



Does Multi-Actuator Vibrotactile Feedback Within Tangible Objects Enrich VR Manipulation?

Pierre-Antoine Cabaret, Thomas Howard, Guillaume Gicquel, Claudio Pacchierotti, Marie Babel, Maud Marchal

► To cite this version:

Pierre-Antoine Cabaret, Thomas Howard, Guillaume Gicquel, Claudio Pacchierotti, Marie Babel, et al.. Does Multi-Actuator Vibrotactile Feedback Within Tangible Objects Enrich VR Manipulation?. IEEE Transactions on Visualization and Computer Graphics, 2023, pp.1 - 13. 10.1109/tvcg.2023.3279398 . hal-04109929

HAL Id: hal-04109929

<https://inria.hal.science/hal-04109929>

Submitted on 30 May 2023

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.

Does multi-actuator vibrotactile feedback within tangible objects enrich VR manipulation?

Pierre-Antoine Cabaret[†], Thomas Howard[†], Guillaume Gicquel,
Claudio Pacchierotti, Marie Babel and Maud Marchal

Abstract—Rich, informative and realistic haptic feedback is key to enhancing Virtual Reality (VR) manipulation. Tangible objects provide convincing grasping and manipulation interactions with haptic feedback of e.g., shape, mass and texture properties. But these properties are static, and cannot respond to interactions in the virtual environment. On the other hand, vibrotactile feedback provides the opportunity for delivering dynamic cues rendering many different contact properties, such as impacts, object vibrations or textures. Handheld objects or controllers in VR are usually restricted to vibrating in a monolithic fashion. In this paper, we investigate how spatializing vibrotactile cues within handheld tangibles could enable a wider range of sensations and interactions. We conduct a set of perception studies, investigating the extent to which spatialization of vibrotactile feedback within tangible objects is possible as well as the benefits of proposed rendering schemes leveraging multiple actuators in VR. Results show that vibrotactile cues from localized actuators can be discriminated and are beneficial for certain rendering schemes.

Index Terms—Haptics, Vibrotactile, Tangible, VR, Manipulation

1 INTRODUCTION

IN Virtual Reality (VR), haptic feedback is desirable to explore, grasp, and manipulate virtual objects. It has been shown to improve user performance, perceived realism [1], and user immersion [2]. Haptic feedback for VR is often focused on single specific interactions or sensations [1], and the design of devices capable of rendering rich and realistic touch in a variety of interactions is a relatively recent development in the field [3]. In this paper, we investigate the use of mixed localized vibrotactile and tangible haptic feedback for object manipulation interactions in VR (see Figure 1).

Without some form of actuation, tangible props are limited to static physical properties and thus can neither render a wide range of virtual object physical properties, nor properties of contact with other virtual elements felt through a grasped virtual object or tool.

Conversely, vibrotactile feedback (VF), can be used to render a range of sensations going from simple contact [4] to texture properties [5]. However, this type of haptic feedback suffers from two major drawbacks. The first one is that vibrotactile actuators (hereafter referred to as vibrotactors) are usually mounted in permanent contact with the user's skin, introducing unwanted sensations of contact independent of the virtual contacts. The second is that while tactile feedback devices are capable of providing cues indicating the onset of virtual contacts, they put up no physical resistance to user motions, introducing potential mismatches between a user's tactile and kinesthetic perception during interaction [6]. This paper explores the possibilities for enhancing VR

manipulation by incorporating vibrotactile feedback into tangible objects (see Figure 1). Our contributions are:

- Three perception studies exploring the feasibility and benefit of localized vibrotactile feedback within a tangible object (section 3).
- The design of mixed haptic rendering schemes for virtual object and contact properties (section 4), validated by a set of human subject studies (section 5) and demonstrated by use-cases in VR (subsection 5.5). These properties cover contact, impacts, macroscopic texture and roughness, object contents, and the sensation of handling vibrating or pulsating objects.

2 RELATED WORK

2.1 Mixed haptics for VR interaction

There has been a recent push towards mixing multiple haptic actuators into single devices to support VR interaction [3]. This development is driven by the fact that rich haptic interaction relies on conveying multiple physical properties of virtual objects, while human haptic perception relies on multiple interconnected sub-modalities each reliant on specific cutaneous and kinesthetic receptors [7].

In this context, vibrotactile feedback has been combined with grounded force-feedback (GFF) [8] and handheld force-feedback [9], [10] devices, to render properties such as stiffness, hardness, and surface texture. Vibrotactors have also been combined with other tactile actuators delivering skin-stretch [9], [11], [12], [13], squeeze [11], [13], [14], thermal [15], [16], [17] and electrotactile [12], [15] cues. These combinations have been investigated in grounded tactile interfaces [16], handheld tactile and force-feedback devices [9], [10], [18] and wearables for the hand [12], [15],

- Pierre-Antoine Cabaret, Thomas Howard, Marie Babel and Maud Marchal are with Univ Rennes, INSA, IRISA, Inria, CNRS – Rennes, France.
- Maud Marchal is also with IUF, France.
- Guillaume Gicquel and Claudio Pacchierotti are with CNRS, Univ Rennes, Inria, IRISA – Rennes, France.
- [†]Both authors contributed equally to this work.



Fig. 1. This paper shows that tangible objects with multiple built-in vibrotactile actuators (A) can effectively render localized vibrotactile cues. These can be used to add spatial information to vibrotactile feedback of interactions with the virtual environment (VE), such as, e.g., feeling textures through a handheld tool (B), feeling contents of a manipulated object (C), or sensing the location of virtual impacts between a virtual handheld object and the VE (D). This opens many perspectives for augmenting passive haptic manipulation in VR.

limbs [11], [13], torso [14] or the head when built into the VR head-mounted display (HMD) [17].

To date, only very few systems which allow making and breaking of contact and grasping of 3D objects in VR provide added VF, and there have been no investigations into such mixed haptic rendering in VR.

2.2 Tangible haptics for object manipulation in VR

Rendering making and breaking of contact is a key feature in enabling haptic exploration and direct manipulation of virtual objects. Arguably, the most effective way to achieve this is by a haptic interface actually making and breaking contact with a user's skin. This can be done by using tangible props representing virtual objects (i.e. passive haptics).

Passive haptics relies on superimposing virtual objects with standalone tangible objects (e.g., [19], [20]). Without adding much complexity to the system, tangibles enable natural object manipulations in VR and provide shape and other physical sensations (e.g., weight, stiffness). They only require adequate tracking of the props [21] or prior calibration and VE modelling [22] combined with tracking of the user. Yet, passive haptics can become unwieldy for representing complex VEs because each virtual object requires its own tangible counterpart. Furthermore, passive haptics cannot adapt to virtual objects which evolve within a VE, or to changing VEs [23].

To counter these limitations, internally actuated tangible props have been investigated for delivering ungrounded force feedback in VR manipulation (e.g., [24], [25], [26]). Beyond this, tangibles have also been combined with GFF devices [27] and ETHDs [28], [29], [30]. These hybrid solutions limit the number of distinct tangible objects required to represent a VE while simultaneously allowing the tangibles' haptic properties to interactively evolve.

2.3 Mixing vibrotactile feedback and tangible haptics

Incorporating vibrotactors into tangible objects to widen the scope of interaction is a concept that has extensively been explored e.g., in tangible user interfaces (TUI) [31], [32] but less so in VR. Here, similarly to recent work by Cabaret et al. [33], we propose to apply this idea to VR and to explore potential benefits to interaction. Vibrotactile feedback in graspable devices in VR has thus far been limited to monolithic vibrotactile feedback (i.e. the entire object vibrates [34]) delivered by permanently held end-effectors in GFF devices [35] or permanently handheld controllers or props [18], [36]. However, localized vibrotactile feedback

has been shown effective in conveying spatial information when used in wearable [34], [37], [38] and surface [39] haptics. Cabaret et al. embedded vibrotactors in a tangible object in VR, focusing on VF for virtual impacts occurring outside the hand [33]. This motivates our investigation into additional possibilities provided by localized VF to modulate the perceived properties of freely manipulated tangible objects.

2.4 Vibrotactile rendering

The simplest application of vibrotactile feedback lies in rendering vibrating virtual objects (e.g., [35]). Vibrotactors have also been successfully used to render contacts with virtual objects [4], [40]. As a consequence, they can indirectly render shape and size, when cues are triggered by onset of contacts between the virtual hand and object (e.g., [41]). VF delivered in this manner was shown to improve user presence and task performance in VR object manipulation and identification tasks [42].

Impacts with virtual objects [38] are usually rendered in a model-based approach, using exponentially decaying sinusoidal vibrations [43], [44] whose parameters can be adjusted to render different virtual object materials [44]. Garcia-Valle et al. demonstrated that VF of collisions enhances the realism of VR serious games [45]. Interaction forces can be rendered by VF by varying a vibration parameter (usually amplitude) proportionally to the reaction forces encountered in a VE [46], [47]. Vibrotactile force feedback has also been applied to render virtual object stiffness. Kildal et al. simulated varying stiffness of a TUI's prop pinched between the fingers through integrated VF [31]. A similar principle was applied in VR by Maereg et al. using finger-mounted wearables [48]. In an alternative approach, asymmetric vibrations have been shown to induce illusions of directional forces acting on the skin (e.g., [49]).

