

Weirding Haptics: In-Situ Prototyping of Vibrotactile Feedback in Virtual Reality through Vocalization

Donald Degraen*
donald.degraen@dfki.de
Saarland University & DFKI, Saarland
Informatics Campus
Saarbrücken, Germany

Emmanouil Potetsianakis
emmanouil.potetsianakis@telecom-
paristech.fr
TNO
The Hague, Netherlands

Bruno Fruchard*
contact@brunofruchard.com
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany

Seref Güngör
s8seguen@stud.uni-saarland.de
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany

Jürgen Steimle
steimle@cs.uni-saarland.de
Saarland University, Saarland
Informatics Campus
Saarbrücken, Germany

Frederik Smolders
info@polygoat.com
Polygoat
Peer, Belgium

Antonio Krüger
krueger@dfki.de
DFKI, Saarland Informatics Campus
Saarbrücken, Germany

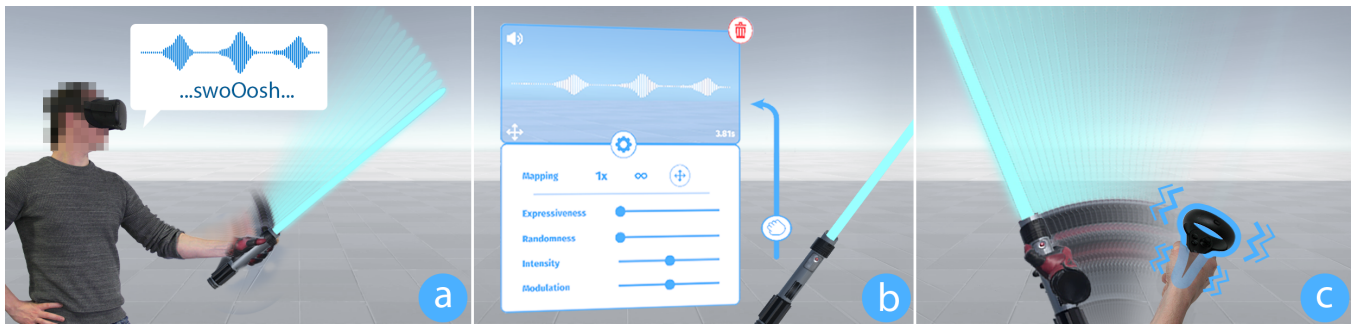


Figure 1: With the Weirding Haptics design tool, one can design vibrotactile feedback in a virtual environment using their voice. a) A designer waves a lightsaber in the air while vocalizing the intended feedback. b) A *vocalization layer* depicts properties of the tactile experience and allows for switching mappings as well as fine-tuning using modifiers (sliders) directly in VR. c) While waving the lightsaber, the vibrotactile feedback can be felt with the VR controller.

ABSTRACT

Effective haptic feedback in virtual reality (VR) is an essential element for creating convincing immersive experiences. To design such feedback, state-of-the-art VR setups provide APIs for programmatically generating controller vibration patterns. While tools for designing vibrotactile feedback keep evolving, they often require expert knowledge and rarely support direct manipulation methods

*Both authors contributed equally to this research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

UIST '21, October 10–14, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8635-7/21/10.

<https://doi.org/10.1145/3472749.3474797>

for mapping feedback to user interactions within the VR environment. To address these challenges, we contribute a novel concept called Weirding Haptics, that supports fast-prototyping by leveraging the user's voice to design such feedback while manipulating virtual objects in-situ. Through a pilot study (N = 9) focusing on how tactile experiences are vocalized during object manipulation, we identify spatio-temporal mappings and supporting features needed to produce intended vocalizations. To study our concept, we built a VR design tool informed by the results of the pilot study. This tool enables users to design tactile experiences using their voice while manipulating objects, provides a set of modifiers for fine-tuning the created experiences in VR, and allows to rapidly compare various experiences by feeling them. Results from a validation study (N = 8) show that novice hapticians can vocalize experiences and refine their designs with the fine-tuning modifiers to match their intentions. We conclude our work by discussing uncovered design implications

for direct manipulation and vocalization of vibrotactile feedback in immersive virtual environments.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**; **Sound-based input / output**; **Virtual reality**; **Human computer interaction (HCI)**.

KEYWORDS

Haptic design, vibrotactile feedback, virtual reality, design tool, voice input, vocalization, direct manipulation

ACM Reference Format:

Donald Degraen, Bruno Fruchard, Frederik Smolders, Emmanouil Potet-sianakis, Seref Güngör, Antonio Krüger, and Jürgen Steimle. 2021. Weirding Haptics: In-Situ Prototyping of Vibrotactile Feedback in Virtual Reality through Vocalization. In *The 34th Annual ACM Symposium on User Interface Software and Technology (UIST '21), October 10–14, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3472749.3474797>

1 INTRODUCTION

Haptic experiences are crucial elements of immersive virtual environments (IVE). Through tactile and kinesthetic feedback, they provide tangibility for visual illusions, increase immersion, support the feeling of presence, and enable users to grasp events experienced in the virtual environment [28, 34, 53]. Previous work underlines the broad variety of potential haptic experiences, including giving tactile feedback through passive physical props [9, 14] or actuated mechanisms [20], rendering kinesthetic feedback [58, 59], and even simulating haptic textures and materials using vibrotactile feedback [12, 23, 33, 46, 47, 54].

However, developing effective and convincing tactile experiences using vibrotactile feedback remains a challenge. State-of-the-art design tools (e.g., [15, 25, 39, 47]) propose to manipulate low-level controllable parameters such as frequency and amplitude, but it remains challenging to transfer such abstract parameters into understandable haptic effects [36, 43]. Moreover, these design tools rarely support fast-prototyping methods nor do they support direct and easy mapping of vibrotactile feedback to users' (spatio-temporal) interactions in VR [19]. Implementing convincing experiences is an even greater challenge for those inexperienced in haptics, e.g., video game programmers who seek to design playful experiences with tactile sensations, students learning haptics through prototyping, or interaction designers wanting to provide tactile feedback in UI widgets [42]. As pointed out by recent work [19, 42], novice hapticians need more *timely, hands-on* interfaces leveraging *direct manipulation* to better grasp the experiences they want to design.

We contribute Weirding Haptics¹, a novel concept for in-situ rapid prototyping of vibrotactile feedback in VR environments. It combines fast and expressive vocalizations with the ease and directness of interaction with virtual objects, to offer a streamlined process for prototyping tactile experiences of virtual objects. Compared to existing approaches, using the voice enables designers to vocalize vibrotactile feedback while interacting with virtual objects inside

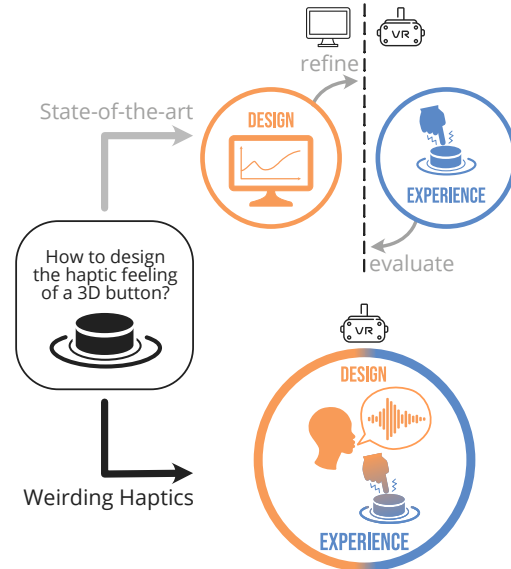


Figure 2: To design the tactile experience of a virtual 3D button, state-of-the-art design tools require designers to go back and forth between designing low-level parameters on a desktop computer and experiencing tactile feedback in VR. In contrast, Weirding Haptics enables designing *inside* the VR environment through vocalizations and supports fast-prototyping through a rapid iterative process.

