications. It should be noted that some of these benefits are not specific to immersive virtual environments and are also discussed in chapters "Augmenting Automotive Gesture Infotainment Interfaces Through Mid-Air Haptic Icon Design" and "Modulation Methods for Ultrasound Midair Haptics" through the lens of touchless user interfaces for interacting with digital content.

In addition to the general benefits of contactless haptic interfaces over contact haptics, UMH devices provide many specific opportunities for VR applications. Ultrasound devices are capable of delivering a variety of sensations ranging from light pressure and vibrotactile stimuli [8] to sensations of airflow [26] and thermal stimuli [39], all within a continuous 3D workspace and with a very high temporal resolution. These high-dimensional stimuli are potentially capable of encoding a lot of information through rapidly and finely re-configurable spatial, temporal and intensity characteristics [61]. The device's workspace and lack of physical constraints on the stimuli means that stimuli do not necessarily need to be restricted to a single point on the skin or even just the hand, but can potentially target any exposed skin on the user's body [24,79]. UMH interfaces are highly interesting in the multi-modal context of augmented and virtual reality interaction as they allow the generation of a wide variety not only of tactile sensations, but also auditory stimuli as well as the levitation of particles. This provides potential for multi-sensory stimulation based on a single device [64,78]. The manner in which UMH can support multimodal interaction are also discussed in detail in chapters "Opportunities for Multisensory Mid-Air Interactions Featuring Ultrasound Haptic Feedback" and "Multimodal Interaction with Mid-Air Haptics". Finally, UMH device have the potential to be embedded within walls, tables

[42], other furniture and objects [82] and even the headset [77] which can be advantageous in terms of footprint, leaving the VR workspace unobstructed.

The present Chapter is intended to provide an overview of the current state of research conducted at CNRS on the subject of integrating ultrasound mid-air haptics to immersive virtual reality. This work is discussed in the context of the broader state-of-the-art for UMH in the hopes of also serving as an introduction to UMH for researchers and designers hoping to apply this technology to immersive interaction with virtual environments.

2 State-of-the-art and challenges for integrating UMH in immersive VR

Some of the earliest proposed use-cases for UMH feedback concerned augmented reality (AR) interactions which share many common aspects and challenges with immersive VR. As early as 2009, Hoshi et al. [28] presented a proof-of-concept integration of multiple focused ultrasound arrays, a hand tracking system and a holographic visual display to add a tactile dimension to holograms. This was later expanded upon in work by Inoue et al. [35] and Kimura et al. [43]. Augmented reality screens with tactile feedback also feature prominently in early use-cases for UMH [86,58,72]. More recently, this concept of touchable holograms has been refined to leverage improved AR HMDs to deliver more complex visual feedback and interface with more diverse external inputs, as seen e.g. in demonstrators by Kervegant et al. [42] or the touchable bio-holograms from Romanus et al. [71]. Finally, a recent development has seen a shift from passive displays to interfaces allowing manipulation with UMH feedback [53,49,43,85]. Please refer to chapter “Touchless Tactile Interaction with Unconventional Permeable Displays” for details on augmented-reality user interfaces as well as chapter “Superimposing Visual Images on Mid-Air Ultrasound Haptic Stimulation” for a detailed discussion of the challenge of congruence between vision and haptics in such systems.

As consumer VR headsets became more widely available, several proof-of-concept interaction use-cases for UMH were developed specifically in VR. Hwang et al. [34] presented the Airpiano, a VR piano playing simulation where key presses were rendered through amplitude-modulated (AM) focal points projected on the user’s fingertips. UMH has also been applied to enhancing VR gaming experiences by providing abstract cues supporting the interaction [23] or sensations of supernatural phenomena [51]. In the domain of medical VR applications, Karafotias et al. [41] explored the potential for using UMH feedback to support applications for pain distraction, while Balint et al. [3] present a proof-of-concept reconfigurable medical palpation simulation based on UMH feedback.

Fig. 1 highlights the categories of challenges relating to the integration of UMH feedback to VR experiences. To synthesise insights from our work on this topic, we can group the challenges we tackled into three categories: *device integration* (red in Fig. 1), *stimulus perception* (blue in Fig. 1) and *rendering* (yellow in Fig. 1).

The challenge of device integration (see Sec. 3) relates to incorporating UMH devices into VR interactions in a way that safely enables full use of the UMH feedback device to support the VR interaction without degrading the interfaces between any other input or feedback devices and the user. Our present work does not explicitly tackle the subject of safety, however details on this can be found in chapters "Safety of High-Intensity Ultrasound" and "Ultrasound Exposure in Mid-Air Haptics".

Along with properly integrating UMH devices into VR interaction, it is necessary to understand how users perceive the generated stimuli (see Sec. 4). This is critical to guide tactile rendering design, overcome inherent limitations to UMH devices, and leverage cross-modal interactions which inevitably occur in the multimodal context of immersive VR.

Finally, these insights into the relationship between stimulus design and perceived haptic properties allow tackling of the challenges of rendering for VR interaction (see Sec. 5). These are two-fold: on the one hand, we seek to understand what properties of virtual objects and contacts can be effectively rendered using UMH, and on the other, we seek to understand how to best render these properties given the UMH stimulus design space and constraints. This discussion on rendering will build on technical aspects of UMH rendering, a large part of which are subsequently discussed in detail in chapters "Introduction to Ultrasonic Mid-Air Haptic Effects", "Modulation Methods for Ultrasound Midair Haptics", "Sound-Field Creation for Haptic Reproduction" and "The Physical Principles of Arrays for Mid-Air Haptic Applications".

While our work has not tackled challenges beyond these three groups, further aspects to consider with respect to integrating UMH feedback to VR interaction have been studied by others, such as the need to understand users' affective responses to UMH stimuli [16,63] and the impact of UMH feedback on sense of agency [13,17] (see also chapter "Opportunities for Multisensory Mid-Air Interactions Featuring Ultrasound Haptic Feedback"). The following three sections discuss each of the listed challenge groups, presenting topic-specific state-of-the-art as well as a survey of the work conducted at CNRS on each of these topics.

3 Integrating UMH devices into the VR ecosystem

As shown in Fig. 1, interaction in immersive VR relies on a set of input devices tracking the users position, state and receiving explicit input actions. These inputs feed into a computer running the simulation of the VE in interactive time (this usually means latency is kept below 75ms [83] and frame rates above 60fps). This simulation in turn generates outputs for the video, audio and haptic feedback devices.

Proper UMH device integration into the interaction loop requires:

- (1) Getting the UMH device to not impede user input and ensuring the input devices provide adequate information for the simulation to generate relevant outputs to the UMH device (see Sec. 3.1).

- (2) Ensuring the UMH device does not break immersion through its physical footprint or the delivery of unwanted stimuli. Ensuring co-location and synchronization between the UMH rendering of a virtual scene with other existing haptic, audio and visual renderings of the same scene, without degrading the interface between any of the other feedback devices and the user.
- (3) Ensuring the UMH devices' rendering workspace adequately covers the VR interaction workspace (see Sec. 3.3).

These challenges are addressed through a combination of hard- and software design as well as adequate interaction design.

3.1 Integrating with input devices

Since the real environment and user's body are hidden from the user in immersive VR, it usually does not pose many constraints on the obtrusiveness of tracking hardware. However, despite it having been shown that UMH feedback can in some cases be delivered through fabric [79], UMH generally requires the user's skin to be unobstructed in order for the stimuli to be effectively delivered. This poses certain constraints on the tracking hardware with regards to obtrusiveness, both of the tracking devices themselves and of their fixation on the user's skin.

Furthermore, the high freedom of spatial reconfigurability of UMH stimuli means that the full potential of the technology is best achievable only if the user's pose is tracked with a high level of detail, i.e. the tracking system provides a high number of degrees of freedom (DoF). This is especially challenging when it comes to tracking the hands, which usually are the focus of interaction with UMH devices, and present a high number of DoFs and relatively small anatomical features. This challenge is more general to immersive VR and is usually tackled through sensorized in-hand controllers (e.g. Valve Index Controllers³), tracking gloves (e.g. ManusVR⁴) or hand exoskeletons (e.g. SenseGlove⁵) which provide accurate high-DoF hand tracking over large VR workspaces. These solutions however are mostly incompatible with UMH as they tend to completely occlude the skin in the interaction region.

Another point to consider is that contrary to e.g. wearable devices, UMH stimuli are generated in the VE world frame of reference. Thus, a relatively high tracking accuracy is required to ensure UMH stimuli are actually delivered to the target area on the user's skin (corresponding to the area of virtual contact). This requirement can be somewhat moderated by taking into account the relatively low spatial precision of UMH stimuli. Focal points are perceptible up to a couple of mm laterally in planes parallel to the array's x-y plane and up to a couple of cm vertically around the nominal focal point position. On the one hand, this property can be an issue for rendering (which we discuss in more detail in Sec. 5.1) but it can also be seen as an advantage with respect to requirements

³ <https://www.valvesoftware.com/en/index/controllers>

⁴ <https://www.manus-vr.com/vr-products>

⁵ <https://www.senseglove.com/product/developers-kit/>

posed for tracking, as there is no reasonable argument for demanding sub-mm accuracy in tracking if the rendering’s accuracy cannot follow.

To satisfy requirements towards integration with tracking hardware, most literature solutions in AR and VR interaction rely on optical tracking systems, as these share a similar requirement for unobstructed line-of-sight to the user’s skin as UMH devices. These range from custom-built (e.g. [28,58]) to commercial (e.g. Leap Motion⁶ used in [35,53,71]) tracking systems, and in the case of AR interaction often leverage inside-out tracking capabilities provided by AR headsets (e.g. the Hololens⁷ used in [42] or Magic Leap⁸ used in [71]).

The majority of work conducted at CNRS presented in this Chapter makes use of the HTC Vive Pro system⁹ which uses outside-in optical tracking. In it, multiple IR lasers projected from base stations placed around the edges of the VR interaction zone scan the workspace and are picked up by active sensors within the headset, controllers and trackers, allowing triangulation of the user’s body parts and tracked objects. This type of tracker is rather bulky and usually used with straps and is thus not ideal for UMH interaction as they require adequate placement of the straps to keep interaction zones on the skin (e.g. the palm and fingertips) unobstructed. However there are possibilities for creating adhesive fixtures on the hairy skin, usually not a target for UMH, such as the one we developed in the work by de Tinguy et al. [80]. Most of our use-cases focused on single hand interactions [50,33,32], sometimes with additional input through a handheld controller in the hand not interacting with UMH. The use of a tracker allows input over a workspace of the size of the VR interaction zone, however does not provide fine pose information for the hand.

Overall, because of the large interaction zones, 3D workspace and freedom of user movement, accurate tracking - in particular of the hands - remains an open challenge for haptics in VR interaction in general and even more so in UMH feedback for VR because of the added constraint of keeping the skin free and unobstructed to receive ultrasound tactile feedback.

3.2 Integrating with feedback devices

UMH feedback in VR appears to yield the most engaging results in congruent multimodal contexts, where the relatively low-fidelity UMH sensations are supported with proper visual and auditory cues to create a compelling experience [23,51] (see also chapters “Opportunities for Multisensory Mid-Air Interactions Featuring Ultrasound Haptic Feedback” and “Multimodal Interaction with Mid-Air Haptics”).