VF is a prime candidate for rendering textures. Texture can be categorized into macro texture (i.e. bumps, surface waviness) and micro texture (roughness) [50]. Macro and micro texture are thought to be primarily mediated by spatial and vibrational cues respectively [51]. Roughness rendering interlinks with the concept of surface friction rendering because the phenomena are physically linked, although perceptually these properties are often considered independent due to the importance of lateral skin stretch in friction perception [50], [52]. Regardless, vibrotactile cues are key components of haptic perception for both roughness [53], [54] and to some extent friction [50]. VF can thus effectively

be applied for both. Macro texture is usually rendered as a series of impacts depending on exploration velocity, texture geometry and applied force [44]. Micro texture rendering has been attempted with model-based approaches [10], [44], [55] as well as data-driven approaches [5], [56]. Model-based approaches are often computationally intensive, and thus require simplifications that compromise realism or simply do not function in real-time VR [8]. Among data-driven approaches, the Penn Haptic Texture Toolkit [5] is an open-source model for procedurally generating VF. We incorporated it into our system and discuss it in subsection 4.4.

Some of the vibrotactile rendering schemes discussed above have already been implemented in immersive VR [38], [48]. However, most vibrotactile rendering schemes from the literature have yet to be demonstrated in this setting. Also, a majority of the work on vibrotactile rendering presented here was conducted with VF as the only source of haptic information, and presented in a monolithic manner.

3 LOCALIZED VIBROTACTILE FEEDBACK

To determine whether localized vibrotactile feedback is a functional and useful complement to tangible haptics in VR manipulations, we begin by investigating how spatialized VF cues within a handheld tangible object are perceived. Here, we detail our prototype design, followed by three user studies assessing whether multiple vibrotactors can provide both monolithic and localized VF, discrimination performance between different localized VF cues, and finally localization performance for VF cues.

3.1 Design

3.1.1 System design

For our prototype and investigations, we used spherical tangible objects into which we embedded multiple voice-coil actuators (see Figure 2). We empirically determined that a 70mm sphere diameter allowed consistent grasping across many different hand morphologies, and thus chose this diameter for our tangible spheres. Our analyses looked into the effects of hand sizes in each of the presented experiments but no significant effects were found.

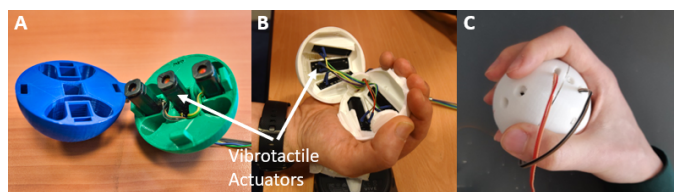


Fig. 2. (A) Prototype tangible sphere containing 3 actuators: the lateral actuators enable localized VF while the central actuator enables monolithic VF for comparison. (B) Prototype tangible sphere containing 5 actuators arranged so that each is located directly beneath a fingertip during grasping. (C) Close-up of a user holding a vibrotactile tangible.

The actuators integrated into the device were Actronika's HapCoil-One voice coil actuators. We chose these actuators despite their rather large form factor and mass because they are able to generate high-bandwidth and arbitrary waveforms with significant intensities, making them ideal for rendering a wide range of sensations [34].

The actuators were driven using audio signals from a USB 7.1 surround sound card (see Figure 3-B). The voltage

from the sound card's audio output was insufficient to drive the vibrotactors, so we used TPA3116D2 dual-channel amplifiers (see Figure 3-B) powered by external 5V power supply units to drive the actuators in pairs. A calibration step was performed to adjust the amplifier gains in order to ensure they produced an equal voltage output amplitude for a given input signal from the sound card. To do this, we connected each amplifier output to an oscilloscope and played back an identical signal on loop, adjusting the gain until the target output acceleration of 2.5g was reached. We assumed that since all actuators used are identical, they also should have a similar response. These values were empirically found to be perceptible, differentiable and comfortable.

All VF stimuli for the experiments were generated with Syntacts [57], an open-source software created specifically for audio-based haptics. This software offers a library of waveforms and modifiers to design haptic cues and provides precise control over latency. A Unity package is available for Syntacts, simplifying its integration within our system. Since this approach is not capable of generating procedural signals from scripts, we used the Unity audio engine for interactions requiring procedural signal generation, i.e. rendering texture roughness (see subsection 4.4).

3.1.2 Experimental setup design

Our first two experiments were dedicated to the use of a pair of localized vibrotactors, as this constitutes the basic building block for multi-actuator spatialization (see Figure 2-A). The third experiment used a sphere fitted with one actuator under each finger, for 5 actuators in total (see Figure 2-B).

All experiments were conducted in VR. The user viewed the VE simulated in Unity3D through an HTC Vive VR headset (see Figure 3-A). A tracker was placed either on the back of the hand or on the forearm, thus leaving the volar face of the hand free of any obstruction or unwanted haptic stimuli. Information from the tracker served to reproduce movements of the user's hand onto the virtual hand while simultaneously providing information about the physical location of the tangible object. During all three experiments, participants used a Vive controller held in their non-dominant hand to answer experimental questions (see

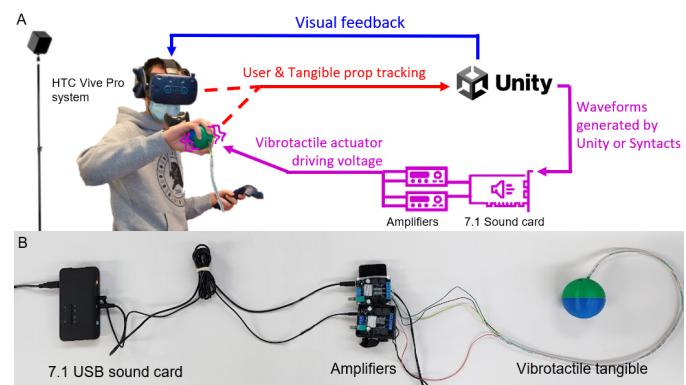


Fig. 3. (A) The user's hand and physical prop are simultaneously tracked using an HTC Vive Tracker. The VE is simulated in Unity3D and rendered to the user's HTC Vive Pro HMD. The simulation outputs vibration patterns for the vibrotactors as audio signals which run through a custom-built amplification stage to drive the actuators. (B) Close-up of the sound card, amplification stage and an actuated tangible sphere.

Figure 4). White noise was played throughout all experiments to mask any sound from the actuators that could bias subject responses. Written informed consent was provided by all participants prior to each experiment. Approval of experimental procedures and protocols was granted by Inria.

3.2 Effectiveness of multi-actuator VF

3.2.1 Research question and hypotheses

To determine whether spatializing VF within a tangible object using multiple actuators is possible, we performed an initial experiment using the sphere fitted with 3 vibrotactors (see Figure 2-A). Specifically, we sought to determine whether in this simplest implementation, localized cues are discriminable from one another, as well as from monolithic vibrations of the tangible. A well-documented phenomenon in vibrotactile feedback using multiple actuators is the phantom sensation [58] or “funneling illusion” [59] wherein two distinct points of vibrotactile stimulation can be perceived as a single vibrotactile stimulus presented at an intermediary location. Because of this, we also aimed to assess whether monolithic single-actuator vibrations of the tangible object could be replaced by multiple localized vibrotactors symmetrically placed around the object and providing simultaneous stimuli at equal amplitudes. Finally, we sought to assess to what extent results are dependent on the stimulus waveform.

Our hypotheses were the following: **(H1)** Localized cues (played on P or T actuators) are distinguishable from one another, regardless of the cue waveform; **(H2)** Localized cues are distinguishable from perceptually amplitude-matched monolithic cues (played on the M actuator), regardless of the cue waveform; **(H3)** Monolithic cues can be simulated by simultaneously triggering symmetrically arranged localized actuators with equal-amplitude cues, regardless of the cue waveform.

3.2.2 Materials and methods

To investigate these questions, we performed a user study involving 16 participants (8F, 8M, all right-handed, ages 21-52 ($M=26.9$, $SD=8.8$)), following an oddball procedure. Participants were recruited in the lab, 12 of them having some experience with haptics. Using the apparatus described in subsection 3.1 above, participants held the spherical tangible in their dominant hand such that the lateral actuators were located respectively between the thumb and index (actuator

location T) and below the pinkie (actuator location P). Visible markings on the sphere allowed the experimenter to ensure that the subject's hand was correctly positioned.

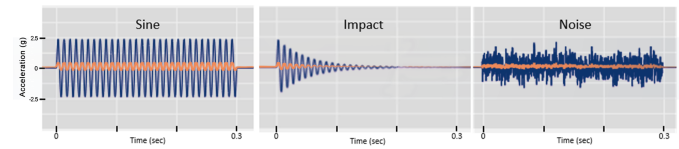


Fig. 5. Waveforms used in the experiment. Weak (0.25g) and Strong (2.5g) levels used in subsection 3.4 are shown in orange and blue respectively.

Stimuli were one of the three representative waveforms (see also Figure 5):

- **Impact (I)**: a 150ms-long 100Hz exponentially decaying sinusoid (see e.g., subsection 4.2 on impact rendering for an example of use),
- **Noise (N)**: 300ms of white noise (see e.g., subsection 4.4 on texture rendering),
- **Sine (S)**: 300ms of a pure sine wave (see e.g., subsection 4.5 on rendering liquid contents).

We considered four stimulus locations (see Figure 4) :

- **Thumb (T)**: Only the vibrotactor on the thumb side of the hand (localized)
- **Pinkie (P)**: Only the vibrotactor on the pinkie side of the hand (localized)
- **Monolithic (M)**: Only the vibrotactor at the center of the object (monolithic)
- **“Pseudo-monolithic” (PM)**: Both thumb- and pinkie-side actuators simultaneously (i.e. monolithic simulated with localized actuators).

Each trial consisted of a sequence of three stimuli of the same waveform, two of which had identical locations and one which had a different location from the others (the oddball, randomly positioned in the sequence). Stimuli were presented in sequence with a 1s pause between them, after which the subject was asked which stimulus felt most different. Using the controller held in their non-dominant hand, participants responded by choosing one of three options (first, second, last) presented as a menu in the VE. In a prior pilot study, stimulus intensities for each type and location were perceptually matched between one another to avoid bias due to unequal perceived stimulus amplitudes.