VR (Figure 2). Moreover, designing with the voice does not require expert knowledge in haptic design. In this work, we answer two main challenges regarding this concept. First, as we directly map vibrotactile feedback to interactions with objects, we identify what kind of spatio-temporal mappings designers need to create convincing haptic experiences that are directly linked to object interactions. Additionally, as we want to assist untrained voices and support a rapid design cycle while immersed in the IVE, we identify modifiers that enable designers to rapidly fine-tune the output generated by their vocalizations in real-time.

Informed by these insights, we contribute an implementation of Weirding Haptics in a VR design tool that enables in-situ fast-prototyping of vibrotactile experiences using vocalizations. Using this tool, one can synchronously vocalize the intended vibrotactile experience of a virtual object during in-situ interaction with objects in VR (see Figure 1a). After sampling frequency and amplitude, the design tool infers how to map these vocalizations to vibrations in space and time based on the designer's interactions. Moreover, the designer can control properties of the tactile experience through a *vocalization layer* inside the VR environment (see Figure 1b). With this layer, the designer can switch between different spatio-temporal mappings and fine-tune the experience in real-time using modifiers. Changes are immediate, the designer can quickly experience the vibrotactile feedback and assess whether it matches their original intention (see Figure 1c). Several vocalization layers can be combined to stack different effects (e.g., background sensation with bursts overlaid) or enable quick comparison. To the best of our

¹Inspired by the weirding module, an object controlled by vocalizing an intention, from the 1984 Dune movie - https://dune.fandom.com/wiki/Weirding_Module

knowledge, this design tool is the first to support direct interaction with objects synchronized with vocalizations to design vibrotactile feedback in-situ.

Through a validation study involving a set of virtual objects, we demonstrate how Weirding Haptics, with only a short training period (~ 15 minutes), supports novice hapticians in designing experiences that match their intentions. Participants designed effective illusions of sand flowing or a rock tumbling inside wooden boxes, a slider providing resistance, surface textures with different roughness sensations, waving a lightsaber in the air, or the realistic sensation of opening and closing metallic and wooden drawers. The results of our study uncovered design implications concerning the in-situ design of vibrotactile feedback with the voice. Future tools must balance the level of fidelity designers require while supporting a fast-prototyping approach, and need to support high spatio-temporal resolutions for synchronizing vocalizations with the user's interactions. We discuss these challenges and conclude that using the voice to design tactile experiences in VR enables novice hapticians to create vibrotactile feedback for virtual objects aligned with their intentions.

2 RELATED WORK

Our contributions relate to prior work on leveraging vibrotactile experiences inside and outside VR, on vocalization of such experiences, and on haptic design and more specifically on design tools in research and industry for producing vibrotactile experiences.

2.1 Vibrotactile Feedback for Virtual Experiences

Haptic feedback provides a sense of tactility in an immersive virtual environment to enhance the experience [45]. Through touch, a user's sense of presence inside the virtual environment is enhanced [14]. Such experiences leverage either passive proxy objects used as physical props [9, 14], active mechanisms such as vibrotactile actuation [23], or mixed approaches aimed at combining the best of both worlds [8, 20, 59]. In this work, we focus exclusively on active vibrotactile feedback.

The pacinian corpuscles, nerve endings situated in the skin, enable us to feel very light vibrations when interacting with objects [3]. When lifting an object, for instance, our skin deforms to adapt to its surface; the perceived vibrations provide information about the object's weight [16]. Moreover, vibrations play an important role in stiffness perception [54]. By controlling the vibrotactile impulses based on the user's actions, one can create various tactile experiences. For instance, one can induce various mechanical properties, making a rigid object feel compliant and deformable [18], and create the illusion of squeezing or shearing an object inside a virtual environment [23]. Using tool mediated approaches such as the use of a stylus, one can also use vibrotactile actuation to create the haptic illusion of compliance [17], or virtual texture exploration [33]. An important aspect of these experiences relies in the way they are designed; through variations in the frequency and amplitude of the produced vibrations, a large variety of tactile illusions can be created [29, 46].

To create convincing vibrotactile experiences, it is mandatory to closely relate them to the user's actions with the object. Strohmeier et al. [46] showed how different mappings between the vibrations and the user's hand movements can impact the perception of the

experiences. Authors used grain-based vibrations, i.e., pulses that happen at certain frequencies based on the user's input. A similar approach was used in Barefoot [47] to create virtual materials below the user's foot. In this case, authors mapped vibrations to the pressure applied on the ground. In terms of multi-modal rendering, virtual walking experiences greatly benefit from the addition of vibrotactile perceptual cues [26]. Heo et al. [12] also leveraged this approach to simulate bending, twisting, and shearing a physical rod in VR.

While approaches in literature often use custom designed actuators, commodity VR controllers often prioritize ease of use over fine grained control. Limitations such as the Oculus SDK's non-buffered haptics [30] which constraints the range of frequency and amplitude, add an extra challenge for designers to create effective tactile effects for virtual experiences. With such restrictions, high fidelity design is translated to a lower fidelity space. As this might influence user experience, in this work we investigate output differences in two modalities, i.e., a commodity VR controller, and soon to be conventional high fidelity actuator.

2.2 Vocalizing and Rendering Tactile Experiences

Associations between language sounds or vocal expressions and perceptual or semantic features have long been understood under the term of sound symbolism [24, 44]. Ideophones are words that depict sensory imagery and cover a wide range of domains, such as motion, texture, and even psychological states [1]. Onomatopoeia are well-known examples and can be found in comics to use language for expressing sounds [11]. As a phonic modality of speech, onomatopoeia and mimetic words are able to communicate perceptual qualities such as the visual appearance of metal textures [57] or tactile sensations [38, 55] and can even be used to transfer embodied expertise [13]. Moreover, iconic vocalizations have been shown to describe tactile sensations [38, 55], and to ground and communicate design intention [2, 5]. While iconic vocalizations show great potential for expressing tactile impressions, they are highly dependent on the cultural background of the speaker [35]. To abstract from these dependencies and provide a universal tool, our work focuses specifically on the acoustic properties of vocalizations.