Integrating UMH into immersive VR likely relies on adequate co-location and synchronisation of the UMH stimuli with the visual (see chapter “Superimposing Visual Images on Mid-Air Ultrasonic Haptic Stimulation”) and auditory

⁶ <https://www.ultraleap.com/product/leap-motion-controller/>

⁷ <https://www.microsoft.com/en-us/hololens/hardware>

⁸ <https://www.magicleap.com/en-us/magic-leap-1>

⁹ <https://www.vive.com/us/product/#pro%20series>

stimuli provided through the HMD as well as other haptic stimuli from various other haptic interfaces. While the extent of this reliance has not been studied for UMH in particular (with the exception of [66], who show that some spatial incongruence may be tolerated), it is reasonable to assume its existence based on prior results on spatially (e.g. [75]) and temporally (e.g. [14]) incongruent stimuli when interacting with virtual environments. Achieving this co-location relies on linking the UMH device frame of reference in which stimuli are positioned with the VE world frame of reference in which other stimuli are positioned, and can be achieved through an initial calibration. While this co-location does not necessarily pose any technical issues with regards to visual or auditory stimuli, co-location with other haptic stimuli can be problematic, either because the actuators delivering those stimuli obstruct the user’s skin at the target location, or because of masking effects. We e.g. regularly observed masking effects when providing simultaneous UMH and wearable vibrotactile feedback, even when the latter was spatially removed from the site of UMH stimulation. We hypothesize this is mainly due to the differences in stimulus intensities that are achievable between both technologies, however we have not yet been able to formally evaluate this hypothesis.

Currently, all software dealing with correctly timing outputs for visual, audio and haptic feedback needs to be developed from scratch, which can prove complex due to the differences in frame rates and requirements between the feedback devices. In order to generate inputs at the frequencies required by UMH devices’ transducers (currently usually 40kHz) and to guarantee the desired temporal modulation of stimulus parameters, the generation of UMH stimuli needs to run on a dedicated thread ensuring output frequency for commands to the device is strictly respected. Considering that in addition, tracking input in immersive VR applications is usually available at frame rates similar to those of the visual feedback, it become apparent that a non-trivial oversampling problem needs to be solved in order to generate UMH feedback that evolves in response to user motion and input. Common solutions to the issue are to generate haptic stimuli independent of user motion, i.e. in a passive manner (e.g. the stimuli used in the perception experiments discussed in Sec. 4.2 [31,61]), to use pre-defined complex stimulus patterns that are triggered by events in the VR interaction (e.g. button-press feedback in [37] or the sensations corresponding to supernatural phenomena in [51]), or to extract frame-rate-independent characteristic values from the tracking input from which UMH feedback can be procedurally generated at whichever frequency is best suited to the UMH device. Such features could e.g. either be user motion speed or the user’s estimated position with respect to pre-defined virtual object features, as is the case in existing works on UMH texture rendering [18,5]).

Beyond co-location and synchronisation of UMH stimuli with other sensory inputs, it is necessary to ensure UMH devices do not generate unwanted stimuli which may break immersion. By their very nature, UMH devices are somewhat at an advantage over contact haptic devices in this regard as their contactless operation guarantees that with adequate placement, they will not generate un-

wanted haptic stimuli through contact with the actuator. However, UMH devices can still generate unwanted haptic sensations through two mechanisms. The first is the unintentional generation of airflow around the focal point. This point has been addressed in early literature on UMH device design [28,86]. The second is through the generation of side lobes, i.e. other potentially perceptible high-pressure areas occurring in the device workspace due to its function principle. This issue can be addressed to a certain extent through UMH device design [68], or by using tracking input to activate UMH feedback only when the user is reasonably close to the intended point of stimulation.

Aside from unintended haptic stimuli, UMH devices may produce unintended auditory stimuli. This issue is somewhat secondary in immersive VR as users generally wear headphones occluding outside sounds which may arise from device operation. It can also be addressed to a certain extent through optimisations to rendering algorithms [27].

It should be pointed out that on the topic of integration with other feedback devices, UMH also brings key opportunities to interaction with virtual environments. For example in the context of AR, the absence of equipment on the user's hand allows seamless integration of the real and virtual visual cues in the interaction area [42,71]. While this does not apply to integration with VR visual feedback, as the user's hand is fully occluded from view, it does apply to concurrent manipulation of tangible objects in a VR setting as the user's hand is kept free from any encumbering device. We plan to investigate this last point in future works.

3.3 Workspace compatibility

Despite the reasonably large workspace of UMH devices and the ability to freely place and move focal points within it, the workspace of ultrasound arrays rarely matches the large size of VR interaction zones. It is therefore necessary to find solutions to make these workspaces compatible.

This can be achieved either through software approaches, hardware approaches, or any combination of both. Software approaches to this problem rely on designing the VR interaction to fit within the rendering area of the focused ultrasound array. This can mean shrinking the effective interaction area (e.g. [23,51]), or leveraging redirection techniques. To the best of our knowledge, the latter have not yet been explored for UMH feedback, but there is extensive literature on the subject relating to passive haptic feedback (e.g. [9]) which may be applied in this context. Hardware approaches on the other hand aim at enlarging the actual workspace within which UMH feedback can be delivered. This has been achieved through the use of multiple synchronised and larger arrays, an approach discussed in detail in chapter "Sound-Field Creation for Haptic Reproduction". For example, Inoue et al. [35] developed a prototype UMH holographic display using four arrays in a cubic arrangement to enable a stronger focus in the centre and multi-directional stimuli. A similar arrangement was used in [43,85], and another with larger arrays was used by [53]. More recently, passive acoustic reflectors have also been proposed to enlarge the workspace of an array without




Fig. 2. The workspace compatibility problem: the rendering workspace of UMH devices is usually small compared to the size of VR interaction spaces, limiting the set of virtual objects for which UMH feedback can be provided (Left). This issue can be tackled through software (Top right), hardware (Bottom right) or combined approaches.

increasing its complexity [2]. Finally, an alternative approach consists in making the UMH arrays mobile, either by mounting them on the user [77] or onto robotic manipulators, in an approach akin to encounter-type haptics [7].

To enlarge the rendering workspace of a UMH device, we have developed an open-source robotic solution for making ultrasound arrays mobile called PUMAH¹⁰ [33].

The device functions as a servo-driven pan and tilt mount for a focused ultrasound array (see Fig. 3). The pan and tilt servomotors are HiTec HS645-MG and HS625-MG¹¹, respectively. We chose these for their low cost, high holding torques of 0.94Nm and 0.66Nm and high rotation speeds of 300°/s and 400°/s, respectively. The servomotors are driven using a Pololu Mini Maestro-24¹² USB servo controller board. The board is powered by a 5V, 2.4A DC power supply and receives position commands generated within a VR application running on a computer connected via USB. To limit unnecessary bandwidth usage and computations on the microcontroller, updates to the target positions for both axes are only sent if they differ from the previous target position, at a rate up to the frame rate of the VR application (approx. 90Hz).

Pan and tilt error angles are computed based on the angular error between the array normal and the vector from the device pivot (intersection of $Pivot_x$ and $Base_z$ in Fig. 3) to a target defined on an application-specific basis. They are then used as commands for the servomotors on each axis. Targets can range from static positions in the VE coordinate frame to dynamic tracking of the user’s hand. There are two main steps to integrate the device into a VR environment: an initial calibration and a definition of the target orientation (e.g. through user hand tracking). For calibration, we drive the device to the $\{0^\circ \text{ tilt}, 0^\circ \text{ pan}\}$ po-

¹⁰ <https://gitlab.com/h-reality/pumah>

¹¹ <https://hitecrd.com/products/servos>

¹² <https://www.pololu.com/product/1356>




Fig. 3. Assembled (left) and exploded (right) view of the PUMAH. The array (1) is mounted on an aluminium holding plate (2) using 3D-printed ABS clips at the corners (3). The plate rotates around the x_{arr} -axis (tilt) within an aluminium tubing and ABS frame (4), which itself rotates around the device vertical axis z_{base} (pan). The axes are driven by HiTec HS645-MG (5) and HS625-MG (6) servomotors. They are mounted on bearings held within aluminium chassis (7), relieving the motor shafts of any radial loads. The complete system is mounted on a 3D-printed ABS foot (8), which can be screwed to a supporting structure or mounted on a tripod using M6 screws. *Adapted from [33].*

sition and acquire a known position and orientation of the device base using an HTC Vive tracker. The virtual model of the device (and the virtual scene elements built around it) are then rotated and translated such that they coincide with positions of their real-world counterparts. Once the virtual and real devices are aligned, a target in the workspace needs to be defined and its position computed or acquired through tracking. The vector between the device pivot and target is computed at every frame. The tilt and pan angular errors can be computed by projecting this vector onto the array y-z plane and device base x-z-plane, respectively. To verify the correct execution of the tracking and angle computations, the virtual model of the device is rotated accordingly, thus acting as an ideal representation of our system’s target behavior. The computed tilt and pan error angles are then translated into target pulse width modulation (PWM) values for the servo-motors based on prior modelling of the relationship between PWM and servo shaft angle.

Our hardware tests and human subject study in an ecological VR setting (reported in [33]) show that a 14-fold increase in workspace volume is achievable, with focal point repositioning speeds over 0.85m/s through device movement alone, and with feedback accuracy below 18mm in the worst case scenario. We also show that through repositioning of the array to better align it with the user’s hand, the PUMAH enables significantly higher stimulus intensities to be delivered throughout the interaction workspace.




Fig. 4. Interaction use-cases developed to showcase PUMAH’s capabilities: (a) Touchable holograms; (b) Virtual solar system; (c) Mid-air haptic fountain; (d) Haptic table-tennis. *Reproduced from [32].*

We designed a set of interaction use-cases for PUMAH leveraging the device’s capabilities in relation to extending the workspace (A and B in Fig. 4) and pro-

viding multi-directional feedback (C and D in Fig. 4) [32]. At the same time, we applied software approaches as previously discussed, designing the interactions to naturally fit within the workspace provided by PUMAH.

Our use-cases are built as four distinct virtual scenes (see Fig. 4). All interactive objects are placed inside the PUMAH’s workspace, and produce various tactile sensations when the user’s hand avatar collides with them. All sensations are generated using Ultrahaptics’ Unity Core Asset¹³. To avoid collision between the PUMAH and the user’s hand, solid virtual objects are placed in the PUMAH’s location in most of the scenes. If the user’s hand gets too close to the array, red bounding boxes are shown around the device position to warn the user.

a) Touchable holograms: A holographic projector displays a circular menu of 4 shapes around it (see Fig. 4-A). Users use a Vive controller to navigate through the menu of holograms. Fixed intensity vibrotactile feedback is provided when the user’s hand collides with the virtual objects. To experience the increase in workspace provided by the PUMAH, the scene also features a button allowing switching between full use of the PUMAH’s workspace and use of only the static array’s workspace.

b) Virtual Solar System: Here, the user can interact with three celestial bodies within a model of the solar system (see Fig. 4-B). The sun is similar to the touchable holograms, but also emits solar flares which can be felt as sweeping vibrotactile sensations. For the earth, different sensations are produced when the ocean or land are touched. Finally, the rings of saturn produce various impact sensations.

c) Mid-air haptic fountain: A fountain is colocated with the PUMAH in the virtual scene, allowing the user’s hand to interact with the flowing water from various angles (see Fig. 4-C). The tactile sensation is continuously updated to give an impression of water flow around the hand. This demo particularly showcases an application for multi-directional feedback enabled by the PUMAH.

d) Haptic table-tennis: Here, the PUMAH provides haptic feedback for a game of VR ping-pong (see Fig. 4-D). Vibrotactile impacts are rendered with an intensity proportional to the ball’s impact force on the hand. Because the virtual impacts are so short, we render impacts as a decaying vibration on the hand over a few frames following the actual virtual impact so as to enhance user’s perception of the stimuli. This demo was meant to showcase an application of multi-directional UMH feedback and the high device repositioning speed.