The experiment was divided into three blocks according to the cue waveform. Within each block, participants were provided with all possible stimulus location combinations (12 possibilities) for all possible oddball positions in the sequence (3 possibilities) in a fully randomized order (i.e. 36 trials per block).

3.2.3 Results

We calculated the oddball correct identification rates for the 6 possible location pairs, for each of the 3 stimulus waveforms (see Table 1).

Oddball identification rates were not normally distributed (Shapiro-Wilk normality tests). In all cases, they were found to be significantly above chance levels (one-sample Wilcoxon signed-rank tests). Stimuli on the lateral actuators (T/P) were correctly differentiated more than 83% of the time for all waveforms. Except for the N waveform,

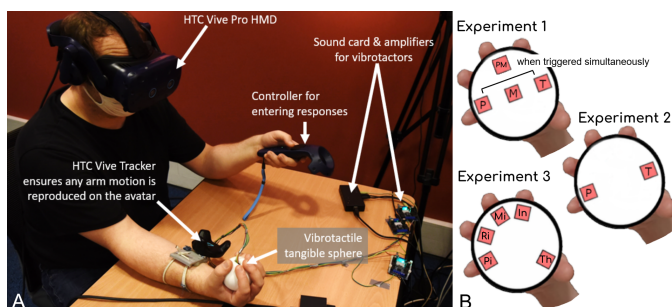


Fig. 4. (A) Experiment setup: participants held a tangible sphere in their dominant hand tracked using a forearm-mounted tracker. They viewed the experimental environment through a HMD and responded to experimental questions with a controller held in their non-dominant hand. (B) Different actuators configurations are used throughout the experiments.

	H1			H2			H3		
	T/P	T/M	P/M	M/PM	T/PM	P/PM			
Impact I	0,833*	1*	0,833*	0,5*	0,5*	0,5*			
	$p = 6,104e-5$ effsize = 0,88 CI = [0,67; 0,92]	$p = 6,104e-5$ effsize = 0,88 CI = [0,75; 1]	$p = 3,05e-5$ effsize = 0,88 CI = [0,83; 0,92]	$p = 0,002$ effsize = 0,72 CI = [0,41; 0,67]	$p = 0,0007$ effsize = 0,77 CI = [0,42; 0,74]	$p = 0,001$ effsize = 0,76 CI = [0,34; 0,58]			
Noise N	0,833*	0,667*	0,917*	0,667*	0,833*	0,667*			
	$p = 6,104e-5$ effsize = 0,87 CI = [0,68; 0,92]	$p = 0,0003$ effsize = 0,85 CI = [0,5; 0,75]	$p = 6,104e-5$ effsize = 0,85 CI = [0,75; 1]	$p = 0,0001$ effsize = 0,84 CI = [0,5; 0,75]	$p = 3,052e-5$ effsize = 0,88 CI = [0,74; 0,91]	$p = 0,0002$ effsize = 0,78 CI = [0,5; 0,76]			
Sine S	1*	0,833*	0,833*	0,5*	0,833*	0,917*			
	$p = 3,052e-5$ effsize = 0,88 CI = [0,75; 1]	$p = 6,104e-5$ effsize = 0,87 CI = [0,67; 0,75]	$p = 3,052e-5$ effsize = 0,88 CI = [0,74; 0,91]	$p = 0,0004$ effsize = 0,79 CI = [0,42; 0,75]	$p = 3,052e-5$ effsize = 0,88 CI = [0,75; 0,92]	$p = 3,052e-5$ effsize = 0,88 CI = [0,74; 1]			

TABLE 1

Median correct oddball identification rates for each location pair within each waveform condition (redder values indicate better performance).

(*) indicates the median rate significantly differs from chance level (0.33) (Wilcoxon signed-rank tests: 95%-confidence intervals (CI) for the median, p -value and effect size are shown below).

the same was the case for differentiation between the localized (T or P) and monolithic (M) actuator locations. The lateral actuator locations were more often confused with the pseudo-monolithic (PM) rendering, but performances were still above chance. Finally, the monolithic (M) and pseudo-monolithic (PM) rendering was often confused, but discrimination rates were still above chance.

We used a generalized linear mixed model to study participants responses with respect to the waveform, the pair of stimulus locations used and the position of the Oddball. Participants were considered as a random effect in the model. Analysis of deviance of the answers showed significant effects of the waveform, location and the oddball position ($p < 0.001$). We did not find any significant interaction effects between the waveform and the stimulus locations. We performed a post-hoc analysis on the different conditions using a Tukey test adapted to the logistic generalized regression model. Trials where the oddball was presented first had higher correct answer rates compared to when the oddball was in second ($Z = 4.02, p < 0.001$) or third position ($Z = 3.14, p = 0.005$). This indicates the task was likely cognitively easier when the oddball came first, however given that the proportion of trials where the oddball came first was identical for all Waveform and Location pair conditions, it is unlikely to affect our conclusions. Correct answer rates were significantly higher when the Sine waveform was used compared to Impact ($Z = -3.38, p = 0.002$) and Noise ($Z = -2.56, p = 0.028$). Participants had a significantly lower rate of correct answers for M/PM compared to all other pairs of stimuli except P/PM ($p < 0.001$), as well as for P/PM compared to T/P and M/P ($p < 0.001$).

3.2.4 Discussion

H1 is clearly supported, as trials combining T and P as the standard and oddball (and vice-versa) yield a very high rate of correct answers, regardless of the cue type. H2 is also strongly supported, with trials respectively combining T or P and M as the standard and oddball (and vice-versa) also yield correct identification rates significantly above chance levels, for all three cue types.

H3 however is not supported. Despite discrimination rates between M and PM being among the lowest for all stimuli, these discrimination rates are still significantly above chance, indicating that PM yields a different sensation to M. While this does not conclusively prove that PM is unusable to provide the sensation of a central monolithic vibration of the object, caution should be exercised when considering substituting M with PM. Furthermore, for impact cues, confusion rates for T and PM as well as between P and PM are significantly higher than those between T and M and P and M respectively. This is likely due to an effect of stimulus duration on discrimination ability.

It is therefore sensible to conclude that spatializing stimuli within a hand-held object is possible using a multi-actuator approach, as cues from localized actuators are readily distinguishable from one another and localized cues are markedly different from monolithic cues. This effect also seems robust to large variations in the stimulus waveforms.

However, it appears that for situations where both monolithic feedback and localized feedback are required, “simulating” monolithic feedback by simultaneously triggering multiple localized actuators may not be a functional alternative. Thus, it is likely that additional actuators may be required in such cases.

3.3 Discrimination between localized VF cues

3.3.1 Research question and hypotheses

We hypothesize that when using localized VF, varying the amplitude distribution across the actuators at the T and P locations (see subsection 3.2.2) could yield an impression of the stimulus location moving from side to side. This idea is based on the “funneling illusion” [58], [59] previously discussed in subsection 3.2.1. To assess the potential effectiveness of this approach, we conduct an experiment aiming to calculate the just-noticeable-difference (JND) in stimulus location based on the balance of actuator amplitudes when a stimulus is simultaneously played on two opposing lateral actuators. Given the difficulty in discriminating between impact stimuli played on a single localized actuator and those played simultaneously on two actuators in the previous experiments (See T/PM and P/PM columns in Table 1), we expect any potentially occurring funneling illusion to be waveform-dependent.

Our hypotheses are the following: (H4) Varying the amplitude balance will lead to the impression of the resulting stimulus being located more towards one side or the other, regardless of the stimulus type. (H5) JNDs for the waveforms N and S will be lower than for I (i.e. more “location levels” will be discernible).

3.3.2 Materials and methods

To investigate this, we varied the amplitude distribution $R = A_P - A_T, R \in [-1; 1]$ between both lateral actuators (T and P locations) whose amplitude is respectively noted A_T and A_P .

We used a two-alternative forced choice (2AFC) procedure based on the method of constant stimuli. In each trial, participants were presented a reference stimulus with $R=0$ (equal amplitude on both actuators) and a test stimulus with R chosen from one of 21 possible levels (0.1 increments

in $R \in [-1; 1]$). They were asked whether the test felt positioned more towards the thumb or towards the pinkie, compared to the reference. The order of test and reference stimuli was randomized within trials. We considered the same set of representative stimulus waveforms and controlled grasp position as in subsection 3.2.2.

The same group of 16 participants (8F, 8M, all right-handed, ages 21-52 (M=26.9, SD=8.8)) took part in the experiment, following the first one. They performed three blocks of 84 trials (4 repeats per value of R), with each block corresponding to one of the three stimulus types \mathbb{I} , \mathbb{N} or \mathbb{S} (see subsection 3.2.2). The order of blocks was counterbalanced across participants. Trial order within blocks was fully random.

After the experiment, participants filled out a short questionnaire assessing their fatigue, perceived task difficulty and performance, and asking them to estimate the number of actuators inside the tangible sphere.

3.3.3 Results

We plotted the proportion of “Test stimulus more towards the pinkie” responses against the amplitude balance R ($R = 1$ indicated a stimulus played entirely on the pinkie-side actuator \mathbb{P}) in Figure 6. We fit cumulative Gaussian distributions to each subject’s data to obtain psychometric functions, from which we derive the 75%-JND.

Poor performance for the impact (\mathbb{I}) stimulus prevented us from computing 75%-JNDs for all participants. For the \mathbb{N} and \mathbb{S} stimuli, mean JNDs were found at 0.27 (Med: 0.26, IQR: 0.12-0.41) and 0.42 (Med: 0.27, IQR: 0.22-0.53) respectively. Individual JNDs were not normally distributed (Shapiro-Wilk tests), and a non-parametric Friedman test showed no significant difference between the JND distributions for \mathbb{N} and \mathbb{S} ($W = 0.14$, $p = 0.13$).