In the field of haptic rendering, acoustic properties of audio recordings, such as frequency and amplitude, are often used to generate convincing effects. In terms of vibrotactile rendering, features of audio signals were used to generate haptic feedback for immersive interaction with the Haptic Cushion [6]. Lee and Choi [22] proposed two real-time audio-to-vibrotactile translation models for enhancing visual and auditory media content, by extracting perception-level metrics from audio signals (i.e., loudness and roughness). In our work, we extract audio features, i.e., frequency and amplitude, from users' vocalizations to output them as vibrations. We envision these acoustic properties of vocalizations as a means for rich, naturalistic and improvisational design of tactile experiences. While Lee and Choi [22] propose a precise automated system to translate perception-level audio features into vibration features, our approach includes users in the translation loop by enabling them to control the translation pipeline with fine-tuning modifiers.

Our concept is inspired by vocal sketching techniques, such as the work of Rocchesso et al. [32] for supporting sound designers and the work of Marino et al. [27]. In the latter, authors proposed to

leverage the user’s voice to imbue a one Degree-of-Freedom actuated robot with emotions by matching vocal features to specific motions. Through co-design studies, they found important parameters users required to design robot motions and create an “alive” feeling. We build on this approach by leveraging the user’s voice to design vibrotactile feedback while manipulating virtual objects to infer spatio-temporal mappings.

2.3 Haptic Design

Designing haptic experiences remains a challenging task. Schneider et al. [36] provide insights into the field of haptic experience design (HaXD). As HaXD becomes part of many designers’ jobs, their work builds an understanding of existing design processes and its specific challenges. Kim and Schneider [19] recently identified design parameters expert hapticians deal with to produce effective experiences (timeliness, intensity, density, and timbre), and outlined the importance of supporting personalization of these experiences by end-users. Similarly, Seifi et al. [42] observed novice hapticians designing haptic experiences. They conclude that haptic design is lacking support at different stages of the design process and that tools to sketch haptics need to provide real-time feedback and enable direct manipulation.

In recent years, novel design tools have emerged that support a large panel of hapticians designing haptic experiences in various contexts. These tools address diverse aspects of the design process: these include designing experiences by manipulating frequency and amplitude curves through time [39], controlling at a low-level how vibrations are mapped to the user’s actions [47], or sharing experiences through visualizations with others [40, 47, 49]. Inspired from other designing approaches, Swindells et al. [50] proposed to design vibrotactile feedback in synchronicity with video by editing haptic events directly in a video editor, and Schneider et al. [37] proposed to design vibrotactile feedback through animations. While previous work proposes more hands-on design approaches such as using instrument-like devices to design vibrotactile feedback [38, 56], state-of-the-art tools remain focused on manipulating abstract parameters [15, 25]. Furthermore, despite some of these tools’ focus on designing experiences for VR [15], they do not support designing inside the virtual environment nor support rapid prototyping.

Our concept and design tool supports the prototyping stage of vibrotactile experiences by providing real-time tactile rendering of users’ designs in VR and enabling them to quickly iterate through different designs.

3 PILOT STUDY ON VOCALIZATION OF TACTILE EXPERIENCES

In order to better understand the efficiency and limitations of vocalizing tactile experiences, we performed a pilot study with nine novice hapticians² in our lab.

3.1 Study Design

While interacting with common physical objects outside of an IVE, we asked participants to vocally express their tactile impressions in as many ways as they saw fit. The goal of this study was to

provide insights into (1) spatial and temporal relationships between vocalizations and the user’s actions, and (2) how a computational tool can support novice hapticians to vocalize tactile experiences. Ethical approval for this study was obtained from the Ethical Review Board of the Department of Computer Sciences at Saarland University (No. 21-01-3).

Experimental Design and Apparatus. Using an iterative process, we aimed to identify parameters that would influence users’ vocalizations. To this aim, rather than selecting a wide range of shapes, we focused on the manner in which the object would be interacted with and included variations in size and texture, as these parameters can highly influence the type of vibrations produced during manipulation. Here, we identified three primary features important to our context, i.e., the *type of action* performed, whether the object is *directly manipulated* or is used as a *tool*, and the *size* of the manipulated object. We consider the two former as independent variables (*Action* and *Tool*) and the latter as a control variable (*Size*). The list of actions, inspired from the haptic exploration procedures by Lederman and Klatzky [21], is as follows: *slide*, *pull*, *push*, *press*, *rotate*, *passive feel*. By crossing the two independent variables, we obtain a set of 10 categories; most tools could be *pushed* and *pulled* so we merge these actions, and the action *passive feel* includes a tool by definition. Each category included a set of 3 to 6 object-action pairs. In total, we evaluated a set of 44 objects that present various characteristics (size, texture, weight), e.g., fabric samples, knobs and buttons offering various levels of resistance, sponges and elastic bands, or an electrical toothbrush, see Appendix A. While we tried to cover a broad range of object-action tasks to include tactile experiences produced by object deformation (e.g., stretching an elastic band or squeezing a sponge), object actuation (e.g., sliding a camera trolley), or the object’s texture (e.g., rubbing fabric samples with various roughness), this list is not exhaustive.

We counterbalanced the categories to avoid any order effect. To allow simple comparisons between similar objects like fabric samples with various roughness or elastic bands with various stiffness, we always use the same order of object-action tasks for a given category.

Participants. We recruited nine novice hapticians (2 identified as female, 7 identified as male) aged between 23 and 35 (median 29) with backgrounds in Computer Science, Media Informatics and Linguistics. Participants had diverse cultural backgrounds and various native tongues such as English, Ukrainian, Russian, Chinese, Hindi, Farsi, and French. Seven of nine (78%) participants indicated to have a background in musical training, while three (33%) had prior experience in voice acting or singing.

Procedure. Before starting the experiment, participants completed a short warm-up task to stimulate their creative skills which consisted in producing as many animal noises as possible under one minute. Afterwards, each participant proceeded with the vocalization tasks. For each task, we asked participants to perform a single action with the object as many time as they wanted. When ready, they were asked to vocalize the tactile experience. In pre-pilots, we noticed participants reproduced the sound that objects would make during manipulation. To avoid confusion and ensure participants would focus on tactile sensations, we instructed them to focus primarily on the tactile sensation while producing vocalizations. However, we did not forbid them to reproduce those noises if they felt the

²we follow the same definition for “novice haptician” as Seifi et al. [42]

sensation matched the noises produced by the objects. We motivated participants to provide as many vocalizations as they could come up with. Should a participant not be able to provide any vocalization, the experimenter would move on to the next task. However, if a participant could not produce the vocalizations they intended due to physical constraints (e.g., frequency too high to produce), we asked them to explain as clearly as possible what they were missing.

Analysis. We video recorded each session and analyzed them using a thematic analysis approach, following an inductive process [4]. Our focus was to observe in detail how participants mapped these vocalizations to their interactions with the objects, and where they would need features their voice could not produce. Data consisted of short video clips of participants vocalizing tactile feedback while manipulating an object and occasional remarks. In a first round, the two first authors watched the video recordings and coded the temporal and spatial features of the vocalizations (instantaneous, repetitive, continuous, interaction bound, random) as well as the kind of sounds produced (pitch-based, amplitude-based, onomatopoeia, etc.). In a second round, they refined the set of codes used to reach agreement, and used it to generate a set of themes.

3.2 Results

From our analysis, we identified four themes of spatio-temporal mappings between the vocalization and the action performed. We complement these themes with a set of challenges participants faced during vocalization.