4 Perception of UMH stimuli

Because UMH interfaces are so novel, relatively little is known about the specifics of how UMH stimuli are perceived, nor about the relationship between a stim-

¹³ <https://developer.ultrahaptics.com/downloads/unity-plugin/>

ulus’ parameter and its perceptual qualities. Understanding both these points is a critical challenge for all UMH applications, immersive VR included, for three main reasons. First, it is a key step in determining the usable ranges and discernible levels of stimulus parameters for rendering. Second, it can provide crucial insights into ways for improving the perception of stimuli and overcoming certain limits of UMH technology such as e.g. weak perceived intensity and poor spatial definition of the stimuli. Finally, in the context of immersive VR or other multi-modal applications, it is essential to understand how UMH stimuli interact with other haptic, audio and visual stimuli in order to effectively design interactions.

Early work on perception of UMH stimuli has found that UMH devices are most effective at stimulating the glabrous skin, where they primarily stimulate the Pacinian corpuscles, and as such UMH devices act essentially as vibrotactile devices. There is some evidence that with the proper choice of stimulus parameters, UMH devices are also capable of stimulating non-glabrous skin [24], and also activate Meissners corpuscles [24,62] as well as slow-adapting tactile receptors [36]. Rakkolainen et al. [69] provide an overview of the state-of-the-art on UMH stimulus perception which goes into more detail on these aspects.

To determine the usable ranges and discernible levels of UMH stimulus parameters, psychophysics studies are conducted with the aim of determining detection and discrimination thresholds. Several studies have focused on this, with some of our own work on the topic [31,50,61] discussed in more detail in the present section, however given the complexity of UMH stimuli there is still much to be discovered on this subject. We provide an in-depth overview of the state of the art on the subject in the work published by Mulot et al. [61].

We have not directly tackled the topic of overcoming limitations of UMH devices with regards to the perceived qualities of the stimuli. However, the framework described in Sec. 4.1 is intended to simplify studies on the topic, such as e.g. those conducted by Frier et al. [19,20] on the impact of sampling strategies on perceived stimulus intensity.

Finally, the study of multimodal integration of UMH stimuli has not yet been tackled by our team or most others working on the subject. A notable exception to this is a recent publication by Pittner et al. [66] which investigated the effect of presenting spatially incongruous visual and UMH feedback in VR, showing that it can be used to generate the illusion of avatar ownership despite the incongruence.

4.1 DOLPHIN: A framework for designing and studying perception of UMH stimuli

A key takeaway from our work on perception of ultrasound tactile feedback was that the high dimensionality of the stimulus design space makes the design and study of ultrasound haptic stimuli a very complex task. Indeed, the choice of an UMH stimulus’ shape, spatial, temporal and modulation parameters all interact and impact a user’s perception of the stimulus properties as well as perceptual thresholds. To efficiently explore perceptual implications of choices made within

this design space, we developed DOLPHIN¹⁴ [61], an open-source framework facilitating the design of UMH stimuli and designed to interface with 3rd party rendering applications as well as perceptual evaluation tools (see Fig. 5).

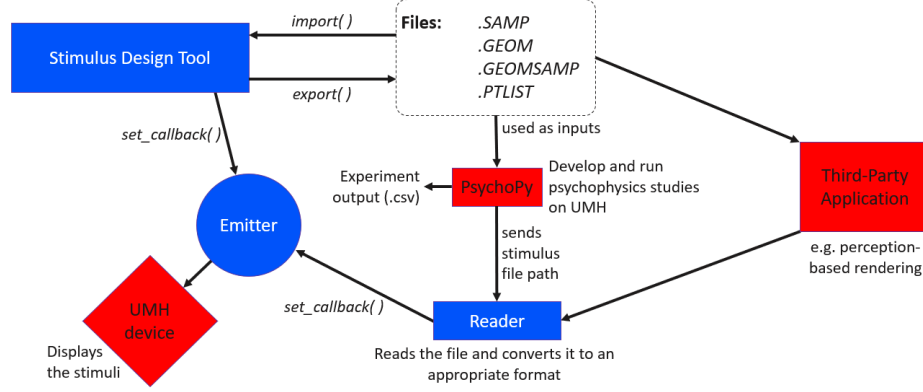


Fig. 5. The DOLPHIN software framework for designing and studying perception of UMH stimuli. Framework components are highlighted in blue, 3rd party components in red. Arrows and their corresponding text show how the different parts are linked. The core of the framework is a stimulus design tool providing full control over spatial, temporal and intensity aspects of the stimuli in an accessible fashion. Export from the design tool is possible in four formats corresponding to the shape (.GEOM), the sampling strategy (.SAMP), or both (.GEOMSAMP, .PTLIST) with the generic .PTLIST format intended for use in external tools. Stimulus files can be used in experiments or third-party applications thanks to a reader program which sends data to the emitter for rendering. *Re-drawn from [61].*

At the core of this framework is a stimulus designer tool (see Fig. 5) which helps users configure UMH stimuli, controlling both their geometric shape properties as well as the sampling strategy used to render them. Expanding on the concept introduced by Frier et al. [20], we define the combination of a spatial discretization of an abstract shape and a set of rules for the temporal display of order and intensity modulation of the resulting points as a sampling strategy (see Fig. 6). Since all UMH stimuli based on focal points can be defined as a temporal sequence of focal point locations in space, DOLPHIN’s approach to stimulus design based on discretized geometries remains extremely general and adaptable to almost any UMH rendering scenario.

As touched upon previously, an UMH stimulus has two main aspect. The first concerns the stimulus’ spatial properties, i.e. its geometry and position in space (see Fig. 6-A). Most of the time the shape can be represented as a parametric function $f : [0, 1]^k \rightarrow \mathbb{R}^3$ whose number of parameters differs depending on the specifics of the shape. For time-dependant shapes such as a rotating line, a time

¹⁴ <https://gitlab.com/h-reality/dolphin>

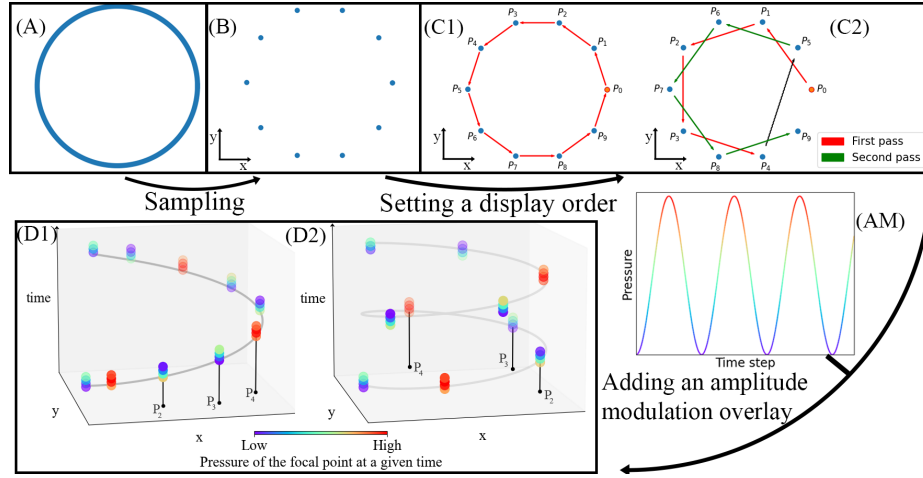


Fig. 6. A sampling strategy represents the process of transforming an abstract geometric shape (A) into a UMH stimulus (D1, D2) through 3 successive steps: spatial discretization (B), definition of a temporal display order (C1 and C2 show two examples of this for the same spatial discretization), and overlay of a time-dependent amplitude modulation (AM). D1 and D2 show the respective application of AM to C1 and C2, with the vertical axis representing time, weak intensities shown in blue and high intensities shown in red. In this example, the dwell time is a third of the AM period. *Reproduced from [61].*

parameter can be included here. The second aspect of the UMH stimulus is its sampling strategy (see Fig. 6). We define a set of N sampled points $P^s = \{q_i, 0 \leq i < N\}$ s.t. $\forall i, \exists t_{1,i}, \dots, t_{k,i}, q_i = f(t_{1,i}, \dots, t_{k,i})$ as a point on the shape. Based on this, the temporal sampling order can be formally defined as a finite series $P^{s,o} = ((P_j, T_j))_j$ where $P_j \in P^s$ for all j and T_j is the dwell time on point P_j . A fully defined configuration is obtained for each timestep $C_k = (P_{j_k}, I_k)$ where $0 \leq j_k < N$ and $I_k = \frac{p(P_j)}{p_{peak}(P_j)}$ represents the pressure to peak pressure ratio for the device at the position P_j , with $p : \mathbb{R}^3 \rightarrow \mathbb{R}^+$. The sampling strategy encompasses all information about how to sample a shape, transforming a potentially continuous shape into a discrete and finite representation. Since the sampling strategy itself is independent of the underlying abstract shape, a same sampling strategy can be applied to different shapes.

DOLPHIN’s designer tool provides the user with information about the physical attributes of the designed stimulus, allows haptic previews of the designed stimuli as well as numerous stimulus import and export functions for incremental design and interfacing with 3rd party applications. A more detailed description of the software’s inner workings and the interface with 3rd party applications is presented in [61]. The framework’s objective is thus not only to accelerate the design and psychophysics evaluation process, but also to lower the barriers of entry to the field of UMH, ease communication in the field by standardizing the stimulus design workflow, and ease replication of UMH perceptual experiments.

4.2 Perceptual thresholds

This section presents a summary of our team’s work on perceptual thresholds in UMH, focusing on the need to determine usable ranges and discernible levels of stimulus parameters for rendering.

All perception experiments described in the following used an identical set-up, where subjects were seated in front of the array, with one of their hands supported by an armrest, placed at a fixed height above the array (usually 10cm to achieve maximum possible intensity or 20cm to allow the largest horizontal extent of perceptible stimuli) with the palm parallel to the array plane. The array used was an Ultraleap STRATOS Explore device, for which changes in intensity ranging between 0% and 100% correspond to acoustic pressure at the focal point between approx. 0–1125 Pa for stimuli centered in a plane 20cm above the array. Subjects always wore noise-cancelling headphones playing pink noise to mask the device operating noise. Instructions were displayed on a computer screen placed directly in front of them, and they provided responses to the experimental questions through a keyboard operated with their non-dominant hand.