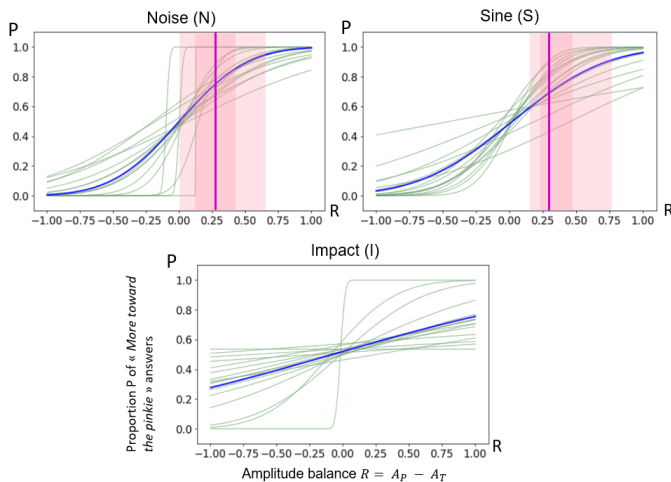


Fig. 6. Psychometric functions for each cue type. Individual functions are shown in green and the mean psychometric function in blue. Dark red bars show the median JND, surrounded by areas respectively marking the interquartile range and full range of individuals’ JNDs in increasingly lighter shades of red.

3.3.4 Discussion

When looking at subject performances for $R = 1$ or $R = -1$, we note that the results appear coherent with data for the \mathbb{P}/\mathbb{PM} and \mathbb{T}/\mathbb{PM} location pairs from subsection 3.2.

H4 is not supported, as for \mathbb{I} the effect is minimal and insufficient to calculate a 75%-JND value. It appears that the stimuli with longer durations (\mathbb{N} and \mathbb{S}) are better suited to achieving the desired effect, since they both allow approx. 7 virtual stimulus locations to be rendered by varying the amplitude balance, while \mathbb{I} allows 2 to 3 locations to be rendered. H5 is therefore also supported. On average, questionnaires showed that participants estimated that the tangible sphere contained between 4 and 5 actuators (min: 1; max: 9). This confirms that varying R often yielded the illusion of a change of stimulus location within the object which is coherent with the calculated discriminable locations.

This means that for applications relying on longer stimuli, it seems reasonable to count on the funneling effect to use a minimal number of actuators to provide a localized sensation. In our case, this translates to 2 actuators possibly providing up to 7 discriminable virtual vibration locations within the sphere. For applications relying on shorter stimuli (e.g., impacts) where more than 3 distinct perceived vibration locations are desired, it seems that more actuators would be required.

3.4 Vibrotactor location identification

3.4.1 Research question and hypotheses

The goal of the present experiment was to assess the ability of participants to determine the location of one among many vibrotactors used alone inside spherical tangible object, as well as possible effects of stimulus intensity on location discrimination performance. Prior experiments showed that approaches with many actuators may be particularly relevant to use-cases involving short stimuli (see subsection 3.3).

With the design of vibrotactile tangibles discussed in subsection 3.1.1, the actuators are not mechanically isolated from one another in any particular way. Therefore vibrations from one actuator may propagate throughout the object, stimulating the whole hand. This may prove problematic in the context of using many localized actuators to render short stimuli.

Because of the similar results observed for \mathbb{S} and \mathbb{N} stimuli in both previous experiments, we omitted the sine stimuli and chose to focus on impact stimuli at two different frequencies as well as noise stimuli for the present experiment. Our first hypothesis is (H6): Stimulus waveform may affect location discrimination accuracy.

Tactile sensitivity varies across the hand. For example, the best vibrotactile acuity (in the millimeter range) is observed on the volar face of the fingers [60], [61], against ranges going up to a few centimeters on the palm [62]. Thus, it seems reasonable to assume that actuator location with respect to the hand may affect location discrimination, leading us to hypothesize (H7): Stimulus locations relative to the hand will affect location discrimination accuracy.

We suppose that intensity is a major factor in discriminating the location of a vibrotactor, and thus hypothesize (H8): Stimulus intensity will affect location discrimination accuracy.

3.4.2 Materials and methods

This experiment used the spherical tangible capable of holding 5 vibrotactors (see Figure 2-B). The vibrotactors are

located beneath each finger, in what we hypothesize is the best-case scenario for vibration source location discrimination as stimuli are generated close to the fingertips, i.e. the most sensitive part of the hand [60], [61], [63]. Markings on the surface of the sphere allowed the experimenter to ensure participants' fingers were properly placed above the actuators.

We considered three experimental variables:

- **Wav** - The impact stimulus waveform (White noise, 50Hz exponentially decaying sine or 100Hz exponentially decaying sine).
- **Pos** - The pair of reference and comparison stimulus locations on the fingers.
- **Lvl** - The stimulus intensity level (weak or strong).

The weak intensities were chosen as the lowest mean intensity at which 100% of participants could detect a stimulus, as determined in a prior pilot study. The strong intensity was chosen as 10 times the weak intensity level.

The experiment involved 18 right-handed participants (13 M, 5 F, ages 21-33 ($M=24.3$, $SD=3.8$)).

3.4.3 Procedure

Participants performed a series of trials during which vibrations were played on a random actuator, after which they were asked to indicate where they felt the vibration using a continuous circular slider representing the tangible sphere (see Figure 7).

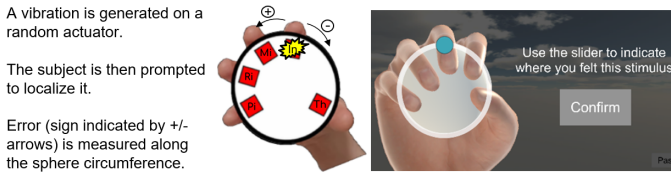


Fig. 7. Location discrimination experiment: After a vibration, the subject must locate it within their hand using an on-screen circular slider.

The experiment was divided into three blocks. Each block used a different type of stimulus played for 0.3 seconds: an exponentially decaying sinusoid at 50Hz, another at 100Hz or white noise. Each one of the 25 ordered pairs of stimulus locations inside the sphere was repeated three times, but at each of the two possible intensity levels for **Lvl**, for a total of 150 trials in each block. All conditions were balanced between participants to avoid any order effects. Participants' responses (i.e. slider position values) were collected for each trial.

The sphere having a circumference of 22cm, we measure localization error mean localization errors ranging from -11cm to 11cm. Negative error values indicate a bias in the pinkie toward thumb direction, while positive values indicate a bias in the opposite direction (see Figure 7). The mean distance between finger pairs on the circular response slider was 3cm. In the following, we label the actuator positions as **Th** (Thumb), **In** (Index), **Mi** (Middle finger), **Ri** (Ring finger) and **Pi** (Pinkie) (see Figure 7).

3.4.4 Results

We used a linear mixed model on the collected data with respect to the three conditions (**Wav**, **Pos** and **Lvl**) to study participant responses with respect to the stimuli. Participants were considered as a random effect in the model.

Measures of mean location discrimination errors are reported in Table 2. An analysis of deviance for the slider position answer showed a significant effect on the **Pos** condition ($p < 0.001$). We did not find any significant effect for **Lvl** ($p = 0.29$) and **Wav** ($p = 0.45$). However, we observed an interaction effect between **Lvl** and **Pos** ($p < 0.001$). We performed a post-hoc analysis on the different conditions using a Tukey test adapted to the logistic generalized regression model. Regarding **Pos**, we found a significant effect between all the locations ($p < 0.001$), except between the **Th** and the **In** locations ($p = 0.56$) as well as between the **Th** and **Ri** location ($p = 0.13$) on the other side.

	Th	In	Mi	Ri	Pi
Mean error \pm SD (cm along the circumference)	-1,3 \pm 2,6	-1,5 \pm 3,7	0,4 \pm 2,4	0,2 \pm 2,9	1,9 \pm 1,6

TABLE 2
Mean error in localization of the vibration source.

3.4.5 Discussion

Participants were generally accurate in estimating the position of the vibration source, placing it within a single between-finger distance from its actual position. The spread of errors is similar across fingers except for the index and pinkie which show higher spreads indicative of a higher uncertainty in subject responses.

H6 was not verified as **Wav** did not affect localization errors. Our results tend to support H7, as localization error significantly varies with respect to stimulus position. Furthermore, we observe a bias towards localizing stimuli at **Th** and **In** more in the direction of the thumb whereas stimuli at **Mi**, **Ri** and **Pi** are localized more towards the pinkie. Localization error for **Th** and to a lesser extent **In** could partly be explained by vibrations propagating to the nearby thenar eminence, which could skew the perception of location towards the side of the thumb. H8 was not verified, as localization errors were unaffected by **Lvl**.

With respect to the design of vibrotactile tangibles, these results indicate that applications requiring fine vibration location rendering on the object surface could likely use upwards of 5 vibrotactors. However, space and adequate mechanical isolation rapidly become challenging with an increased number of vibrotactors.

4 MULTI-ACTUATOR VIBROTACTILE RENDERING

Based on the literature on vibrotactile rendering presented in subsection 2.4, we implemented feedback for sensations arising from manipulating vibrating or pulsating objects, impacts between a hand-held object and the environment, textures of the environment and forces or impacts resulting from the contents of a manipulated container. The present section details the algorithms for each of these haptic effects and how they can leverage the potential offered by stimulus spatialization. An evaluation of the added benefit of spatialization is discussed in section 5. In the proposed VR interactions, we use the same system setup as previously described in subsection 3.1, with a spherical tangible object containing three built-in vibrotactors (at locations **T**, **P** and **M**). This allows each of the proposed rendering schemes to generate either localized (using actuators at locations **T** and **P**) or monolithic (using actuator location **M**) VF.