3.2.1 Spatio-Temporal Mappings. Participants mapped their vocalizations to specific events during the object manipulations, relating to either time or space:

Instantaneous. An instantaneous vocalization described short events in the tactile experiences such as flicking a marble, pressing a button, or closing scissors. This was sometimes combined with *continuous* vocalizations as a way to express a stronger signal, like a bump when reaching the end of a rotatory knob. All participants produced such a mapping at least once during the experiment.

Repetitive. A repetitive mapping consists of a sequence of similar instantaneous vocalizations. A repetitive vocalization strongly relates to the speed of the user's action as its frequency increases or decreases accordingly. All participants used repetitive mappings during the experiment to describe bursts in the tactile experience, such as a knob producing distinct positional clicks during rotation.

Continuous. A continuous mapping represents a smooth, invariable experience while performing an action with an object. Participants unanimously used such a mapping when passively feeling the vibrations of an active electric toothbrush. A majority of participants also used such mappings while sliding a coin over a table or sliding their finger on smooth or rough uniform pieces of fabric.

Action-related. While all mappings relate to the actions performed, some mappings are tightly connected to the movement or the force exerted on the object. By varying their pitch or volume based on the change in tactile sensations, participants used these mappings when squeezing or pulling deformable objects (e.g., sponge or elastic band), or when sliding actuated parts of a tool (e.g., sliding a bike pump or a drawer).

3.2.2 Challenges in Producing Vocalizations. Participants sometimes faced difficulties producing vocalizations. For instance, a participant remarked “*the pitch [for the hard sponge] should be higher than [the pitch for] the soft sponge*” (P6) and “*I would like a more grainy voice*” when pressing a finger on a rough sponge. Another participant produced a vocalization and specified it should be “*with a higher pitch, very high*” (P3). In general, participants acknowledged the complexity of the tactile experience and the various layers they comprehend. In this regard, one participant (P6) made interesting remarks while vocalizing the rotation of a stepper motor. They wanted to superimpose two different “*tracks*”: a repetitive pattern and “*some random stuff*”. While this is challenging in terms of vocalizations, a design tool could provide support here.

3.3 Design Implications

The themes extracted from this study alongside participants' comments enable us to infer design implications for design tools implementing Weirding Haptics.

1. *Infer spatio-temporal mappings and enable designers to switch between them.* We observed four distinct spatio-temporal mappings adopted by the participants to vocalize tactile experiences. A design tool should infer such mappings and let users decide upon them in a fine-tuning stage of the design process. For example, a system might infer an *instantaneous* mapping based on the user's actions, while they actually intended to use a *repetitive* mapping with the same vocalization (e.g., pouring water out of a bottle).

2. *Adapt vocalizations to the user's movement speed.* To not bias participants' natural ways to manipulate an object, we intentionally did not constrain the speed of their actions. This speed, however, has a great impact on the vocalization process for repetitive and action-based mappings; stretching an elastic band rapidly or slowly results in the elastic band wiggling or not. Features of a vocalization should adapt based on the speed of the action performed.

3. *Enable fine-tuning vocalization features and support untrained voices.* Vocal skills of users are limited, particularly when their voice is not trained. Our observations showed the necessity to provide such virtual augmentations in various scenarios. For instance, participants needed to produce random patterns to generate noise that would better match the tactile experiences, or to increase their frequency range at multiple occasion. A design tool should provide fine-tuning functionalities that one can use to compensate for imprecise vocalizations and provide computer-supported functionalities.

4. *Instantiate vocalizations as layers to enable combining them.* Our analysis showed users may need to decouple their voice in several layers to reach a desired experience. This is a common process in audio producing to combine various effects and set unique properties to each. Users should be able to iterate over the design of a vibrotactile experience and stack various layers together to create a complete tactile experience.

4 DESIGNING VIBROTACTILE FEEDBACK IN VR

Based on the design implications from our pilot study, we created the Weirding Haptics design tool, which transforms users' vocalizations into vibrotactile feedback inside a virtual environment. Designers can rapidly record vocalization while interacting with virtual objects




Figure 3: Core principles of the Weirding Haptics design tool. Through a *in-situ design* approach, designers can vocalize vibrotactile feedback while interacting with virtual objects in VR. The *spatio-temporal mapping* is inferred based on the movement magnitude and recording duration. To support rapid iterative design, vibrotactile feedback can be fine-tuned using modifiers.

inside VR and iterate over different designs quickly. To design effective vibrotactile feedback, the design tool takes an *in-situ design* approach based on direct manipulation methods, infers *spatio-temporal mappings* based on object interactions, and enables *fine-tuning* output from vocalizations through real-time *modifiers* (see Figure 3). The Weirding Haptics design tool is built on top of the Unity³ game engine and is available here: <https://github.com/darty/wh>. We used the SteamVR Unity plugin [51] to make this tool adaptable to various VR setups. We designed it with the Oculus Quest 2 setup.

4.1 In-Situ Design Process

Our aim is to enable the designer to easily and rapidly prototype expressive vibrotactile feedback for virtual object interactions. The design tool builds on vocalizations, i.e., audio signals produced by the user’s voice, recorded while directly manipulating objects inside the virtual environment. From these recordings, we extract audio features, i.e., frequency and amplitude, map them to vibrotactile feedback, and enable the designer to control their processing during the playback pipeline.

To normalize the range of vibrotactile feedback everyone can produce, the design tool allows designers to calibrate their voice with two simple tasks to set the bounds of their vocal range in terms of amplitude and frequency. The tool only considers normalized features in the playback pipeline.

In the current version, we focus on the design of haptic feedback for interaction with single objects. To avoid conflicts between recording and experiencing vibrotactile feedback, we leverage a bi-manual interaction design with the non-dominant hand controlling the context and the dominant hand interacting with objects [10]. The non-dominant hand is used to arm and possibly stop the recording process, while the dominant hand is used to perform interactions with objects inside the scene.

The recording process starts once the user has armed the recording and starts directly when manipulating an object. It ends as soon as the user stops interacting with the object, or releases the trigger used to arm the recording. The vocalization is automatically mapped to the respective object and interaction, and immediately visually

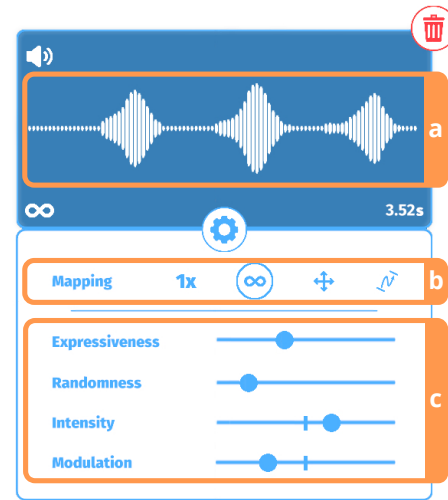


Figure 4: Graphical representation of a vocalization layer. a) The *vocalization* is represented as a stylized waveform. b) The designer can switch between different *spatio-temporal mappings* to try out different experiences, c) or fine-tune the experience by manipulating the *modifiers* using sliders.

depicted in the VR scene as a *vocalization layer* (see Figure 4). The designer can directly experience the vibrotactile output by again performing the interaction with the object. To support fast prototyping of vibrotactile designs, vocalizations can easily be added, removed and fine-tuned.