Intensity Detection Thresholds All our experiments aimed at determining detection thresholds used 1-up, 1-down staircase methods, with the threshold estimate taken as the mean of the last 8 of 9 reversals. 15 healthy subjects (8 male, 7 female; ages 22-32 (M=25); 2 left-, 13 right-handed) participated in a first experiment aimed at evaluating the 50%-detection threshold for a single

focal point generated 20cm above the array using AM, both in a passive touch (immobile hand) and an active touch (moving hand) condition (see Fig. 7-A). The focal point’s amplitude was modulated at a frequency of 200Hz as this is close to peak sensitivity for vibrotaction in humans.




Fig. 7. (A) Schematic of the stimulus and user hand placement. A fixed AM focal point was generated at (0,0,20cm) in the array’s frame of reference, and the user’s hand was either placed with the center of the palm fixed at the focal point position (passive touch condition) or free to explore the plane at $z=20\text{cm}$ in the array’s frame of reference. (B) Mean subject 50%-detection threshold estimates for both conditions. Blue dots show the actual values, boxes the interquartile range (IQR), the blue horizontal line shows the median while the red lines show the mean plus or minus one standard deviation. Adapted from [31].

We determined the 50%-detection threshold for a single focal point generated with AM, finding it to be largely independent of hand movement and on average between 48.7% and 49.5% intensity, i.e., approx. 560 Pa peak acoustic pressure (see Fig. 7-B).

Our second experiment evaluated the detection threshold for a 15cm line pattern rendered through spatio-temporal modulation (STM)[40]. Given that user motion had no impact on AM focal point detection, we hypothesized this to hold true for STM and only investigated an active condition. Instead, we chose to investigate possible effects of line orientation relative to the hand. Subjects were presented with lines aligned with the device x- or y-axis at a height of 20 cm (see Fig. 8-A). We selected a focal point movement speed of 7 m/s, which is close to the best perceived intensity [19]. The experiment was performed in two parts: the first one investigated the detection threshold regardless of the perceived pattern (50%-detection), while the second one investigated the intensity required for users to reliably feel the displayed pattern as a line (50%-identification threshold). Subjects responded to the binary questions “Did you feel a stimulus?” and “Did you feel a line?” for the detection and identification experiments, respectively. 7 healthy subjects (5 male, 2 female, ages 22-32 (mean: 27.2), all right-handed) participated in the STM detection thresholds experiment, and 12 different healthy subjects (9 male, 3 female, ages 22-47 (mean:

28.3), 11 right handed, 1 left handed) participated in the STM identification thresholds experiment.

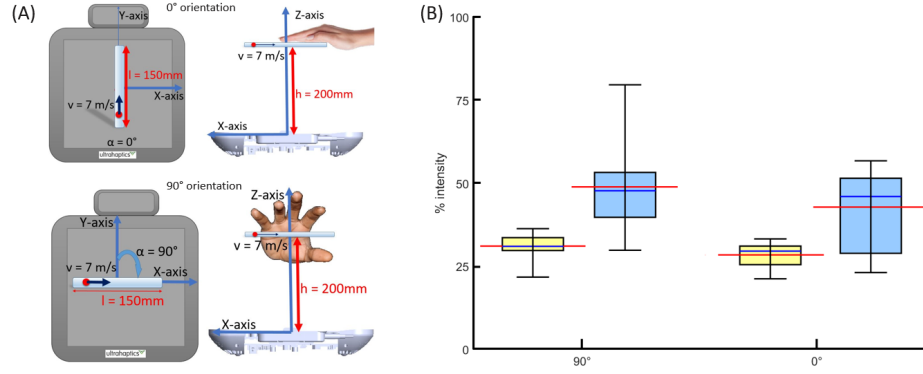


Fig. 8. (a) Schematic of the stimuli and user hand placement. 15cm STM lines were displayed in a plane 20cm above the array in one of two possible orientations relative to the user's hand. (b) Mean subject 50% detection (yellow) and identification (blue) thresholds. The boxplots show the median, IQR and total spread of values while the horizontal red lines show the mean for all subjects. *Adapted from [31].*

For both the 0° and 90° orientations, the measured identification thresholds are well above the detection thresholds (median 46% for identification vs. median 29.6% for detection in 0°, median 47.8% for identification vs. median 31% for detection in 90°). Pairwise comparisons between these thresholds in the 0° and 90° orientations show these differences to be significant with $p < 0.05$ and $p < 0.01$ respectively (Wilcoxon rank-sum tests).

Intensity discrimination We conducted an intensity discrimination experiment for STM stimuli on the palm in the context of the evaluation of an approach for rendering virtual object stiffness with UMH (see also Sec. 5.2) [50].

Subjects performed an experiment in VR which consisted in comparing the stiffness of two virtual pistons. Each virtual piston was modeled as a 1-D spring following Hooke's law. Whenever a user touched the piston, the system simulated a spring-like feedback, where the pressure generated at the focal point by the array was defined by $p = k(z_0 - z) + p_0$ if the user touched the piston, 0 otherwise. k is the simulated stiffness of the piston (in Pa/m, sound pressure over displacement), z the current stroke of the piston, $z_0 = 30\text{ cm}$ its free length, $\Delta z = z_0 - z$ its current travel, and $p_0 = 146.87\text{ dB SPL}$ (441 Pa) the absolute detection threshold we registered at 30 cm (when $\Delta z = 0$). The pistons were fully compressed at $z = 20\text{ cm}$ ($\Delta z_{\max} = 10\text{ cm}$). The stimulus was rendered as a small STM circle parallel to the array plane centered on the centroid of the

piston's upper plate (see Fig. 9-(A)). When the user interacts with the piston, this point coincides with the center of the user's palm.

Twenty subjects (16 males, 4 females, ages 21-29 (M=24), all right-handed) participated in the experiment. They viewed the virtual environment, composed of a virtual piston placed on a black table, through an HTC Vive VR headset, and used a HTC Vive controller held in their non-dominant hand to enter responses to the experimental question (see Fig. 9-(B)).

Three differences of piston stiffness $|k_{ref} - k_{test}|$ were possible: 0.001176 N/mm, 0.001764 N/mm and 0.002353 N/mm, corresponding to the absolute values of the difference of the possible stiffness of the test piston vs. those of the reference piston.

The order of presentation of the two pistons was counterbalanced: every couple of pistons was presented in all orders. The starting reference was also alternated to ensure that fatigue did not influence the last block. Thus, subjects were presented with 90 trials per reference stiffness (270 in total), divided into 5 blocks of 6 trials in a randomized order for each block.

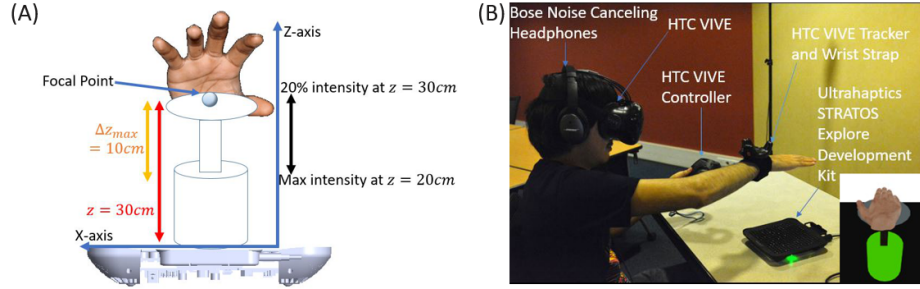


Fig. 9. Intensity discrimination experiment: (A) Schematic of the feedback design, the user feels an STM stimulus centered on their palm as they interact with a virtual piston. The intensity goes from its minimum values at a height of 30cm above the board to its maximum at a height of 20cm above the board. (B) Experimental set-up and view of the VE through the HTC Vive HMD. Users hold their dominant hand outstretched over a table holding the UMH device to press on the virtual piston, while wearing noise cancelling headphones to avoid any bias due to device operating noise. *Reproduced from [50].*

Fig. 10 shows the psychometric functions obtained for each of the three considered references stiffnesses. For the reference stiffness K_{ref1} , we obtained a 75%-Just Noticeable Difference (JND) value of 20% and the Point of Subjective Equality (PSE) at +2.16%. For K_{ref2} , the 75%-JND was of 32% and PSE of +3.65%. Our results for K_{ref3} differed from the others as we were not able to reach proportions of correct answers close to 1 on the right-hand side of the curve. This could have been due to the pressure being delivered at K_{ref3} being close to the device maximum output (at 80% of the maximum). Also in this condition, subjects often reported feeling the displayed shape changing from a

circle into something else. We therefore evaluated the psychometric only taking into account the stiffness intensities for which users reported feeling a circular shape (blue points in Fig. 10), yielding a 75%-JND of 18% and PSE of +0.58%.

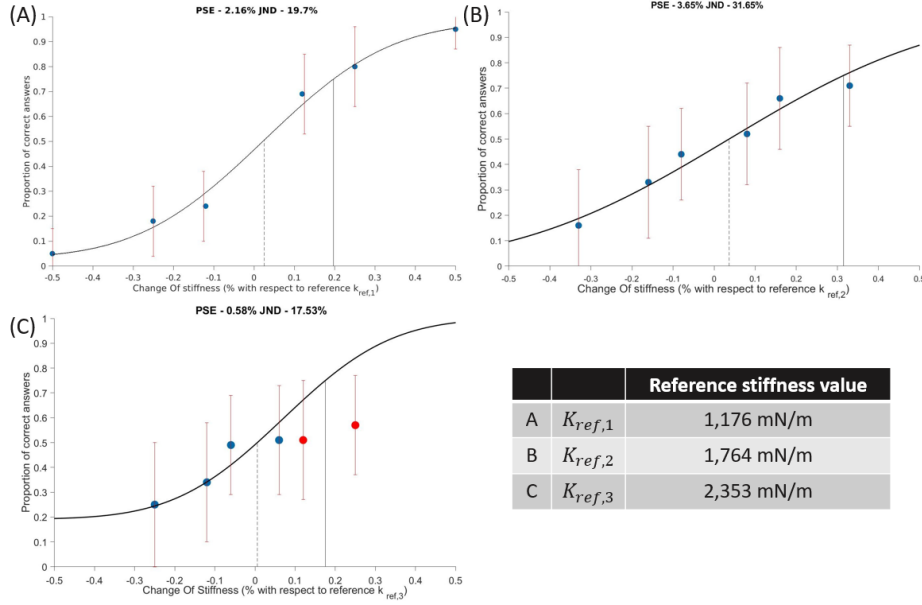


Fig. 10. Psychometric curves for the three reference stiffness values, fitting a cumulative Gaussian to the data. We plot the proportion of correct answers as a function of the percentage increase in stiffness with respect to the reference $k_{ref,i}$. The vertical dashed and solid lines represent the PSE and the 75%-JND, respectively. Error bars represent standard deviation. *Adapted from [50].*

Line orientation discrimination In our experiments, subjects repeatedly reported a certain degree of confusion regarding the orientation of STM lines, especially those rendered at lower intensities. This prompted us to investigate line orientation discrimination performance and the effect of intensity on it. Although subject's hand movement was not found to impact detection thresholds, we hypothesized that it may impact discrimination of orientations.

Subjects were presented with a randomized sequence of lines in four different orientations, $\alpha = 0^\circ; 45^\circ; 90^\circ; 135^\circ$, at three different intensity levels, a low-intensity level (54% of maximum intensity), an intermediary intensity level (77% of maximum intensity), and a high intensity level (100% of maximum intensity). For each stimulus they were asked to attribute it to one of the four possible orientations. The low intensity level was chosen to be above the 50%-identification threshold of 90% of the sample population (see Sec. 4.2).