4.1 Vibrating and pulsating objects

Vibrotactile feedback can serve to render properties of manipulated objects themselves, in particular vibrations coming from an object held in hand (e.g., a motorized tool). In this case, there is no generally applicable rendering model as the vibrations depend on the object to be rendered.

Given that rendering object vibrations relies on prolonged sinusoidal stimuli (either pure sines or a superposition of sines), results from our prior experiments indicate that spatialization of these cues along an axis crossing the handheld object could be effective by continuously varying the balance of amplitudes between off-center vibrotactors (see subsection 3.3). In this manner, a virtual vibration source could be displayed more or less off center with respect to the position at which virtual object is grasped.

In addition to vibration, another phenomenon termed *haptic beats* [64] can be used to render pulsing virtual objects held in hand. This phenomenon occurs when two sine waves with a small difference in frequency are delivered to two separate but close locations on the skin, introducing a “beating” sensation. It is especially perceivable under the fingertips and the palm [64]. We reproduced this effect within the spherical end-effector by playing different frequency vibrations on two distinct vibrotactors. A similar but less marked effect arises when superimposing sinusoidal waveforms on a single actuator, allowing for a monolithic alternative.

4.2 Rendering impacts

When exploring a VE with their hand or interacting with the VE with a grasped object or tool, the user can impact virtual objects or surfaces, generating vibrations depending on impact location, speed and material of the objects involved. As discussed in subsection 2.4, vibrotactile impact rendering commonly uses an exponentially decaying sinusoid model [43], [44]. The impact waveform used to generate the signal itself is defined by $x(t) = A(v)e^{-\beta t} \sin(\omega t)$, with A the amplitude, a function of the impact velocity, and β , ω depending on the simulated material. Here, we used material parameters identified by Choi et al. [18] and Okamura et al. [44]. We consider two impact scenarios.

A. A small object held in the hand (e.g., a baseball) impacts another (e.g., a table). In this situation a uniform vibration is generated within the object, and the impact can be rendered either using a single monolithic actuator or multiple localized actuators delivering simultaneous equal-amplitude stimuli (“pseudo-monolithic” case described in subsection 3.2). Alternatively, the impact location may be precisely rendered by varying the balance of amplitudes between simultaneously activated localized actuators, as described in subsection 3.3.

B. A large object held in hand (e.g., a bat) impacts another one (e.g., a ball) away from the hand. In this case, the impact should produce more intense vibrations on the side towards which the impacts occurs. Thus, a straightforward approach consists in displaying the impact on the closest localized actuator. However, since impact waveform location discrimination is not as good as for other waveforms (see subsection 3.2 and subsection 3.3), it may also be sufficient

to use a monolithic or pseudo-monolithic approach. This question is further addressed in section 5.

4.3 Macroscopic texture

As discussed in subsection 2.4, macroscopic texture features are mediated predominantly by spatial cues, and often rendered as a series of impacts as users scan the virtual surface. Since setting up a collider to trigger impacts for every bump of a textured surface is rather cumbersome, we opt for a simplified approach where textured surfaces were modeled by two parameters: distance between bumps and bump orientation. As the user dragged a sphere over the virtual surface, we compute the estimated frequency of impacts from the velocity at the contact and adjust the rate at which impact stimuli are triggered based on the result.

Just like impacts at the level of the hand (case A in subsection 4.2), macro texture can be rendered using a monolithic vibrotactor or a pseudo-monolithic substitute. However, we hypothesized that triggering the impact waveforms on the actuator located in the direction towards which the user is moving the handheld object across the surface may reinforce the illusion, in which case localized rendering would be beneficial.

4.4 Texture roughness

We chose to adapt the open-source Penn Haptic Texture Toolkit [5] to our needs in the present system. This toolkit follows a data-driven rendering approach based on real-world contact force and acceleration data during free exploration of textures. Each recorded texture is modeled as a collection of autoregressive processes stored in a Delaunay triangulation according to the scanning speed and normal force associated with them. At runtime, the speed of the user over the virtual texture is used to determine which models to interpolate. The resulting coefficients are used to generate the signals sampled at 10kHz for the actuators.

Since procedural generation of audio signals is not possible using Syntacts, we used the Unity audio engine, through which custom filters can be applied over an audio source. This however offers little control over the haptic rendering latency.

Texture roughness could be rendered using a monolithic actuator, however we hypothesize that similarly to macro texture, vibrations on the actuator located towards the direction of motion may also reinforce the illusion. We further investigated this idea in section 5.

4.5 Rendering contents of a grasped object

Previous interactions focused on rendering information about contacts and haptic exploration interactions. Vibrotactile feedback provided inside a tangible object can however also serve to enrich manipulation interactions by providing physical information on the manipulated object itself.

In this context, we explored rendering virtual contents of a grasped object, such as the movement of a liquid in a glass or the movement of dices shaken inside a cup. Rendering small objects such as dice or marbles inside a cup is done using the physics simulation of the Unity engine. For every impact between the object in hand (the “container”) and

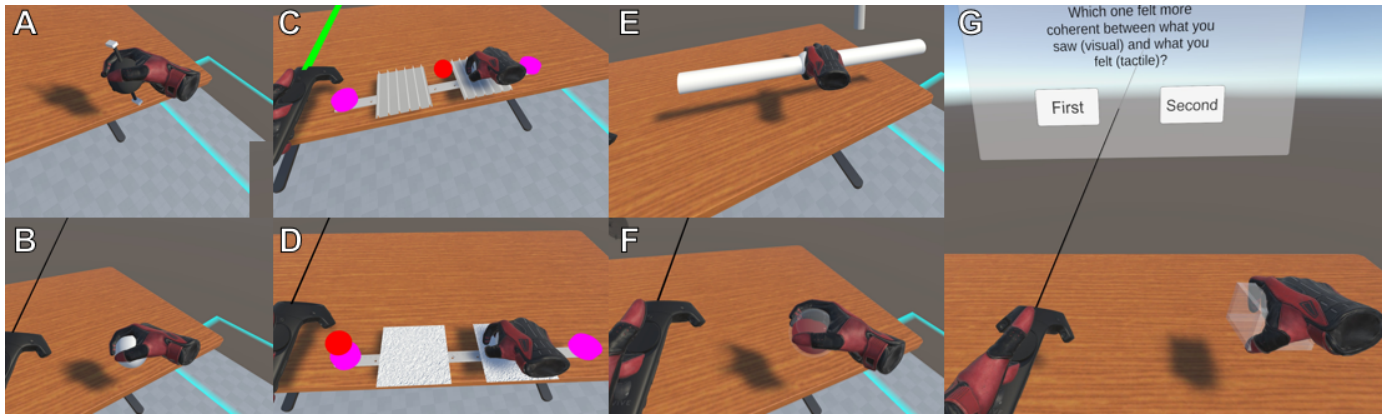


Fig. 8. The experiment assessing perceived coherence of localized vs. monolithic VF focused on 7 different properties: (A) Vibration of objects, (B) pulsating objects, (C) macro texture (bumps), (D) texture roughness, (E) impacts between a handheld object and the environment, (F) liquids contained in a handheld container and (G) solids contained in a handheld container. Subject responses were recorded for each stimulus pair (G).

an object within it (the “contents”), the distance between the impact location and the position of all actuators is calculated. The closest to the virtual impact is then selected to play the impact effect. A higher number of actuators would thus theoretically enable a higher resolution, within the limits of users’ ability to discriminate vibration locations (see subsection 3.4.1). We compare this spatialized approach to a simple monolithic rendering of solid content impacts, wherein the monolithic actuator is triggered every time the contents impact the container, in section 5.

To give the illusion of a fluid inside the object, we attach a virtual mass to the center of the virtual object with an underdamped spring, leading the mass to lag behind the object during movement and to oscillate when the movement stops. The amplitude of vibration displayed on each actuator is inversely proportional to the distance of the mass from said actuator, generating a sensation of motion between actuator locations. This approach can be seen as the vibrotactile equivalent to the purely mechanical rendering of fluids inside a tangible proposed by Sagheb et al. [24]. We compare this approach to monolithic rendering of fluids, wherein a central actuator’s stimulus amplitude is proportional to the lateral displacement of the mass away from the equilibrium point in section 5.

5 DOES LOCALIZED VF PROVIDE A BENEFIT OVER MONOLITHIC VF IN CERTAIN INTERACTIONS?

5.1 Research questions and hypotheses

In this section, we investigate whether using multiple actuators for the rendering schemes discussed in section 4 provides added benefits over single-actuator monolithic feedback. We asked the same research question for each of the feedback schemes, i.e. whether users perceived the localized or monolithic stimuli provided as more coherent with the associated visual feedback.

Our hypotheses are the following: **(H9)** For pulsating objects (haptic beats), the stronger effect that is achievable by using two actuators rather than a superposition of waveforms on a single actuator will lead to the multi-actuator rendering being perceived as more coherent; **(H10)** For vibrating objects and impacts, the additional spatial information provided by localizing the VF towards the location of the virtual source of the vibration will lead to the

combined visual and haptic feedback as being perceived as more coherent than monolithic VF; **(H11)** For cues generated in response to user motion (micro and macro texture, object contents), localizing the VF towards the direction of motion (of the handheld object or its contents) will reinforce the illusion of motion, leading to the feedback to be perceived as more coherent than monolithic VF.

5.2 Materials and methods

We determined the level at which T, P and PM vibrations were perceived at the same level as M through a pilot experiment using a method of adjustment. For each stimulus type and location, participants adjusted the intensity until it perceptually matched that of a monolithic reference stimulus. We used the same sphere and controlled grasp position as in subsection 3.2.2.