Multiple vocalization layers can be superimposed on an object, to create more complex vibrotactile feedback of multiple discrete vocalizations. For instance, a designer could design the continuous humming vibrations of a lightsaber as one layer, and more intense bursts when waving it as another. Superimposed layers are experienced synchronously during interaction by compounding them and playing the maximums of their amplitudes and frequencies. While recording new vocalizations, all vibrotactile feedback from already existing layers for the same object interactions can be felt to ensure alignment.

4.1.1 Implementation. Once a vocalization is recorded, the system will optionally timescale the vocalization, and pre-process the signal by sampling frequency and amplitude.

Time-scaling. In order to support vocalizations that are directly linked to positions in space, it is important to produce the correct vibrations at defined landmarks. As the speed of interaction when experiencing the feedback is unknown at the time of recording, we normalize the audio recording based on the recording speed. Therefore, for positional recordings, we record the user’s hand position through time on a one-dimensional line segment. Using this information, we apply a non-linear time-scaling algorithm using the python Rubber Band Library [31]. The time-scaling ensures signal features remain correctly aligned to the position at which they were recorded. This phase is only applied once for objects able to receive a positional mapping.

³Unity Real-Time Development Platform - <https://unity.com/>

Sampling. As processing delays negatively influence the timeliness of the provided feedback, the system pre-processes recorded vocalizations to extract frequency and amplitude. For amplitude extraction, we first calculate the envelope of the signal⁴ and sample the result in 50ms intervals. Frequency extraction is done directly on the signal using a Yin pitch recognizer [7] with 100ms intervals. The interval for frequency sampling is longer than the interval of the amplitude sampling, as frequency sampling on shorter signals provides incorrect or no results.

4.2 Spatio-Temporal Mappings

Each vocalization layer is assigned a spatio-temporal mapping which defines how the signal is sequenced and repeated in relation to the interaction. Based on the insights gained from the pilot study, we distinguish between four different types of mappings, i.e., instantaneous, continuous, motion, and positional.

With an *instantaneous* mapping, the vocalization is experienced exactly once. This mapping supports, for example, designing the tactile experience of tapping on a surface or pressing a button. Using the *continuous* mapping, the layer is played repetitively as long as the user interacts with the object. This mapping supports, for example, constant humming vibrations produced by an active object like an electrical toothbrush. Similar to the continuous mapping, the *motion* mapping will keep iterating over the vocalization as long as the user is interacting with the object, such as moving while grabbing or touching the object. In this case, however, we multiply the vocalization amplitude and frequency by the designer’s movement velocity. As a consequence, the designer will not feel the vocalization when idling, while the intensity will increase with speed. This mapping supports, for example, exploring surface textures or particles moving inside a container. The *positional* mapping is used to map parts of the vocalization at given spatial positions. This mapping supports, for example, interacting with a drawer that provides various resistance throughout its path or when opening a door that creaks at given landmarks.




Figure 5: The design tool infers the spatio-temporal mapping of a vocalization directly after recording by considering the user’s movement magnitude and the recording duration. D (0.5 m) and T (2.5 s) are empirically determined constants.

4.2.1 Implementation. After recording a vocalization, the spatio-temporal mapping is inferred based on the user’s movement magnitude while manipulating the object as well as the recording duration. We assess the movement magnitude based on the SteamVR Unity

⁴we implemented the Shockley diode algorithm used in the audio_dspy python library - https://github.com/jatinchowdhury18/audio_dspy




Figure 6: The designer can interact with the object to evaluate a vibrotactile feedback (left), and leverage the rapid prototyping cycle provided by the system to fine-tune it (right).

plugin’s controller velocity for free movement and angular velocity for rotational interaction. Our current implementation infers the mapping based on two thresholds established during informal testing of the system: $D = 0.5$ m for the magnitude, and $T = 2.5$ s for the duration (Figure 5). Of note, our concept is not limited to this approach and compatible with future, more advanced algorithms for automatically inferring the mapping.

When interacting with the object, the features of the recorded signal are extracted during the sampling phase. Based on the spatio-temporal mapping assigned to the vocalization layer, the sampling location within the signal is determined. For all the mappings excluding the positional one, the system samples the signal based on time. For the positional mapping, the system samples the signal based on the current position of the user’s hand on a line segment. This limitation was chosen to lower the complexity of the mapping between recorded signal and the generated output. Our framework easily allows us to extend the positional mapping to any segment in 3D space. Once the sampling interval within the signal is determined, we get the amplitude (A_{raw}) and frequency (F_{raw}) for the current interval. We then normalize the vocalization features in real-time based on the designer’s calibration. Once the features are normalized, we obtain A_{norm} and F_{norm} .

4.3 Rapid Iterations with Fine-Tuning Modifiers

Enabling fine-tuning of vocalization layers is crucial in our approach: it allows tweaking subtle parameters of the vibrotactile feedback that would be hard to control with the voice, supports designers with computer-generated modifications, and supports untrained designers’ voices. We propose modifiers for simple and direct fine-tuning of a vibrotactile experience, by adjusting sliders on the vocalization layer (see Figure 4c and Figure 6). When modifying the vibrotactile feedback with one of the modifiers, it is updated in real-time such that the designer can directly feel the changes.

The design tool proposes an *intensity* and *modulation* modifier, respectively modifying the amplitude and frequency of the vocalization. The former is particularly useful to amplify or dampen a vocalization when it does not match the initial intention of the designer, or to balance several vocalization layers. The latter is useful to change how the vibrotactile feedback can be perceived, as shown by previous work [29].

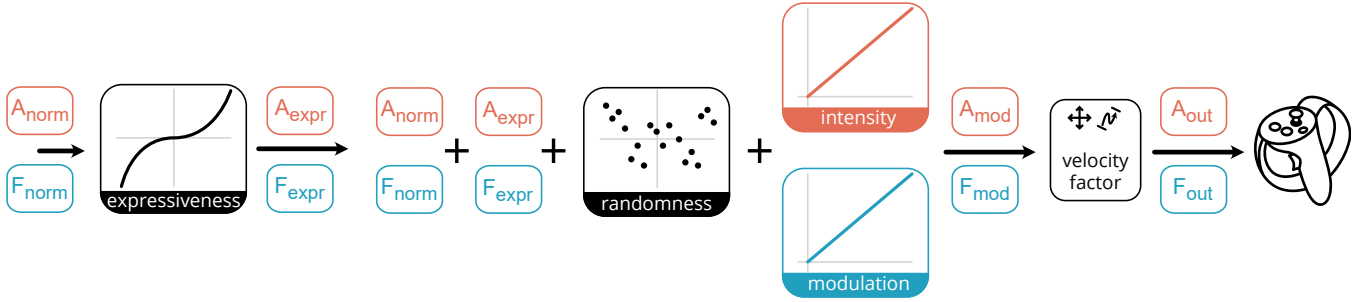


Figure 7: The playback pipeline. Based on the normalized samples produced by the mapping used, we first compute the *expressiveness*, then add the *randomness*, and later apply the *intensity* and *modulation* respectively. The resulting features are multiplied by the relative velocity of the user’s movement for motion or positional mappings.