Subjects performed two trial blocks with 120 trials each (10 for each pair of intensity and orientation levels), one block in a passive touch condition, one in an active touch condition. 10 healthy subjects (8 male, 2 female, ages 22-44 ($M=27.2$), all right handed) participated in the study.

We obtained confusion matrices for each intensity level (see Fig. 11). Overall, the majority of orientations were correctly identified. 0° lines were occasionally confused with 45° or 135° , but never with 90° lines (both in active and passive conditions). The most confused orientation is 135° with 90° . A 2-way ANOVA showed neither orientation, nor intensity, nor interactions between both significantly affect correct identification rates.

Analysing the effect of intensity and condition (passive vs. active) on correct identification rates for the 0° orientation, a 2-way ANOVA showed a significant effect of condition ($p < 0.01$) on discrimination accuracy. No other significant effect were observed, except for the confusion rate between orientations 0° and 90° , where a significant effect of the condition on the confusion rate ($p < 0.05$) was found. The confusion rate between orientations 90° and 135° was also significantly affected by the interaction between intensity and condition ($p < 0.1$).

		Perceived				
		0°	45°	90°	135°	
PASSIVE CONDITION	Actual	0°	99.3 %	0.7 %	0 %	0 %
		45°	4 %	94 %	2 %	0 %
		90°	0 %	3.4 %	95.6 %	1 %
		135°	4 %	1 %	5.7 %	89.3 %

		Perceived				
		0°	45°	90°	135°	
	Actual	0°	99.3 %	0.7 %	0 %	0 %
		45°	6 %	93 %	1 %	0 %
		90°	0 %	3.9 %	95.1 %	1 %
		135°	2 %	1 %	3.1 %	93.9 %

		Perceived				
		0°	45°	90°	135°	
	Actual	0°	99.2 %	0.8 %	0 %	0 %
		45°	7 %	89 %	2 %	2 %
		90°	0 %	5.3 %	92.7 %	2 %
		135°	3 %	0 %	1 %	96 %

		Perceived				
		0°	45°	90°	135°	
ACTIVE CONDITION	Actual	0°	97 %	2 %	0 %	1 %
		45°	2 %	96 %	2 %	0 %
		90°	2 %	2 %	96 %	0 %
		135°	2 %	1 %	1 %	96 %

		Perceived				
		0°	45°	90°	135°	
	Actual	0°	93 %	5 %	0 %	2 %
		45°	3 %	97 %	0 %	0 %
		90°	1 %	4 %	95 %	0 %
		135°	3 %	2 %	1 %	94 %

		Perceived				
		0°	45°	90°	135°	
	Actual	0°	95 %	3 %	0 %	2 %
		45°	2 %	95 %	0 %	3 %
		90°	1 %	2 %	92 %	5 %
		135°	1 %	1 %	5 %	93 %

Minimum intensity (54%)

Intermediary intensity (77%)

Maximum intensity (100%)

Fig. 11. Mean confusion matrices for each condition. Correct identification rates are highlighted in green. The most prevalent confusions are highlighted in orange, the second most prevalent in yellow. *Reproduced from [31].*

Curvature discrimination We conducted a user study on the effect of sampling strategy parameters on curvature discrimination [61] for dynamic tactile pointers [25] displaying arcs of four possible curvature radii across the width of

the subject's hand for a fixed duration of two seconds. 19 volunteers (2 f., 16 m., 1 non-binary, mean age \pm std = 23 ± 3.4), all right-handed) participated in the experiment. The four different ratios of curvature radii to hand width were equal to $\frac{1}{\pi}$ (semi-circle), $\frac{3}{2\pi}$, $\frac{5}{2\pi}$, ∞ (straight line). For the computations, the radius to hand width ratio for the straight line was set to 1.99 which is the smallest ratio ensuring the height difference at the edge of the hand remained smaller than the radius of a focal point (see Fig. 12).

We evaluated five sampling strategies with different numbers of spatial sampling points ($n = 10, 50, 200, 400$, and 800), resulting in dwell times on each point of $t = \frac{2}{n}$ seconds. Based on Hajas et al.'s work [25], the static amplitude modulation frequency was set to 200 Hz. Each stimulus was displayed 10 cm above the array and the tactile pointer always moved from the thumb side towards the pinkie.




Fig. 12. Curvature discrimination experiment stimuli: (A) We used four different curvature radii for curves scanning across the hand from the thumb towards the pinkie (as indicated by the red arrow). The curvatures were defined in relation to subject hand width w_{hand} using a ratio k such that $r = k \cdot w_{hand}$. (B) In theory, the flat line stimulus would require $k = +\infty$, however in practice we consider a curve to be effectively flat if k is sufficiently large so as to ensure that the difference in height between the centre and end of a curve is inferior to the diameter of a focal point $\varnothing FP$ (yielding $k \geq 1.99$). Each curve was uniformly divided into n spatial sampling points ($n = 10, 50, 200, 400$, and 800). In the interest of readability, only the sampling points for stimuli with $n = 10$ are graphically represented by the purple dots.

The experiment was divided into five blocks corresponding to each of the five studied sampling strategies. Each block followed a two-alternatives forced choice protocol where pairs of stimuli with different curvatures were presented with a 1.5 s break between them, after which users had to indicate which felt closest to a straight line. Stimulus pairs were repeated three times per block, yielding a total of thirty-six trials per block. The pair order within blocks and the order of blocks were randomized.

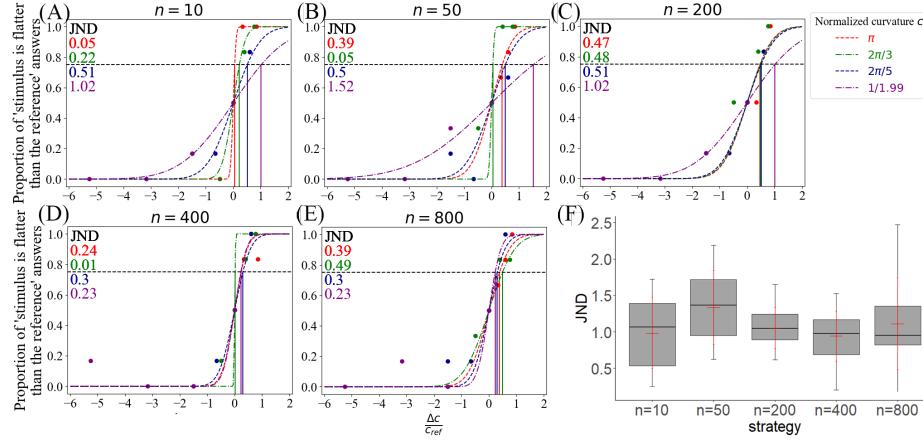


Fig. 13. Experiment results. (A-E) Examples of cumulative Gaussian curves fit to a single subject’s results for each strategy, with one curve fit per reference stimulus. The proportions of “stimulus flatter than reference” answers from the user are plot against the Weber fraction in curvature. For a stimulus with radius r_i , the curvature c_i is defined as $\frac{1}{r_i}$ and the Weber fraction with stimulus 1 as a reference is defined as $\frac{c_2 - c_1}{c_1}$. JND estimates for this strategy are reported on the left. (F) Box plot of the mean JNDs obtained for each strategy, using the average of the four JNDs corresponding to the different reference stimuli for each user and strategy. Red lines represent the mean and one standard deviation. *Adapted in part from [61].*

For each user and each strategy, we plotted the proportion of ‘stimulus is flatter than the reference’ answers against the relative difference in curvature for each of the four references, as shown in Fig. 13. For a reference stimulus with radius r_1 and test stimulus with radius r_2 , the x value corresponding to the pair was $\frac{c_2 - c_1}{c_1}$, where $c_i = \frac{1}{r_i}$ corresponds to the curvature of an arc with radius r_i . We added the hypothetical point at (0, 0.5) corresponding to a fully random answer when both stimuli are identical. We then fitted the set of observations with a cumulative gaussian centered on 0 and used the curves to calculate the 75%-JND estimate in curvature for this user, strategy and reference. The obtained JND is thus expressed as the Weber fraction in curvature (see Fig. 13).

Outlier JND values (values greater than $Q_{75\%} + 1.5IQR$ where $Q_{75\%}$ is the 75% quartile and IQR is the interquartile) were removed, yielding mean subject

JND distributions for each strategy (see Fig. 13-F). These distributions were assumed to be normal (Shapiro-Wilk test did not reject H_0 with $p > 0.05$ for all five distributions). We therefore performed an ANOVA which revealed no significant difference between strategies in terms of mean resulting JND in curvature ($F = 1.676$, $p = 0.165$). We conclude that when designing dynamic tactile pointers, the number of points used for the sampling strategy does not have a significant impact on users' ability to discriminate arc curvatures. The mean JNDs obtained were 0.98, 0.98, 1.33, 1.05 and 1.11 for the strategies with $n = 10, 50, 200, 400$ and 800 points respectively.

The presence of success rates above 50% in most cases shows that curvature discrimination is indeed possible for UMH dynamic tactile pointers, regardless of the sampling strategy used for rendering. However, curvature discrimination still appeared rather complex to perform on the haptic stimuli alone, echoing prior results on shape identification and discrimination for UMH. The number of spatial sampling points (and thus the dwell time per point) did not appear to significantly impact discrimination performance, meaning that UMH stimulus designers have more freedom when designing such stimuli as the strategy will likely not adversely impact stimulus discrimination.

5 UMH Rendering

Challenges relating to UMH rendering for immersive VR are two-fold. On the one hand, it is necessary to understand what properties of the VE and interaction between user and VE can be rendered using UMH, in a way that supports interaction by providing immediate, informative feedback and increases realism by accurately representing the VE's underlying processes. And on the other, using insights gleaned from the study of stimulus perception, it is necessary to understand how to best render said properties.

As discussed previously, a core component of UMH stimuli is their spatial and geometric properties (see Sec. 4.1). In the context of 3D immersive VR interaction, a fundamental aspect of rendering virtual objects with UMH is thus to render their 3D shape [47,45,44,52] and position in space, i.e. to render spatial information about contacts.

Beyond this, UMH has also been applied to rendering tactile surface properties such as textures [59,18,5]. In the first example [59], the authors modulate the waveform for an AM focal point projected on the finger to provide different textures to holographic UI elements. Freeman et al. [18] presented a method for rendering macroscopic texture features by varying the motion path of an STM focal point, along with its intensity, modulation waveform and frequency as well as the STM shape draw frequency. In the last examples [5], the authors present a method for extracting STM draw frequency and intensity parameters from graphics textures. These are then applied to a small circle projected on the fingertip as the user explores the visually rendered textures in AR or VR to provide different levels of perceived roughness.

It is important to note that existing work on texture rendering focuses on congruence between tactile and visual stimuli rather than on perceived realism of the tactile textures per se. As such, an application of these methods to VR should always consider the limitation that while UMH texture feedback may help create a coherent multi-modal experience, it may not have much in common with the represented real-world equivalent texture.