To answer our main experimental question, we evaluated the subjective perception of 7 different types of haptic sensations. The experiment is a collection of 6 independent component experiments, each following a within-subject design. This makes up 6 blocks, each dealing with a specific type of haptic sensation (i.e. vibrating object, pulsating object, impacts, macro and micro texture, solid content, liquid content). The order of blocks was counterbalanced across participants using a Latin square design. Within each block, trials consisted in participants being presented with pairs of sensations (one monolithic and one localized). The collected data were participants’ subjective assessment of perceived coherence between visual and tactile stimuli in each condition, when asked in VR *which one felt more coherent between what they saw (visual) and what they felt (tactile)*.

21 participants performed the experiment (8F, 13M, 20 right-handed, ages 22-33 (M=24.6, SD=3.2), 15 having some experience with haptics). Blocks contained 6 repetitions of each stimulus pair and the pair presentation order was counterbalanced across participants.

Vibrating objects: In the vibrating object block (Figure 8-A), participants held a virtual sphere with a visible eccentric rotating mass on either side. In each trial, one mass would spin while either the central monolithic actuator or the closest localized actuator played an associated vibration.

Pulsating objects: In the pulsating object block (Figure 8-B), participants held a sphere whose visual diameter periodically grew and shrunk at the specified pulse rate while

a haptic beats effect was generated at the same rate. This beating effect was either generated by superimposing two different frequency sines on the monolithic actuator or by distributing each sine on one of the localized actuators.

Textures: In the textures block (Figure 8-C,D), the tangible sphere was fixed to a linear guide rail constraining its motion along one axis. Pairs of textured surfaces were displayed in front of the user, who was instructed to drag a virtual sphere across them. A visual guide was used to ensure consistency in hand motion speeds across all trials. The actuator location was switched between virtual surfaces, with playback occurring either on the monolithic central actuator or on the localized actuator to the side towards which motion occurred.

Impacts: In the impacts block (Figure 8-E), participants held a virtual stick extending symmetrically around the hand. An object would then fall and randomly impact the stick on one side. A corresponding vibration was either played on the monolithic actuator or on the localized actuator to the side of the virtual impact.

Contents: In the solid and liquid contents blocks (Figure 8-G,F), participants held a virtual transparent container with a visible marble inside. Participants were instructed to shake the container from side to side for 5 seconds. A ghost hand performing the desired motion was displayed to ensure consistency between trials. Corresponding VF was either played back from the monolithic actuator or using the pair of localized actuators as described in subsection 4.5.

5.3 Results

For each condition, we calculated the proportion of trials in which participants respectively expressed a higher perceived coherence for the monolithic or localized versions of the vibrotactile feedback (see Table 3). We compared each proportion of perceived coherence to the case where localized and monolithic feedback were perceived as equally coherent using exact binomial tests.

Vibrating and pulsating objects: Regardless of the side of the hand on which the virtual vibration occurs, participants perceive the localized VF of object vibration as more coherent than the monolithic alternative. We therefore conclude that H10 is verified for this rendering scenario. However, H9 does not appear to be verified since we fail to show any significant difference in perceived coherence for localized or monolithic feedback of object pulsation. It is worth noting that in this rendering case, localized feedback is not detrimental to perceived coherence despite it not showing any significant benefit.

Impacts: When rendering impacts on a virtual handheld object, participants seem to show no difference in perceived coherence towards localized or monolithic feedback of impact vibrations. This appears coherent with prior results showing poorer impact localization performance (see subsection 3.2 and subsection 3.3). A finer analysis of VF coherence as a function of the side to which the impact occurs reveals no significant difference from the overall trend. For rendering impacts, it thus appears that localized vibrotactile feedback provides neither a benefit nor a disadvantage. H10 is thus not verified in this case.

Macro Texture (bumps): Overall, we found no difference in perceived coherence towards localized or monolithic VF

	Vibrating	T/M	P/M	Pulsating	Impact	T/M	P/M
Localized	63% \uparrow	60% \uparrow	66% \uparrow	53%	47%	45%	49%
Monolithic	37% \downarrow	40% \downarrow	34% \downarrow	47%	53%	55%	51%
Statistics	$p = 2.18e-5$ $CI: (0.57)$	$p = 0.025$ $CI: (0.51)$	$p = 2.29e-4$ $CI: (0.58)$	$p = 0.533$ $CI: (0.61)$	$p = 0.413$ $CI: (0.54)$	$p = 0.327$ $CI: (0.54)$	$p = 0.929$ $CI: (0.58)$

	Bumps	T/M	P/M	Texture	T/M	P/M	Liquid	Solid
Localized	51%	40% \downarrow	63% \uparrow	46%	42%	51%	75% \uparrow	46%
Monolithic	49%	60% \uparrow	37% \downarrow	54%	58%	49%	25% \downarrow	54%
Statistics	$p = 0.753$ $CI: (0.57)$	$p = 0.025$ $CI: (0.49)$	$p = 0.005$ $CI: (0.71)$	$p = 0.396$ $CI: (0.54)$	$p = 0.155$ $CI: (0.53)$	$p = 0.913$ $CI: (0.62)$	$p = 9.61e-9$ $CI: (0.83)$	$p = 0.42$ $CI: (0.55)$

TABLE 3

Percentage of perceived coherence between localized and monolithic stimuli for each interaction. For vibrating objects, impacts, macro and micro texture, T/M and P/M respectively show the detailed proportion of perceived coherence for localized stimuli presented on the thumb and pinkie side of the hand respectively. (*) indicates the distribution is different from a binomial distribution of parameter $p = 0.5$ (Exact binomial test: 95%-confidence intervals (CI) for the true proportion of perceived coherence towards localized VF are shown in brackets, p-value is shown above). Green indicates significant preferences.

for macroscopic texture rendering, contradicting hypothesis H11. Since macroscopic texture is very similar to impacts, this appears coherent with our other results. An analysis of coherence as a function of the motion direction revealed that participants preferred localized feedback when the motion occurred towards the thumb (H11 supported), but preferred monolithic feedback when the motion occurred in the opposite direction (H11 contradicted).

Micro Texture (roughness): Overall, we see no difference in perceived coherence towards localized or monolithic VF for texture roughness. H11 is therefore not supported for rendering of texture roughness.

Contents: When rendering solid contents of a handheld container, participants again did not perceive localized or monolithic VF as more coherent with the VE. This is coherent with prior results on impact rendering for which H10 was not verified. For liquid contents however, participants rate localized feedback as significantly more coherent than monolithic VF. H11 is therefore partially supported.

5.4 Discussion

Both H10 and H11 are partly supported by our results.

Our preliminary investigation highlighted that using the rendering approaches proposed in section 4, users expressed a clear preference for localized VF for both the vibrating objects scenario (involving a prolonged sinusoidal stimulus without leveraging a “funneling” effect) and the liquid contents scenario (also involving a prolonged sinusoidal stimulus, this time leveraging a “funneling” effect). For such effects, using a localized VF setup within the tangible object will provide a benefit in terms of the experience being perceived as more coherent.

In most other scenarios, be they scenarios involving short stimuli (impacts and solid contents) or prolonged stimuli (pulsating objects, texture roughness), the proposed localized rendering approaches did not show any significant benefit in terms of perceived coherence. We can conclude that, for these effects, a localized setup is not necessary, but also not detrimental. Basically, if these effects are to be used alone in a VR scenario, using monolithic feedback would



Fig. 9. VR use-case showcasing the different possible interactions. Immersed users can grasp and manipulate objects (F), feel impacts (B,C,D) and textured surfaces through an object held in hand (e.g., sand and wood in A and B), manipulate vibrating tools (the mixer in D) and feel the dynamics of object contents (e.g., wine in the bottle in E). A video of this use-case is available as supplemental material.

make sense given the lower system complexity. However, if used in a scenario where other effects which can benefit from localized VF are being used, a localized VF setup is a viable option for the entire set of considered stimuli.

It is also possible that using different waveforms could unlock untapped potential for localization in these scenarios, although this would require further investigation. Our perceptual studies have shown that funneling likely cannot be used in conjunction with shorter stimuli (impacts), which somewhat limits the scope of possibilities. There may however be potential benefits to be gained from modifying the proposed rendering approaches for e.g., pulsating objects or texture roughness in order to leverage “funneling” effects, to e.g., provides sensations of a moving origin of the pulsation within the object or a changing contact point between object and textured surface.

The results obtained in the macro texture are ambiguous. There seems to be a preference towards monolithic feedback when motion is done towards the thumb side of the hand, but for localized feedback when motion is done towards the pinkie side. The thumb side of the hand is more sensitive to VF due to the thenar eminence. This may cause the localized sensation on that side to be perceived as stronger, or may lead to the bumps being felt too much on the side of the object, rather than towards the tip of the lower hemisphere of the object, where the virtual physical contact is expected to occur. This could have caused a higher perceived incoherence in the localized condition, for motions going towards the thumb side of the hand. Further investigation at different intensity levels and using the non-dominant hand could be interesting to verify this.

For this result, it is not possible to give clear guidelines for designers. If monolithic feedback is indeed preferred in a specific situation this would mean that applications combining multiple effects would require a combination of localized and monolithic central actuators, which would be selected based on how appropriate they are to a given stimulus to be rendered.

5.5 Use-Cases

Since our experiment showed that multi-actuator localized rendering schemes were either beneficial over or equivalent to monolithic VF, we developed an VR use-case using only

two lateral actuators and demonstrating the inclusion of all rendering schemes discussed in section 4 into a single coherent VE. A video of these use-cases is available as supplementary material.

In the VE, users can explore an outdoor scene, manipulating various objects such as e.g., a ball (see Figure 9-A,B), a bat (see Figure 9-C), a hand-held mixer (see Figure 9-D), a wine bottle (see Figure 9-E) while receiving tangible haptic feedback. We render impacts (Figure 9-B,C,D), surface texture (Figure 9-A,B,C), vibrating handheld objects (Figure 9-D) as well as liquid contents (Figure 9-E) of manipulated objects as per the approaches described in section 4.