In the pilot study, some participants remarked they would like to create random patterns for certain tactile experiences they could not produce directly with their voice. Previous work showed the interest of such patterns when generating virtual material to “evoke natural experiences” [47], as well as when designing robotic movements to create more natural behaviors [27]. Therefore, the design tool proposes a *randomness* modifier to introduce noise in the feedback.

Lastly, the system provides an *expressiveness* modifier that enables contrasting peaks in the signal. Based on the pilot study, we noted participants repeating certain onomatopoeia, e.g., saying “tick”, to describe singular bursts in the tactile sensation. Such bursts can be smoothed out based on the vocalization rhythm and the tool frame rate, thus the expressiveness modifier enables designers to control their attack and intensity directly, whereas boosting up the intensity would uniformly change the vocalization.

4.3.1 Implementation. Each modifier controls a given variable (E for expressiveness, R for randomness, I for intensity, and M for modulation) that has a unique impact on the output signal. We detail how we take into account these variables in the following, along with explanations about how the vocalization-to-vibration algorithm iterates for each update cycle (50ms update rate), see Figure 7.

Modifiers. We first apply the *expressiveness* modifier. It builds on a sigmoid function using the normalized features to contrast peaks in the vocalizations: $A_{expr} = A_{norm} \times \left(\frac{2}{1 + e^{-E(A_{norm} - A_{median})}} - 1 \right) \times e$ for $E \in [0, 10]$, e an empirically informed constant, and A_{median} being the median between the minimum and maximum amplitude of the current vocalization. We apply the same procedure for the frequency. Both the amplitude and frequency results are added to the normalized features. We then add the *randomness* modifier, i.e., we add a random value $R \times r$ with $R \in [0, 1]$ and $r \in [-0.5, 0.5]$. As a last step, we add the intensity value $I \in [0, 1]$ to the amplitude, and the modulation value $M \in [0, 1]$ to the frequency.

Velocity Factor. For motion and positional mappings, once all modifiers have been applied to the amplitude and frequency values, we multiply the ratio between the velocity of the user’s action and the average recording velocity with A_{mod} and F_{mod} .

After all the processing stages, the amplitude (A_{out}) and frequency (F_{out}) values are ready to be sent to the controller for actuation.

4.4 Haptic Output

Varying both frequency and amplitude is essential to build a large gamut of tactile experiences [29, 46, 47]. While the SteamVR plugin proposes to input both frequency and amplitude to control VR controller vibrations, it seems to not correctly control these two variables for all commodity controllers. To verify this, we performed hardware tests by placing an IMU on the Oculus controller and swiping through frequencies using the same amplitude, which confirmed changing frequencies to have no effect. To tackle this issue, we devised two possible solutions. Firstly, for commodity controllers with frequency limitations, we compound the F_{out} result as a factor to the A_{out} , $A_{out} = A_{out} + (F_{out} \times 2 - 1) \times f$, $f = 0.25$ being an empirically informed constant. This method compensates for the limitations by varying vibrations based on frequency modulation.

Secondly, to avoid restraining our vocalization approach with a low-resolution output, we built an alternative hardware setup for more advanced vibrotactile rendering that can be attached to the VR controller. We use a similar actuator (AFT14A903A) as the one found in Oculus Touch controllers⁵. To compensate for natural resonant frequencies inherent to such a device, we normalize the amplitude response of frequencies between 150 Hz to 300 Hz to produce a 1G output approximation. We chose this frequency range as pacinian cells are most sensitive to ranges around 220Hz [29, 47, 48, 52], but the design tool can adapt to any kind of frequency range as it solely deals with relative values. The final device is controlled by an ESP32-DevKitC V4 microcontroller board, connected via USB. The input frequency, amplitude and duration values sent by the host are used to create fitting sine waves to avoid artifacts created by switching frequencies too abruptly. A digital potentiometer (AD5280BRUZ50) applies the amplitude to the signal, and a Class D amplifier (PAM8403) amplifies the result for playback by the actuator. We attach this *custom* actuator on the Oculus controller and directly use the A_{out} and F_{out} values for vibrotactile actuation.

5 ASSESSING THE USABILITY OF WEIRDING HAPTICS

We conducted a user study with novice hapticians in which they designed vibrotactile sensations inside VR using their voice.

⁵<https://www.ifixit.com/Teardown/Nintendo+Switch+Teardown/78263>




Figure 8: Virtual objects used in the study to inspire participants to design vibrotactile feedback: (a) boxes of different sizes, (b) surface textures, (c) slider, (d) slider with textures on sliding rod, (e) medieval sword and lightsaber, (f) textured walls, and (g) metallic and (h) wooden drawers.

5.1 Study Design

The focus of this study was to assess the usability of the Weirding Haptics concept and design tool, and better understand the research challenges related to in-situ design of haptic feedback. We used a think-aloud process with open-ended tasks. Ethical approval for this study was obtained from the Ethical Review Board of the Department of Computer Sciences at Saarland University (No. 21-03-9).

Participants. We recruited a total of 8 novice hapticians (4 identified as male, 4 identified as female) aged between 23 and 33 (avg. 27) with backgrounds in Computer Science, Media, Microbiology and Linguistics. Three of these participants had participated in the pilot study, while 5 were newly recruited⁶. One participant indicated to regularly use VR for work purposes, 4 participants indicated to have experienced VR a few times, 1 participant indicated to have experienced VR only once, while 2 participants did not have any VR experience at all. All participants indicated to be novices in the field of haptic design. All participants provided consent to record both video and audio during their session.

Apparatus. To provide participants with a varying and attractive virtual environment, we created a scene with 6 different types of objects, see Figure 8. These objects were partly inspired from the pilot study, leveraged various types of interactions, and provided different appearances, thus could convey various properties such as roughness, weight, or uniformity. Three surface textures and two walls were stationary and allowed the user to design vibrotactile feedback to touch events. Two sliders with different designs and two drawers of varying materials invited participants to design positional mappings. Lastly, three boxes of different sizes and two types of swords supported vibrotactile feedback for touch interaction or for moving them. Vibrotactile feedback was first provided with the native actuators used in the Oculus Quest 2 controllers, followed by the custom actuator. Their comparison was aimed at understanding the effect of high-resolution frequencies on the vocalization design process.

Procedure. Before starting the experiment, participants were explained the concept of designing haptics through vocalizations. Using an example of a vocalization layer, see Figure 4, we explained the procedure of creating, experiencing, fine-tuning and deleting layers. After this introduction, participants were asked to attach a clip-on microphone and enter the virtual environment by putting on the Oculus Quest 2 HMD.

Once inside the virtual environment, a training scene was loaded. Here, participants calibrated the framework to their vocal range. We then introduced participants to the controls and guided them in creating a vibrotactile design for touching a surface and for grabbing a box. As soon as participants felt they could control the system efficiently, the next scene was loaded. On average, this training task lasted around 14 minutes ($\sigma = 2$ min).