Because of the freely movable nature of UMH stimuli, early research on rendering also focused on rendering motion. This has been achieved by sequentially displaying AM focal points (e.g. [29,30,84]), adequate selection of STM draw frequency (e.g. [22,25]) or even tactile illusions [65]. While each of these methods has been shown to be effective, they are not specific to immersive VR. Interaction with virtual objects in motion can be relatively easy to design from scratch based on the research cited above or by e.g. making use of Ultraleap’s Sensation Editor¹⁵ which provides preset sensations with motion patterns. These stimuli can be easily ported to VR environments using Ultraleap’s Unity Core Asset¹⁶. Stimuli incorporating motion have been showcased in VR demos using UMH feedback [23,51].

Since UMH relies on a fluid coupling between the actuator and the skin (vibrations and pressure are transmitted through the air), physical properties that can be rendered in an ecological manner encompass interactions with fluids ([4,38]). Barreiro et al. [4] simulate the pressure field on the surface of the skin as a result of the hand coming into contact with virtual fluids, then use a custom algorithm to extract an STM stimulus path and intensity modulation from the data in interactive time. While this example is implemented with the virtual scene displayed on a screen, there are no inherent limitations to porting the visual display to a HMD and thus applying them to immersive VR. Jang et al. [38] propose an alternative particle-based fluids simulation from which they extract modulation parameters for multiple AM focal points projected on the user’s hand. They demonstrate the function of this approach in an immersive VR use-case where users interact with various liquid streams, sprays, rainfall or fountains.

Finally, UMH feedback has also largely been applied to the rendering of abstract information to support interaction. Some of these cues are purely abstract (e.g. [15,6]), others are designed to be evocative of underlying virtual mechanism (e.g. mid air buttons [37]). Dzidek et al. [15] design and evaluate a set of abstract UMH stimuli to support a wide range of interactions in an AR demonstrator. Van den Bogaert et al. [6] present a study on the selection of abstract UMH patterns intended for providing users with feedback when interacting with menus in mid-air in AR. Their study demonstrate a notable preference for relatively simple stimulus design when the aim is to support interaction, with a strong focus on the use of stimulus intensity and temporal parameters (continuous vs. discrete) as rendering parameters. In a VR application, Frutos-Pascual et al. [21] demonstrate the usefulness of abstract UMH feedback cues to the palm to

¹⁵ <https://developer.ultrahaptics.com/downloads/sensation-editor/>

¹⁶ <https://developer.ultrahaptics.com/downloads/unity-plugin/>

indicate successful grasping of virtual objects, in particular during manipulation of smaller objects.

Building on the understanding of stimulus perception discussed in Sec. 4, we are interested in understanding what properties of interactions with virtual objects can be effectively rendered with mid-air ultrasound haptics. On this topic, we have explored the use of stimulus intensity to encode information about shape and stiffness of virtual objects.

5.1 Rendering Local Shape

As previously discussed, the ability to rapidly move the focal point in the UMH device’s 3D workspace makes it an ideal candidate for rendering shape and geometric features. As such, STM patterns displayed in planes parallel to the device can be perceived relatively accurately, as the extent of the perceptible high pressure region around the focal point is usually considered to be on the order of one wavelength of ultrasound in air, i.e. approx. 8mm (see Fig. 14-left). However, due to the function principle of focused ultrasound arrays, these perceptible high pressure regions around the focal point tend to extend vertically over heights up to several cm, meaning that rendering small changes in height of a shape above the board is compromised, as close to the focal point path, these are completely masked by the vertical elongation of the focal point (see Fig. 14-right). It should also be noted that even in the most favorable case, i.e. rendering a shape in a plane parallel to the device, shape identification performance has been shown to be quite poor [74], although some methods have been proposed to improve it (e.g. [25]).

We therefore proposed rendering such local shape features along the array’s z-axis by keeping the focal point path in a plane parallel to the device, but varying the stimulus intensity along that path to generate the illusion of bumps or holes as the user explores the stimulus [31] (see Fig. 15-left).

We investigated the effectiveness of this rendering approach by performing a shape identification experiment in which users had to identify a given pattern among 5 possible shapes: a bump, a hole, and three lines of different intensities (low, medium and high) acting as confounders. Similar to our perception experiments on STM shape detection and identification thresholds (see Sec. 4.2), we once again performed a preliminary investigation into potential effects of line orientation on shape identification performance (see Fig. 15-right).

For this initial investigation, we chose to use the largest possible difference between baseline and peak intensity for the bump and hole patterns to maximise chances of correct identification.

12 healthy subjects (8 male, 4 female; ages 23-30 ($M = 25.3$); 10 right-handed, 2 left-handed) took part in the experiment. They performed two blocks of 90 trials, with blocks divided according to pattern orientation, and 30 trials occurring for each possible shape. The order of trials was fully randomized and the block order counterbalanced.

Overall, our results showed that identification rates were well above chance (30% correct) for all shapes in both considered orientations, indicating that

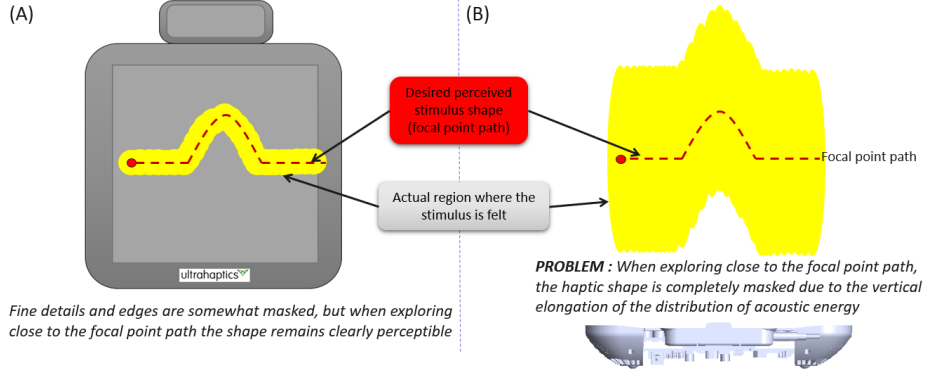


Fig. 14. Examples of shapes rendered in (A) a plane parallel to the device and (B) a plane perpendicular to the device. The red dotted line shows the focal point path while the yellow area schematically represents the area in which a stimulus is perceptible. As the user's hand explores the shapes close to the focal point path, it becomes obvious that lateral deviations of this path within a plane parallel to the device (A) can relatively easily be detected whereas vertical deviations within a plan perpendicular to the device (B) cannot.




Fig. 15. (Left) Proposed rendering concept for small changes in height of a shape above the board: shapes are rendered as STM patterns within a plane parallel to the device, while the intensity of the focal point is varied as a function of the linear coordinate along the shape. Drops in intensity are expected to be perceived as holes while increases in intensity are expected to be perceived as bumps. (Right) Stimuli used for our user study evaluating this approach: Subjects were presented with one of five patterns (hole, 3 lines of different intensities, bump) in one of two possible orientations. *Adapted in part from [31].*

subjects were capable of correctly attributing the stimuli to the intended shapes. However confusion rates between presented shapes ranged from 9.7% (confusion between bumps and holes) to as high as 25.3% (confusion between bumps and lines), indicating that the rendering approach is currently insufficient to reliably generate the desired percepts, at least with UMH feedback alone. Similarly to our prior study on STM line detection thresholds, shape orientation did not appear to have a significant impact on identification or confusion of shapes. A detailed analysis of identification and confusion rates is presented by Howard et al. ([31]).

While these results highlight certain limitations of this rendering approach, they are nonetheless encouraging as performances were well above chance level in the studied haptic-only identification task. In the future we wish to investigate how congruent visual stimuli may affect perceived realism of the haptic shapes and possibly increase shape identification performance.

5.2 Rendering Stiffness

We applied insights into the discrimination of intensities discussed in section 4.2 to rendering stiffness of virtual elements through the variation of stimulus intensity proportionally to the virtual reaction force. It should be kept in mind that this approach to rendering stiffness differs from actual force feedback as one would obtain from interacting with a real stiff object or e.g. a grounded force feedback device, simply because there is no solid boundary resisting the user’s motion. Similar to the prior discussion on texture roughness rendering, this approach to rendering stiffness aims at creating a perceptually coherent evocation of stiffness rather than rendering it in a manner true to life. However, when asked whether the sensation delivered by the UMH device in the study discussed in Sec. 4.2 resembled stiffness, 80% of subjects responded that it did. The remaining 20% could not clearly describe the physical nature of the stimulus but still recognized an increase in force. When asked to describe the sensations they felt, answers ranged from “feeling a real piston” to feeling a “stream”, “circular air flows”, and “some kind of resistance”. It therefore appeared reasonable that this approach to rendering could elicit a perceptual experience which was coherent with the idea of virtual object stiffness.

The perceptual study conducted on intensity discrimination (see Sec. 4.2) showed that users could not only discriminate between three and five intensity levels within the perceptible range, but also correlated changes in perceived intensity during pressing of a virtual object with the object’s stiffness. This rendering approach was therefore implemented in a 4-part VR use-case (see Fig. 16). The first depicted a scene at a carnival fair featuring a pump for inflating virtual balloons (see Fig. 16-A). Users inflated the balloon by repeatedly pressing on the pump. A small STM circle of varying intensity was displayed at the contact location between the virtual piston’s handle and the user’s hand, with the intensity gradually increasing as the pressure inside the balloon increases, providing the user with the feeling of a pump with increasing stiffness.




Fig. 16. Four VR scenes exemplifying our approach to rendering stiffness of virtual objects: (A) Users inflate balloons at a carnival fair by pressing on a piston whose stiffness increases proportionally to the balloon’s inflation; (B) Users play a virtual piano. The action of pressing down on the keys is rendered using an STM stimulus, and the key’s stiffness (a function of its place in the register) affects the stimulus intensity; (C) Virtual medical palpation: users are instructed to find stiffer parts on a virtual patient’s abdomen based on the STM stimulus used to render the contact; (D) Virtual puzzle game requiring users to press down on identical-looking blocks in ascending order of stiffness to unlock a treasure. *Adapted from [50].*

The second use-case implemented a scenario similar to that presented by Hwang et al. [34], presenting a small 1-octave virtual piano whose octave could be lowered or raised using a pair of virtual buttons on the side (see Fig. 16-B). Real piano keys are generally weighted, with a higher key stiffness for the lower register and lower stiffness for the higher register. We render four octaves, each with a variable degree of stiffness rendered through a small STM circle projected on the user’s fingertip when it contacts a key.

The third use-case presents a proof of concept of an application to virtual medical palpation (see Fig. 16-C), where users can palpate a virtual character’s abdomen, sensing regions of different stiffness in a task where they are instructed to point out the stiffer regions than the surrounding tissue.

Finally, a last use case presents a virtual treasure-hunt game in a fantasy setting where users must press on four blocks in order of increasing stiffness in order to unlock the treasure (see Fig. 16-D). The blocks are visually identical, forcing the user to rely purely on the haptic feedback to solve the puzzle.