6 CONCLUSION AND PERSPECTIVES

This paper investigates the inclusion of vibrotactile feedback (VF) within tangible objects to enrich haptic manipulation in VR. In particular, we explored the use of multiple vibrotactors to render spatial properties of virtual interactions. We demonstrate that this allows a multitude of dynamic haptic properties to be rendered during interaction with the virtual environment (VE). These include rendering vibrating and pulsating objects, impacts, texture components as well as rendering virtual contents of handheld containers. Through a series of perception studies, we show that multi-actuator localized VF can yield sensations which are perceivably localized within a handheld object, and can in some cases be used to create a haptic funneling effect within a handheld tangible object. We then assess the added benefit of multi-actuator VF in terms of perceived coherence between visual and haptic feedback through a human subject experiment. In certain circumstances, multi-actuator localized vibrotactile feedback is shown to yield more coherent haptic feedback than monolithic object vibration, and is almost never detrimental to perceived coherence.

These studies show the potential of multi-actuator VF within tangible objects in VR, but currently only yield a limited range of scenarios where this approach is beneficial. Thus, we plan to study additional rendering schemes to determine whether they could benefit from this approach. We also plan to investigate different vibrotactor arrangements, tangible object geometries, and grasp types as these may impact the benefit provided by multi-actuator VF.

The approach we investigate has the major advantage of not encumbering the user and avoids any potentially disruptive haptic stimuli caused by wearables not matching expected stimuli from the virtual environment. However, this approach is not without challenges that still need to be overcome. First of all, actuating tangible objects can require comparatively more actuators and thus a higher system complexity and cost. Second, for our approach to work, the tangible object requires an internal power source or has to be wired. Both challenges may limit freedom of manipulation. Because of this, the advantages of our approach are particularly prominent for interfaces and systems where interaction is constrained to a single tangible object (e.g., [19]). They can however be extended to conventional passive haptics using multiple tangible objects (e.g., [20]), encounter-type haptic devices supporting multiple end-effector tangibles (e.g., [28]), or even grounded force feedback devices with tangibles as end-effectors (e.g., [27]). Finally, our current rendering schemes ignored vibration propagation and mechanical effects which may impact the quality of the perceived stimuli. A challenge for future work will be to factor in these effects. Whether this is sensible given the added required effort remains to be investigated.

Our investigations were done with the tangible object grasped in hand, so it may be interesting to investigate cases where grasp and release occur, or where the hand only touches the tangible, as these could further extend the applicability of the proposed approach. This initial study only investigated subjective preference of rendering schemes. It may be useful to look at impacts on task performance in the future to get a more nuanced view of what spatialization of VF cues can achieve. Finally, extended investigations with larger and more diverse sample populations could be beneficial in ensuring these preliminary results are generalizable.

Mixed haptics approaches open up vast possibilities for enriching interaction through the complementary nature of tangible and VF, which particularly fit the multi-sensory context of immersive VR. Our present contributions focused on augmenting interaction realism by rendering richer haptics, but our system design also offers possibilities for providing abstract information to support interaction performance, convey affective information and support communication in VR. Because of all this, we believe that mixed haptics has the potential to greatly enhance user engagement, performance, immersion and presence in VEs.

ACKNOWLEDGMENTS

This project has received funding from the Inria Défi project "DORNELL", the ANR project "MIMESIS", and from the European Union's Horizon 2020 programme under grant agreement No 801413; project "H-Reality".

REFERENCES

- [1] H. Culbertson, S. B. Schorr, and A. M. Okamura, "Haptics: The Present and Future of Artificial Touch Sensation," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 385–409, 2018.
- [2] G. V. Popescu, G. C. Burdea, and H. Trefftz, "Multimodal interaction modeling," in *Handbook of Virtual Environments*, 2002, pp. 475–494.
- [3] D. Wang, K. Ohnishi, and W. Xu, "Multimodal haptic display for virtual reality: A survey," *IEEE Trans. Industrial Electronics*, vol. 67, no. 1, pp. 610–623, 2019.
- [4] E. Pezent, A. Israr, M. Samad, S. Robinson, P. Agarwal, H. Benko, and N. Colonnese, "Tasbi: Multisensory squeeze and vibrotactile wrist haptics for augmented and virtual reality," in *IEEE World Haptics Conf.*, 2019, pp. 1–6.
- [5] H. Culbertson, J. Jose Lopez Delgado, K. J. Kuchenbecker, L. Delgado, and J. Jose, "The Penn Haptic Texture Toolkit for Modeling, Rendering, and Evaluating Haptic Virtual Textures," Tech. Rep., 2014. [Online]. Available: http://repository.upenn.edu/meam_papers/299
- [6] C. Pacchierotti, S. Sinclair, M. Solazzi, A. Frisoli, V. Hayward, and D. Prattichizzo, "Wearable haptic systems for the fingertip and the hand: taxonomy, review, and perspectives," *IEEE Trans. Haptics*, vol. 10, no. 4, pp. 580–600, 2017.
- [7] H. P. Saal and S. J. Bensmaia, "Touch is a team effort: interplay of submodalities in cutaneous sensibility," *Trends in neurosciences*, vol. 37, no. 12, pp. 689–697, 2014.
- [8] H. Culbertson and K. J. Kuchenbecker, "Importance of matching physical friction, hardness, and texture in creating realistic haptic virtual surfaces," *IEEE Trans. Haptics*, vol. 10, no. 1, pp. 63–74, 2016.
- [9] I. Choi, H. Culbertson, M. R. Miller, A. Olwal, and S. Follmer, "Gravity: A wearable haptic interface for simulating weight and grasping in virtual reality," in *Proc. 30th Annual ACM Symposium on User Interface Software and Technology*, 2017, pp. 119–130.
- [10] I. Choi, E. Ofek, H. Benko, M. Sinclair, and C. Holz, "Claw: A multifunctional handheld haptic controller for grasping, touching, and triggering in virtual reality," in *Proc. CHI Conf. Human Factors in Computing Systems*, 2018, pp. 1–13.
- [11] N. Dunkelberger, J. Bradley, J. L. Sullivan, A. Israr, F. Lau, K. Klumb, F. Abnoui, and M. K. O'Malley, "Improving perception accuracy with multi-sensory haptic cue delivery," in *Int. Conf. Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 2018, pp. 289–301.
- [12] V. Yem and H. Kajimoto, "Wearable tactile device using mechanical and electrical stimulation for fingertip interaction with virtual world," in *IEEE Virtual Reality (VR)*, 2017, pp. 99–104.
- [13] M. Aggravi, F. Pausé, P. R. Giordano, and C. Pacchierotti, "Design and evaluation of a wearable haptic device for skin stretch, pressure, and vibrotactile stimuli," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2166–2173, 2018.
- [14] A. Delazio, K. Nakagaki, R. L. Klatzky, S. E. Hudson, J. F. Lehman, and A. P. Sample, "Force jacket: Pneumatically-actuated jacket for embodied haptic experiences," in *Proc. CHI Conf. Human Factors in Computing Systems*, 2018, pp. 1–12.
- [15] C. V. Keef, L. V. Kayser, S. Tronbolli, C. W. Carpenter, N. B. Root, M. Finn III, T. F. O'Connor, S. N. Abuhamdi, D. M. Davies, R. Runser et al., "Virtual texture generated using elastomeric conductive block copolymer in a wireless multimodal haptic glove," *Advanced Intelligent Systems*, vol. 2, no. 4, p. 2000018, 2020.
- [16] M. B. Khoudja, M. Hafez, J. Alexandre, A. Kheddar, and V. Moreau, "Vital: A vibrotactile interface with thermal feedback," *Journée scientifique Int.e IRCICA, Lille*, 2004.
- [17] D. Wolf, M. Rietzler, L. Hnatek, and E. Rukzio, "Face/on: multimodal haptic feedback for head-mounted displays in virtual reality," *IEEE Trans. visualization & computer graphics*, vol. 25, no. 11, pp. 3169–3177, 2019.
- [18] I. Choi, E. Zhao, E. J. Gonzalez, and S. Follmer, "Augmenting perceived softness of haptic proxy objects through transient vibration and visuo-haptic illusion in virtual reality," *IEEE Trans. Visualization and Computer Graphics*, 2020.
- [19] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson, "Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences," in *Proc. CHI Conf. on human factors in computing systems*, 2016, pp. 1968–1979.
- [20] L.-P. Cheng, L. Chang, S. Marwecki, and P. Baudisch, "iturk: Turning passive haptics into active haptics by making users reconfigure props in virtual reality," in *Proc. CHI Conf. Human Factors in Computing Systems*, 2018, pp. 1–10.
- [21] H. G. Hoffman, "Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments," in *Proc. IEEE Virtual Reality Ann. Int. Symp.*, 1998, pp. 59–63.
- [22] H. Brument et al., "Pyramid escape: Design of novel passive haptics interactions for an immersive and modular scenario," in *Proc. IEEE Conf. Virtual Reality & 3D User Interfaces (VR)*, 2019, pp. 1409–1410.