The study scene consisted of all the objects depicted in Figure 8. We asked participants to vocalize haptic experiences for these objects in a given sequence. To help them to get inspired, we started the tasks with the three boxes and gave specific instructions for each of them: we tasked participants to create the illusion of sand flowing, pebbles moving, or a rock tumbling inside the small, medium, and large box respectively. For the other objects, they were free to interact with each object as they saw fit and were allowed to create, fine-tune and delete as many experiences as they wanted. Their initial designs were first experienced with the native controller actuators for each object. Once they were satisfied with the result, participants were asked to rate their creations based on their *initial vocalization intention* on a 7-point scale, where 1 meant “not as intended”, and 7 meant a “perfect fit”. If something was not aligned with how they wanted it to be, participants were asked to elaborate on which aspects of the resulting experience felt close or near to their intentions. After each rating, the haptic output was provided using the custom actuator. Here, participants were able to experience their designs again and were free to fine-tune or even re-design the experience if necessary. Again, they were asked to rate and elaborate on the experience based on their intention. After completing the experiment, participants filled out a demographics form, and were asked about their experience of designing vibrotactile feedback using vocalizations.

On average the study lasted 84 min ($\sigma = 14$) and breaks were issued at the halfway point, or upon request by the participant.

⁶This partial overlap was caused by the difficulties of recruiting participants during the COVID-19 pandemic. Due to the time in-between both studies (2 months) and their difference in nature, we did not find a reason for this to bias participants' performance.




Figure 9: Occurrences of the participants' ratings for each object used in the validation study.

5.2 Results

After the experiment, the video recordings and remarks provided by participants were analyzed. Based on our observations, we assess the usability of our approach, reflect on the impact of the Weirding Haptics core principles on the tactile experiences created, and discuss the limitations of our current approach for designing vibrotactile feedback in VR.

5.2.1 In-Situ Design of Vibrotactile Feedback. After a relatively short training session, all participants were able to design and fine-tune vibrotactile experiences regardless of their previous experience with VR. On average, the designed feedback was indicated to be close to their intention ($\bar{x} = 5.49$), see Figure 9. The highest average rating was given for feedback designed for the lightsaber ($\bar{x} = 6.25$), followed closely by the textured slider ($\bar{x} = 6.13$), the illusion of a rock inside the large box ($\bar{x} = 6.13$), and opening or closing the wooden drawer ($\bar{x} = 6.06$). The lowest ratings were given to the designed feedback of touching the right wall ($\bar{x} = 4.63$), and the right texture ($\bar{x} = 4.21$). Participants commented that the initial experience felt “surprisingly good” (P3) and rated their overall experience positively (P4: “it was fun”; P8: “it’s really cool, it’s fun”). Interestingly, we observed participants being sometimes satisfied by the result of a vocalization on the very first iteration, i.e., recording one vocalization that would instantly match their intention (Figure 10a). This indicates novice hapticians can produce vocalizations that are accurate enough to match their intention using an in-situ design approach.

On average, participants used 1.11 layers ($\sigma = 0.37$) for each tactile experience. Most of the designed tactile experiences consisted of a single vocalization layer (81.21%), while the use of 2 (6.84%) or even 3 layers (1.71%) was infrequent. While some participants intuitively used this feature to enhance experiences, others needed to be reminded of its availability or we had to explain to them how a certain effect they described could be implemented using this feature. Designs using multiple layers were mostly aimed at adding more detail to a signal. Additionally, by using layers, participants were able to separately design individual features of the tactile experience. For example, while looking at the lightsaber, P3 said “OK, I’m going to do two layers” and proceeded to create two layers with two different mappings, one for the background humming and another for the waving interactions (Figure 10b).

5.2.2 Spatio-Temporal Mappings Usage. From the created designs, we saw that both the *motion* and *positional* mappings were most commonly used, respectively 63 (50%) and 34 times (27%), while the *instantaneous* and *continuous* mapping were less common, respectively 17 (13%) and 13 times (10%). This underlines the relevance of mapping vibrotactile feedback to spatial and temporal interactions in the VR environment. Figure 10b shows an interesting example of a participant leveraging two layers to produce an experience combining *continuous* and *motion* mappings. Figure 10c depicts an example where a participant designed vibrotactile actuation for opening a metallic drawer. As noted by the participant, it was important to have a smooth experience that would provide an intense “boom” when the drawer reached the end. While the mappings available in the design tool enabled participants to create various experiences, participants were sometimes missing added temporal or spatial resolution. For instance, in our design tool vocalizations start as soon as the user starts interacting, however, this was not always matching the user’s intention (P1: “the sword, for instance, I had some vibrations when I was taking it, but not when I was putting [it] back on the table”). Also, some participants wanted to have more degrees of freedom on the action performed. For instance, one participant wanted to trigger different vocalizations based on the direction of the movement while stroking a texture (P4: “one [vocalization] I’d like to have when I move to the right, and one [vocalization] when I move to the left”) and another wanted to feel sand moving inside the box only when rotating it, but not when translating it (P8: “I would have found it cooler to use the rotation as the playback [...] now it does not really matter how you move it, since I am moving it will produce this [output]”).

5.2.3 Relevance of the Fine-tuning Modifiers. Generally, participants felt more inclined to alter the experience with modifiers than to re-record their vocalization, as explained by P2: “for me, it’s better to adjust the first recording than to record over and over again”. We observed participants manipulating modifiers extensively to refine a vocalization layer for better matching their intentions. Out of 117 effective vocalization layers, only 19 (16%) were left with modifiers untouched, which indicates the importance of the fine-tuning phase. *Expressiveness* was used 58 times (50%) and helped, for example, to make the rock illusion more convincing (P1: “[expressiveness] makes the rock feel bigger”). *Randomness* was used 28 times (24%)




Figure 10: Examples of vibrotactile experiences designed by the participants during the study: (a) the illusion of pebbles inside a box when shaking, (b) two layers with different mappings for isolating tactile experiences of a lightsaber, (c) emphasis on end of movement with positional mapping when interacting with a drawer. More examples can be found in Appendix A.

and helped, for example, to create more sand grains in the small box (P7: “[randomness] helps to make [the objects inside] feel less dense.”, or make a surface rougher (P6: “Randomness [...] makes [the slider] rougher.”). *Intensity* was used often (66 times, 56%) to amplify a vocalization, and sometimes to balance various layers. However, *modulation* was used less (25 times, 21%) as its effect only slightly impacted the experience using the commodity controller. When using the custom actuator, *modulation* was used more often to fine-tune experiences as here the effect was clearly noticeable. Overall, these results indicate fine-tuning modifiers are essentials to design vibrotactile feedback in-situ through vocalizations. However, several participant wished a more pronounced visual feedback on the effect of modifiers, as part of the waveform representing the vocalization.

While Weirding Haptics promotes direct manipulation and rapid prototyping through a quick vocalization process rather than through extensive post-processing functionalities, participants pointed out an important trade-off in the design process: they were missing audio processing functionalities. For example, they wanted to *trim* or *dampen* part of the vocalization to better control the timing of the vocalization or to fine-tune it (P2: “can I somehow delete one part of the audio, because [then the experience] would be perfect”, P5: “It would be very good if I can directly edit this [waveform]. [...] I would like to remove this part”). Several participants wanted to apply the *modifiers* to certain parts rather than the entire vocalization (P2: “I would like to [apply] intensity only on this part of the audio, but if I adjust the intensity, everything will change”). Some participants wondered whether they could speed up or down the vocalization they recorded to better match the movements performed in the playback phase (P8: “can I speed it up? That would be great”).