While functional, this approach is still somewhat limited by multiple factors. First, the usable intensity ranges for UMH stimuli is generally quite restricted. For example, with the Ultraleap Stratos Explore¹⁷ device used in our work, given the detection thresholds (see Sec. 4.2) and JNDs for stiffness (see Sec. 4.2), it is

¹⁷ <https://www.ultraleap.com/product/stratos-explore/>

expected that only up to five levels of stimulus intensity can be distinguished. In the future it may be possible to increase this resolution by using more powerful devices or possibly using other parameters such as e.g. STM draw frequency [19] to increase the stimulus's apparent intensity. An additional limitation comes from the fact that for stimuli whose base intensity was above 162.43 dB SPL, users no longer appeared to be able to discriminate intensity in a way that fit observations made at lower intensities as the perceived nature of the sensation changed from appearing as a small circle to something entirely different [50]. Further investigation into the acoustic phenomena occurring in this situation may provide solutions for overcoming this issue.

6 Conclusion and Perspectives

This Chapter presented a brief overview of research conducted at CNRS-IRISA on the topic of bringing ultrasound mid-air haptic (UMH) feedback to virtual reality interactions, within the context of the broader state of the art of UMH for 3D user interfaces both in augmented and in virtual reality.

We tackled three groups of challenges, relating to device integration, stimulus perception and rendering, with the aim of better understanding the contributions UMH can make to immersive VR experiences, as well as to propose and evaluate novel approaches for designing these interactions.

This work has yielded a novel robotic system aimed at enlarging a UMH interface's workspace, making it more compatible with the requirements of immersive VR [33]. It also led to the development of a software framework facilitating the design and evaluation of UMH stimuli [61] as well as numerous insights into UMH stimulus perception [31,50,60,61]. Finally, some novel rendering approaches for virtual object properties based on stimulus intensity variation were also proposed [31,50].

There is still much to be learned on the topic of integrating UMH into immersive VR. Fundamentally, the absence of physical contact in UMH interaction as well as the characteristics of UMH stimuli provide both limitations that need to be worked around and opportunities for designing unique immersive VR interaction techniques.

We plan to extend our work on adapting the workspace of UMH devices to VR interactions, through hard- and software solutions that will enable e.g. novel bi-manual and fine manipulations with UMH feedback. There are also many exciting avenues for exploration on the use of UMH with concurrent haptic feedback, and manipulation combining UMH with tangible objects. Approaches of this type could allow to complement UMH stimulus properties and overcome the technology's limitations, however they pose major challenges both in terms of device integration and rendering. Finally, there is also a need to investigate solutions to effectively overcome issues associated with the physical footprint of UMH devices, through e.g. approaches to collision avoidance in VR and visual occlusion in AR. Many of the technical aspects of rendering and integration

touched upon in this chapter will become clearer in the upcoming chapters focused on the theory, technological implementation and safety of UMH.

On the topic of UMH stimulus perception, there is still much to be learned from the psychophysical study of UMH stimuli. Even more importantly, it will be necessary to study cross-modal interactions between UMH and other haptics, as well as between UMH, visual and auditory feedback, as immersive VR interaction is inherently multi-modal. The next chapter (chapter “Multimodal Interaction with Mid-Air Haptics”) will tackle this aspect in more detail.

These insights will lead to the development of many novel approaches and application scenarios for UMH rendering. Furthermore, there is an urgent need for investigating novel interaction techniques and metaphors to leverage the unique properties and compensate for the limitations of UMH technology.

The development of more advanced VR interaction techniques leveraging UMH will open up perspectives for studying the impact of UMH on virtual interaction performance, user immersion and presence, as well as embodiment towards virtual avatars and factors relating to user engagement and affect.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 programme under grant agreement No 801413; project “H-Reality”.

References

1. Arafsha, F., Zhang, L., Dong, H., El Saddik, A.: Contactless haptic feedback: State of the art. In: 2015 IEEE International Symposium on Haptic, Audio and Visual Environments and Games (HAVE). pp. 1–6. IEEE (2015)
2. Ariga, K., Fujiwara, M., Makino, Y., Shinoda, H.: Midair haptic presentation using concave reflector. In: International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. pp. 307–315. Springer (2020)
3. Balint, P., Althoefer, K.: Medical virtual reality palpation training using ultrasound based haptics and image processing. *Proc. Jt. Work. New Technol. Comput. Assist. Surg* (2018)
4. Barreiro, H., Sinclair, S., Otaduy, M.A.: Path routing optimization for stm ultrasound rendering. *IEEE transactions on haptics* **13**(1), 45–51 (2020)
5. Beattie, D., Frier, W., Georgiou, O., Long, B., Ablart, D.: Incorporating the perception of visual roughness into the design of mid-air haptic textures. In: ACM Symposium on Applied Perception 2020. pp. 1–10 (2020)
6. Van den Bogaert, L., Geerts, D.: User-defined mid-air haptic sensations for interacting with an ar menu environment. In: International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. pp. 25–32. Springer (2020)
7. Brice, D., McRoberts, T., Rafferty, K.: 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. pp. 120–137. Springer (2019)

8. Carter, T., Seah, S.A., Long, B., Drinkwater, B., Subramanian, S.: Ultrahaptics: multi-point mid-air haptic feedback for touch surfaces. In: Proceedings of the 26th annual ACM symposium on User interface software and technology. pp. 505–514 (2013)
9. Carvalheiro, C., Nóbrega, R., da Silva, H., Rodrigues, R.: User redirection and direct haptics in virtual environments. In: Proceedings of the 24th ACM international conference on Multimedia. pp. 1146–1155 (2016)
10. Chatzopoulos, D., Bermejo, C., Huang, Z., Hui, P.: Mobile augmented reality survey: From where we are to where we go. *Ieee Access* **5**, 6917–6950 (2017)
11. Cirio, G., Marchal, M., Otaduy, M.A., Lécuyer, A.: Six-dof haptic interaction with fluids, solids, and their transitions. In: 2013 World Haptics Conference (WHC). pp. 157–162. IEEE (2013)
12. Corenthy, L., Giordano, M., Hayden, R., Griffiths, D., Jeffrey, C., Limerick, H., Georgiou, O., Carter, T., Müller, J., Subramanian, S.: Touchless tactile displays for digital signage: mid-air haptics meets large screens. In: Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems. pp. 1–4 (2018)
13. Cornelio Martinez, P.I., De Pirro, S., Vi, C.T., Subramanian, S.: Agency in mid-air interfaces. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. pp. 2426–2439 (2017)
14. Di Luca, M., Mahnan, A.: Perceptual limits of visual-haptic simultaneity in virtual reality interactions. In: 2019 IEEE World Haptics Conference (WHC). pp. 67–72. IEEE (2019)
15. Dzidek, B., Frier, W., Harwood, A., Hayden, R.: Design and evaluation of mid-air haptic interactions in an augmented reality environment. In: International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. pp. 489–499. Springer (2018)
16. Eid, M.A., Al Osman, H.: Affective haptics: Current research and future directions. *IEEE Access* **4**, 26–40 (2015)
17. Evangelou, G., Limerick, H., Moore, J.W.: I feel it in my fingers! sense of agency with mid-air haptics. *IEEE World Haptics* (2021)
18. Freeman, E., Anderson, R., Williamson, J., Wilson, G., Brewster, S.A.: Textured surfaces for ultrasound haptic displays. In: Proceedings of the 19th ACM International Conference on Multimodal Interaction. pp. 491–492 (2017)
19. Frier, W., Ablart, D., Chilles, J., Long, B., Giordano, M., Obrist, M., Subramanian, S.: Using spatiotemporal modulation to draw tactile patterns in mid-air. In: International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. pp. 270–281. Springer (2018)
20. Frier, W., Pittera, D., Ablart, D., Obrist, M., Subramanian, S.: Sampling strategy for ultrasonic mid-air haptics. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–11 (2019)
21. Frutos-Pascual, M., Harrison, J.M., Creed, C., Williams, I.: Evaluation of ultrasound haptics as a supplementary feedback cue for grasping in virtual environments. In: 2019 International Conference on Multimodal Interaction. pp. 310–318 (2019)
22. Georgiou, O., Biscione, V., Harwood, A., Griffiths, D., Giordano, M., Long, B., Carter, T.: Haptic in-vehicle gesture controls. In: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications Adjunct. pp. 233–238 (2017)

23. Georgiou, O., Jeffrey, C., Chen, Z., Tong, B.X., Chan, S.H., Yang, B., Harwood, A., Carter, T.: Touchless haptic feedback for vr rhythm games. In: 2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). pp. 553–554. IEEE (2018)
24. Gil, H., Son, H., Kim, J.R., Oakley, I.: Whiskers: Exploring the use of ultrasonic haptic cues on the face. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. pp. 1–13 (2018)
25. Hajas, D., Pittera, D., Nasce, A., Georgiou, O., Obrist, M.: Mid-air haptic rendering of 2d geometric shapes with a dynamic tactile pointer. *IEEE transactions on haptics* **13**(4), 806–817 (2020)
26. Hasegawa, K., Qiu, L., Noda, A., Inoue, S., Shinoda, H.: Electronically steerable ultrasound-driven long narrow air stream. *Applied Physics Letters* **111**(6), 064104 (2017)
27. Hoshi, T.: Gradual phase shift to suppress noise from airborne ultrasound tactile display. In: Proc. ACM CHI Workshop Mid-Air Haptics Displays: Syst. Un-instrumented Mid-Air Interact., Session 2: Provide Vis. Haptic Feedback (2016)
28. Hoshi, T., Abe, D., Shinoda, H.: Adding tactile reaction to hologram. In: RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication. pp. 7–11. IEEE (2009)
29. Hoshi, T., Iwamoto, T., Shinoda, H.: Non-contact tactile sensation synthesized by ultrasound transducers. In: World Haptics 2009-Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. pp. 256–260. IEEE (2009)
30. Hoshi, T., Takahashi, M., Iwamoto, T., Shinoda, H.: Noncontact tactile display based on radiation pressure of airborne ultrasound. *IEEE Transactions on Haptics* **3**(3), 155–165 (2010)
31. Howard, T., Gallagher, G., Lécuyer, A., Pacchierotti, C., Marchal, M.: Investigating the recognition of local shapes using mid-air ultrasound haptics. In: 2019 IEEE World Haptics Conference (WHC). pp. 503–508. IEEE (2019)
32. Howard, T., Gicquel, G., Marchal, M., Lécuyer, A., Pacchierotti, C.: Pumah: Pan-tilt ultrasound mid-air haptics. In: WHC 2021-IEEE World Haptics Conference (2021)
33. Howard, T., Marchal, M., Lécuyer, A., Pacchierotti, C.: Pumah: Pan-tilt ultrasound mid-air haptics for larger interaction workspace in virtual reality. *IEEE transactions on haptics* **13**(1), 38–44 (2019)
34. Hwang, I., Son, H., Kim, J.R.: Airpiano: Enhancing music playing experience in virtual reality with mid-air haptic feedback. In: 2017 IEEE World Haptics Conference (WHC). pp. 213–218. IEEE (2017)
35. Inoue, S., Kobayashi-Kirschvink, K.J., Monnai, Y., Hasegawa, K., Makino, Y., Shinoda, H.: Horn: the hapt-optic reconstruction. In: ACM SIGGRAPH 2014 Emerging Technologies. pp. 1–1 (2014)
36. Inoue, S., Makino, Y., Shinoda, H.: Active touch perception produced by airborne ultrasonic haptic hologram. In: 2015 IEEE World Haptics Conference (WHC). pp. 362–367. IEEE (2015)
37. Ito, M., Kokumai, Y., Shinoda, H.: Midair click of dual-layer haptic button. In: 2019 IEEE World Haptics Conference (WHC). pp. 349–352. IEEE (2019)
38. Jang, J., Park, J.: Sph fluid tactile rendering for ultrasonic mid-air haptics. *IEEE transactions on haptics* **13**(1), 116–122 (2020)
39. Kamigaki, T., Suzuki, S., Shinoda, H.: Noncontact thermal and vibrotactile display using focused airborne ultrasound. In: International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. pp. 271–278. Springer (2020)