- [23] X. de Tinguy *et al.*, "How different tangible and virtual objects can be while still feeling the same?" in *Proc. IEEE World Haptics Conf.*, 2019, pp. 580–585.
- [24] S. Sagheb, F. W. Liu, A. Bahremand, A. Kidane, and R. LiKamWa, "Swish: A shifting-weight interface of simulated hydrodynamics for haptic perception of virtual fluid vessels," in *Proc. Annual ACM Symposium on User Interface Software and Technology*, 2019, pp. 751–761.
- [25] J. Shigeyama, T. Hashimoto, S. Yoshida, T. Narumi, T. Tanikawa, and M. Hirose, "Transcalibur: A weight shifting virtual reality controller for 2d shape rendering based on computational perception model," in *Proc. CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–11.
- [26] M. Sinclair, E. Ofek, M. Gonzalez-Franco, and C. Holz, "Capstan-crunch: A haptic vr controller with user-supplied force feedback," in *Proc. ACM symposium on user interface software and technology*, 2019, pp. 815–829.
- [27] Y. Bae, B. Cha, and J. Ryu, "Calibration and Evaluation for Visuo-haptic Collocation in Haptic Augmented Virtuality Systems," *Int. J. Control, Automation and Systems*, vol. 18, no. 5, pp. 1335–1342, 2020.
- [28] V. R. Mercado, T. Howard, H. Si-Mohammed, F. Argelaguet, and A. Lécuyer, "Alfred: the haptic butler on-demand tangibles for object manipulation in virtual reality using an ethd," in *Proc. IEEE World Haptics Conf.*, 2021, pp. 373–378.
- [29] X. De Tinguy, T. Howard, C. Pacchierotti, M. Marchal, and A. Lécuyer, "WeATaViX: WEearable Actuated TAngibles for Virtual reality eXperiences," in *Haptics: Science, Technology, Applications*, ser. EuroHaptics, vol. 12272, 2020, pp. 262–270.
- [30] R. Kovacs, E. Ofek, M. Gonzalez Franco, A. F. Siu, S. Marwecki, C. Holz, and M. Sinclair, "Haptic pivot: On-demand handhelds in vr," in *Proc. 33rd Annual ACM Symposium on User Interface Software and Technology*, 2020, pp. 1046–1059.
- [31] J. Kildal, "Kooboh: Variable tangible properties in a handheld haptic-illusion box," in *Int. Conf. Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 2012, pp. 191–194.
- [32] M. Sun, W. He, L. Zhang, and P. Wang, "Smart haproxy: A novel vibrotactile feedback prototype combining passive and active haptic in ar interaction," in *IEEE Int. Symp. Mixed & Augmented Reality Adjunct (ISMAR-Adjunct)*, 2019, pp. 42–46.
- [33] P.-A. Cabaret, T. Howard, C. Pacchierotti, M. Babel, and M. Marchal, "Perception of spatialized vibrotactile impacts in a hand-held tangible for virtual reality," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*, 2022, pp. 264–273.
- [34] S. Choi and K. J. Kuchenbecker, "Vibrotactile display: Perception, technology, and applications," *Proc. IEEE*, vol. 101, no. 9, pp. 2093–2104, 2012.
- [35] A. Passalenti, R. Paisa, N. C. Nilsson, N. S. Andersson, F. Fontana, R. Nordahl, and S. Serafin, "No strings attached: Force and vibrotactile feedback in a virtual guitar simulation," in *IEEE Conf. Virtual Reality and 3D User Interfaces (VR)*, 2019, pp. 1116–1117.
- [36] H. Culbertson, J. M. Romano, P. Castillo, M. Mintz, and K. J. Kuchenbecker, "Refined methods for creating realistic haptic virtual textures from tool-mediated contact acceleration data," in *IEEE Haptics Symposium (HAPTICS)*, 2012, pp. 385–391.
- [37] V. de Jesus Oliveira *et al.*, "Designing a vibrotactile head-mounted display for spatial awareness in 3d spaces," *IEEE TVCG*, vol. 23, no. 4, pp. 1409–1417, 2017.
- [38] O. B. Kaul, K. Meier, and M. Rohs, "Increasing presence in virtual reality with a vibrotactile grid around the head," in *IFIP Conf. Human-Computer Interaction*. Springer, 2017, pp. 289–298.
- [39] L. Pantera *et al.*, "Multitouch vibrotactile feedback on a tactile screen by the inverse filter technique: Vibration amplitude and spatial resolution," *IEEE Trans. Haptics*, vol. 13, no. 3, pp. 493–503, 2020.
- [40] F. Chinello, C. Pacchierotti, M. Malvezzi, and D. Prattichizzo, "A three revolute-revolute-spherical wearable fingertip cutaneous device for stiffness rendering," *IEEE Trans. Haptics*, vol. 11, no. 1, pp. 39–50, 2017.
- [41] X. Zhong, J. Wu, X. Han, and W. Liu, "Mobile terminals haptic interface: a vibro-tactile finger device for 3d shape rendering," in *Int. Conf. Intelligent Robotics and Applications*, 2017, pp. 361–372.
- [42] J. Kreimeier, S. Hammer, D. Friedmann, P. Karg, C. Bühner, L. Bankel, and T. Götzelmann, "Evaluation of different types of haptic feedback influencing the task-based presence and performance in virtual reality," in *Proc. 12th ACM Int. Conf. Pervasive Technologies Related to Assistive Environments*, 2019, pp. 289–298.
- [43] P. Wellman and R. D. Howe, "Towards realistic vibrotactile display in virtual environments," in *Proc. ASME Dynamic Systems and Control Division*, vol. 57, no. 2, 1995, pp. 713–718.
- [44] A. M. Okamura, J. T. Dennerlein, and R. D. Howe, "Vibration feedback models for virtual environments," in *Proc. IEEE Int. Conf. Robotics and Automation*, vol. 1, 1998, pp. 674–679.
- [45] G. García-Valle, M. Ferre, J. Breñosa, and D. Vargas, "Evaluation of presence in virtual environments: haptic vest and user's haptic skills," *IEEE Access*, vol. 6, pp. 7224–7233, 2017.
- [46] L.-T. Cheng, R. Kazman, and J. Robinson, "Vibrotactile feedback in delicate virtual reality operations," in *Proc. ACM Int. Conf. Multimedia*, 1997, pp. 243–251.
- [47] I. Herbst and J. Stark, "Comparing force magnitudes by means of vibro-tactile, auditory, and visual feedback," in *IEEE Int. Workshop on Haptic Audio Visual Environments and their Applications*, 2005, pp. 5–pp.
- [48] A. T. Maereg, A. Nagar, D. Reid, and E. L. Secco, "Wearable vibrotactile haptic device for stiffness discrimination during virtual interactions," *Frontiers in Robotics & AI*, vol. 4, p. 42, 2017.
- [49] T. Amemiya, H. Ando, and T. Maeda, "Virtual Force Display: Direction Guidance using Asymmetric Acceleration via Periodic Translational Motion," in *Proc. IEEE World Haptics Conf.*, 2005, pp. 7–10.
- [50] S. Okamoto, H. Nagano, and Y. Yamada, "Psychophysical dimensions of tactile perception of textures," *IEEE Trans. Haptics*, vol. 6, no. 1, pp. 81–93, 2012.
- [51] M. Hollins and S. R. Risner, "Evidence for the duplex theory of tactile texture perception," *Perception & psychophysics*, vol. 62, no. 4, pp. 695–705, 2000.
- [52] K. Ito, S. Okamoto, Y. Yamada, and H. Kajimoto, "Tactile texture display with vibrotactile and electrostatic friction stimuli mixed at appropriate ratio presents better roughness textures," *ACM Trans. Applied Perception*, vol. 16, no. 4, pp. 1–15, 2019.
- [53] S. J. Lederman and M. M. Taylor, "Fingertip force, surface geometry, and the perception of roughness by active touch," *Perception & Psychophysics*, vol. 12, no. 5, pp. 401–408, 1972.
- [54] M. Hollins, S. J. Bensmaïa, and S. Washburn, "Vibrotactile adaptation impairs discrimination of fine, but not coarse, textures," *Somatosensory & motor research*, vol. 18, no. 4, pp. 253–262, 2001.
- [55] C. G. McDonald and K. J. Kuchenbecker, "Dynamic simulation of tool-mediated texture interaction," in *IEEE World Haptics Conf.*, 2013, pp. 307–312.
- [56] V. L. Guruswamy, J. Lang, and W.-S. Lee, "Iir filter models of haptic vibration textures," *IEEE Trans. Instrumentation and Measurement*, vol. 60, no. 1, pp. 93–103, 2010.
- [57] E. Pezent, B. Cambio, and M. K. O'Malley, "Syntacts: Open-Source Software and Hardware for Audio-Controlled Haptics," *IEEE Trans. Haptics*, pp. 1939–1412, 2020.
- [58] D. S. Alles, "Information transmission by phantom sensations," *IEEE transactions on man-machine systems*, vol. 11, no. 1, pp. 85–91, 1970.
- [59] G. von Békésy, "Funneling in the nervous system and its role in loudness and sensation intensity on the skin," *J. of the Acoustical Soc. of America*, vol. 30, no. 5, pp. 399–412, 1958.
- [60] C. A. Perez, C. Holzmann, and E. Sandoval, "Two point vibrotactile spatial resolution as a function of pulse frequency and pulse width," in *Proc. Int. Conf. of the IEEE Engineering in Medicine & Biology Society*, vol. 5, 1998, pp. 2542–2545.
- [61] C. Perez, C. Holzmann, and H. Jaeschke, "Two-point vibrotactile discrimination related to parameters of pulse burst stimulus," *Medical & Biological Engineering & Computing*, vol. 38, no. 1, pp. 74–79, 2000.
- [62] C. E. Sherrick, R. W. Cholewiak, and A. A. Collins, "The localization of low- and high-frequency vibrotactile stimuli," *The Journal of the Acoustical Society of America*, vol. 88, no. 1, pp. 169–179, 1990.
- [63] S. Weinstein, "Intensive and extensive aspects of tactile sensitivity as a function of body part, sex and laterality," *The skin senses*, 1968.
- [64] S. Yang, K. Tippey, and T. K. Ferris, "Exploring the emergent perception of haptic beats from paired vibrotactile presentation," in *Proc. Human Factors and Ergonomics Society*, no. 1, 9 2014, pp. 1716–1720.