40. Kappus, B., Long, B.: Spatiotemporal modulation for mid-air haptic feedback from an ultrasonic phased array. *The Journal of the Acoustical Society of America* **143**(3), 1836–1836 (2018)
41. Karafotias, G., Korres, G., Teranishi, A., Park, W., Eid, M.: Mid-air tactile stimulation for pain distraction. *IEEE transactions on haptics* **11**(2), 185–191 (2017)
42. Kervégant, C., Raymond, F., Graeff, D., Castet, J.: Touch hologram in mid-air. In: *ACM SIGGRAPH 2017 Emerging Technologies*, pp. 1–2 (2017)
43. Kimura, Y., Makino, Y., Shinoda, H.: Computer-created interactive 3d image with midair haptic feedback. In: *International AsiaHaptics conference*. pp. 491–494. Springer (2016)
44. Korres, G., Aujeszky, T., Eid, M.: Characterizing tactile rendering parameters for ultrasound based stimulation. In: *2017 IEEE World Haptics Conference (WHC)*. pp. 293–298. IEEE (2017)
45. Korres, G., Eid, M.: Haptogram: Ultrasonic point-cloud tactile stimulation. *IEEE Access* **4**, 7758–7769 (2016)
46. Kreimeier, J., Hammer, S., Friedmann, D., Karg, P., Bühner, C., Bankel, L., Götzelmann, T.: Evaluation of different types of haptic feedback influencing the task-based presence and performance in virtual reality. In: *Proceedings of the 12th ACM International Conference on Pervasive Technologies Related to Assistive Environments*. pp. 289–298 (2019)
47. Long, B., Seah, S.A., Carter, T., Subramanian, S.: Rendering volumetric haptic shapes in mid-air using ultrasound. *ACM Transactions on Graphics (TOG)* **33**(6), 1–10 (2014)
48. Lugo-Villeda, L.I., Frisoli, A., Pabon, S., Padilla, M.A., Sotgiu, E., Bergamasco, M.: Light-exoskeleton and data-glove integration for enhancing virtual reality applications. In: *2009 International Conference on Advanced Robotics*. pp. 1–6. IEEE (2009)
49. Makino, Y., Furuyama, Y., Inoue, S., Shinoda, H.: Haptoclone (haptic-optical clone) for mutual tele-environment by real-time 3d image transfer with midair force feedback. In: *CHI*. pp. 1980–1990 (2016)
50. Marchal, M., Gallagher, G., Lécuyer, A., Pacchierotti, C.: 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*. pp. 297–306. Springer (2020)
51. Martinez, J., Griffiths, D., Biscione, V., Georgiou, O., Carter, T.: Touchless haptic feedback for supernatural vr experiences. In: *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. pp. 629–630. IEEE (2018)
52. Martinez, J., Harwood, A., Limerick, H., Clark, R., Georgiou, O.: Mid-air haptic algorithms for rendering 3d shapes. In: *2019 IEEE International Symposium on Haptic, Audio and Visual Environments and Games (HAVE)*. pp. 1–6. IEEE (2019)
53. Matsubayashi, A., Makino, Y., Shinoda, H.: Direct finger manipulation of 3d object image with ultrasound haptic feedback. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. pp. 1–11 (2019)
54. Meli, L., Pacchierotti, C., Salvietti, G., Chinello, F., Maisto, M., De Luca, A., Prattichizzo, D.: Combining wearable finger haptics and augmented reality: User evaluation using an external camera and the microsoft hololens. *IEEE Robotics and Automation Letters* **3**(4), 4297–4304 (2018)
55. Meli, L., Scheggi, S., Pacchierotti, C., Prattichizzo, D.: Wearable haptics and hand tracking via an rgb-d camera for immersive tactile experiences. In: *ACM SIGGRAPH 2014 Posters*, pp. 1–1 (2014)

56. Mercado, V., Howard, T., Si-Mohammed, H., Argelaguet, F., Lécuyer, A.: Alfred: the haptic butler - on-demand tangibles for object manipulation in virtual reality using an ethd. *IEEE World Haptics* (2021)
57. Mestre, D., Fuchs, P., Berthoz, A., Vercher, J.: *Immersion et présence. Le traité de la réalité virtuelle*. Paris: Ecole des Mines de Paris pp. 309–38 (2006)
58. Monnai, Y., Hasegawa, K., Fujiwara, M., Yoshino, K., Inoue, S., Shinoda, H.: Haptomime: mid-air haptic interaction with a floating virtual screen. In: *Proceedings of the 27th annual ACM symposium on User interface software and technology*. pp. 663–667 (2014)
59. Monnai, Y., Hasegawa, K., Fujiwara, M., Yoshino, K., Inoue, S., Shinoda, H.: Adding texture to aerial images using ultrasounds. In: *Haptic Interaction*, pp. 59–61. Springer (2015)
60. Mulot, L., Gicquel, G., Frier, W., Marchal, M., Pacchierotti, C., Howard, T.: Curvature discrimination for dynamic ultrasound mid-air haptic stimuli. In: *WHC 2021-IEEE World Haptics Conference* (2021)
61. Mulot, L., Gicquel, G., Zanini, Q., Frier, W., Marchal, M., Pacchierotti, C., Howard, T.: DOLPHIN: A Framework for the Design and Perceptual Evaluation of Ultrasound Mid-Air Haptic Stimuli. In: *ACM Symposium on Applied Perception* (2021)
62. Obrist, M., Seah, S.A., Subramanian, S.: Talking about tactile experiences. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. pp. 1659–1668 (2013)
63. Obrist, M., Subramanian, S., Gatti, E., Long, B., Carter, T.: Emotions mediated through mid-air haptics. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. pp. 2053–2062 (2015)
64. Ochiai, Y., Hoshi, T., Suzuki, I.: Holographic whisper: Rendering audible sound spots in three-dimensional space by focusing ultrasonic waves. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. pp. 4314–4325 (2017)
65. Pittera, D., Ablart, D., Obrist, M.: Creating an illusion of movement between the hands using mid-air touch. *IEEE transactions on haptics* **12**(4), 615–623 (2019)
66. Pittera, D., Gatti, E., Obrist, M.: 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*. pp. 1–15 (2019)
67. Popescu, G.V., Burdea, G.C., Trefftz, H.: Multimodal interaction modeling. In: *Handbook of Virtual Environments*, pp. 475–494. CRC Press (2002)
68. Price, A., Long, B.: Fibonacci spiral arranged ultrasound phased array for mid-air haptics. In: *2018 IEEE International Ultrasonics Symposium (IUS)*. pp. 1–4. IEEE (2018)
69. Rakkolainen, I., Freeman, E., Sand, A., Raisamo, R., Brewster, S.: A survey of mid-air ultrasound haptics and its applications. *IEEE Transactions on Haptics* **14**(1), 2–19 (2020)
70. Ramsamy, P., Haffegge, A., Jamieson, R., Alexandrov, V.: Using haptics to improve immersion in virtual environments. In: *International Conference on Computational Science*. pp. 603–609. Springer (2006)
71. Romanus, T., Frish, S., Maksymenko, M., Frier, W., Corenthy, L., Georgiou, O.: Mid-air haptic bio-holograms in mixed reality. In: *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. pp. 348–352. IEEE (2019)

72. Rümelin, S., Gabler, T., Bellenbaum, J.: Clicks are in the air: how to support the interaction with floating objects through ultrasonic feedback. In: *Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. pp. 103–108 (2017)
73. Rutten, E., Van Den Bogaert, L., Geerts, D.: From initial encounter with mid-air haptic feedback to repeated use: the role of the novelty effect in user experience. *IEEE Transactions on Haptics* (2020)
74. Rutten, I., Frier, W., Van den Bogaert, L., Geerts, D.: Invisible touch: How identifiable are mid-air haptic shapes? In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. pp. 1–6 (2019)
75. Saint-Aubert, J., Regnier, S., Haliyo, S.: Cable driven haptic interface for co-localized desktop vr. In: *2018 IEEE Haptics Symposium (HAPTICS)*. pp. 351–356. IEEE (2018)
76. Salazar, S.V., Pacchierotti, C., de Tinguy, X., Maciel, A., Marchal, M.: Altering the stiffness, friction, and shape perception of tangible objects in virtual reality using wearable haptics. *IEEE transactions on haptics* **13**(1), 167–174 (2020)
77. Sand, A., Rakkolainen, I., Isokoski, P., Kangas, J., Raisamo, R., Palovuori, K.: Head-mounted display with mid-air tactile feedback. In: *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology*. pp. 51–58 (2015)
78. Shakeri, G., Freeman, E., Frier, W., Iodice, M., Long, B., Georgiou, O., Andersson, C.: Three-in-one: Levitation, parametric audio, and mid-air haptic feedback. In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. pp. 1–4 (2019)
79. Suzuki, S., Takahashi, R., Nakajima, M., Hasegawa, K., Makino, Y., Shinoda, H.: Midair haptic display to human upper body. In: *2018 57th Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE)*. pp. 848–853. IEEE (2018)
80. de Tinguy, X., Howard, T., Pacchierotti, C., Marchal, M., Lécuyer, A.: Weatavix: wearable actuated tangibles for virtual reality experiences. In: *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. pp. 262–270. Springer (2020)
81. de Tinguy, X., Pacchierotti, C., Marchal, M., Lécuyer, A.: Toward universal tangible objects: Optimizing haptic pinching sensations in 3d interaction. In: *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. pp. 321–330. IEEE (2019)
82. Vi, C.T., Ablart, D., Gatti, E., Velasco, C., Obrist, M.: Not just seeing, but also feeling art: Mid-air haptic experiences integrated in a multisensory art exhibition. *International Journal of Human-Computer Studies* **108**, 1–14 (2017)
83. Waltemate, T., Senna, I., Hülsmann, F., Rohde, M., Kopp, S., Ernst, M., Botsch, M.: The impact of latency on perceptual judgments and motor performance in closed-loop interaction in virtual reality. In: *Proceedings of the 22nd ACM conference on virtual reality software and technology*. pp. 27–35 (2016)
84. Wilson, G., Carter, T., Subramanian, S., Brewster, S.A.: Perception of ultrasonic haptic feedback on the hand: localisation and apparent motion. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. pp. 1133–1142 (2014)
85. Yoshida, K., Horiuchi, Y., Inoue, S., Makino, Y., Shinoda, H.: Haptoclonear: mutual haptic-optic interactive system with 2d image superimpose. In: *ACM SIGGRAPH 2017 Emerging Technologies*, pp. 1–2 (2017)

86. Yoshino, K., Shinoda, H.: Visio-acoustic screen for contactless touch interface with tactile sensation. In: 2013 World Haptics Conference (WHC). pp. 419–423. IEEE (2013)