5.2.4 Output Modalities. To understand the effect of high-resolution frequencies on the vocalization design process, participants were asked to rate all experiences using both *native* and *custom* actuators. On average, the former was rated lower than the latter (native actuator: $\bar{x} = 5.25$; custom actuator: $\bar{x} = 5.75$; Mann-Whitney test:

$p < 0.01, r = 0.361$ - moderate effect). Participants found the custom actuator could “bring out the details” (P1) of the experience and made objects feel lighter or heavier (P6: “it feels heavier and it makes it more realistic kind of”; P8: “feels much more heavy, much more deep, and way too strong”). Outputting frequencies could also produce better precision (P7: “it feels like a smoother experience, like I am pulling the drawer more easily”; P8: “it’s more refined, there is more difference in the signal”), or increase the randomness or sharpness of the experience (P3: “feels more random”). Additionally, frequencies produced rougher (P7: “it feels better, it feels rougher”) and stickier sensations (P8: “it feels like it’s slightly sticky [...] that’s actually not what I wanted, but feels pretty cool”).

However, sometimes the *custom* actuator produced worse experiences. Some participants remarked it could introduce a sense of delay (P3: “weird, it feels delayed”; P6: “it feels more off, it does not really react to my movements”) or dampen vibrations to the extent of canceling out important bits of a vocalization (P3: “it feels like there is a dip [...] like you’re not doing anything”). While producing sharper experiences made some the experiences more convincing, it could also feel off in certain cases (P5: “I think you made it even more random and also sharper, but in this case it should not be sharper”; P7: “it does not feel as smooth”). Similarly, while frequencies induced weight, it sometimes did not align with the intention of the participant (P4: “it feels a bit softer, that’s not how I imagine it would feel”; P8: “too deep kind of”).

6 DISCUSSION AND LIMITATIONS

Here, we discuss lessons learned through the implementation of our design tool and observations from our studies. We reflect on Weirding Haptics’ efficiency to rapidly design vibrotactile feedback using the voice, as well as limitations highlighted during design. Furthermore, we formulate important design implications uncovered by the results of our initial validation study for in-situ fast-prototyping of vibrotactile feedback.

In-Situ Design of Vibrotactile Feedback: Our validation study results show that novice hapticians can use their voice to design a variety of tactile experiences that match their intention, after only a

short training period (~ 15 min). They quickly grasped how to control vibrotactile feedback through vocalizations, and how to interact with virtual objects to map these vocalizations to actions in the VR environment using spatio-temporal mappings. While users have created interesting tactile experiences only with their voice, they fine-tuned most (84%) of the vocalizations with the set of modifiers proposed by our design tool. This highlights the importance of the refinement phase in the design process. Furthermore, while we offered a multi-layered vocalization approach, in most cases (81.20%) a single vocalization layer was sufficient to reach the user's goals.

High-Resolution Spatio-Temporal Mappings: While the number and type of virtual objects used in our study might have constrained the mappings used, all available spatio-temporal mappings were used during the study. The limitation placed on the positional mapping to be constrained to a line segment was not mentioned as an issue by participants. As the visual mapping corresponded to the vibrational mapping, the interaction was similar to how real-world interactions are constrained, e.g., when opening a door using its handle the same path is traversed each time. Future versions of our framework could extend the positional mapping's functionality to consider any segment in 3D space.

Some participants were observed using multiple mappings for a single interaction by leveraging multiple vocalization layers. While participants created convincing tactile experiences with these mappings, they sometimes remarked missing some degree-of-freedom to map their vocalizations to visuo-spatial features or specific actions. For example, some participants wanted to map vibrotactile feedback to visuo-spatial features of a texture (e.g., holes between bricks in a wall), or trigger vibrotactile feedback based on the direction or type of their movements (rotational vs. translational). While such features can be supported by analyzing users' movements in more depth, high resolution spatial alignments remain a challenge. To tackle this, one could identify the correct part of the vocalization, and interpolate the feedback between landmarks of the same object. Our framework is compatible with more advanced mapping schemes as extensions can easily be incorporated.

Trade-off in the Fidelity Level while Fine-Tuning: Our studies highlighted the importance of the fine-tuning process while designing vibrotactile feedback, as also indicated in previous work [41]. With fine-tuning modifiers, our design tool supports rapid iterations for low- and medium-fidelity prototyping. Nevertheless, participants required a finer control of the vocalization audio processing. A recurring request was to provide tools to *trim* the audio signal to compensate for delays and support better timing of the vocalizations, as well as *edit* (with the modifiers or other functionalities) parts of the audio to preserve satisfactory ones. This underlines a limitation in the timeliness aspect of our design tool as well as a trade-off between the fast-prototyping approach supported by our tool and the relatively high-fidelity control participants aimed to have. While we expected users to re-record mistimed vocalizations, they seemed to be willing to spend more time fine-tuning features of their recordings. Future work will need to investigate the inclusion of higher-fidelity editing features with the ease and directness of in-situ design.

Generalizability of Weirding Haptics: With Weirding Haptics, a vocalization is strongly bound to the object interactions performed while recording. A drawback of this principle is that in order to experience exactly the same experience, one might need to execute

exactly the same action (e.g., waving an object using the same timing). Generalizing object interactions to match several interactive scenarios remains a major challenge. To achieve this, a designer's object interactions and vocalizations need to be considered in depth.

Furthermore, the Weirding Haptics design tool only considers single-object manipulation. Yet, many haptic experiences felt in the physical world involve several objects interacting together (e.g., sliding a cup of coffee on a table). Enabling free design of multi-object interactions is not obvious for design tools centered on non-expert users. Open challenges relate to how vocalization layers are assigned in this case, and how to manipulate them.

Actuator's Output Resolution: We compared two levels of resolution with two actuators. While we observed higher resolution output to produce more accurate experiences with respect to users' expectations, we constrained users to design with the low resolution actuator. Future work should investigate differences in design strategies depending on the degrees-of-freedom.

7 CONCLUSION AND FUTURE WORK

We presented Weirding Haptics, a novel concept for in-situ rapid prototyping of vibrotactile feedback in VR environments. Designing such feedback in-situ enables designers to synchronize vocalizations and object interactions using automatically inferred spatio-temporal mappings. Moreover, to provide a rapid design cycle inside VR, we identified fine-tuning modifiers to refine vibrotactile feedback and compensate for vocal limitations. Based on these insights, we presented a VR design tool implementing Weirding Haptics. Through a validation study, we tasked novice hapticians to freely design a set of haptic experiences for several virtual objects. The study results show that Weirding Haptics enables designing vibrotactile feedback that matches the designer's expectations, after only a short training time. Our observations uncovered important research challenges for the design of haptic experiences in-situ.

While Weirding Haptics supports novice hapticians, it is likely also an interesting concept for expert hapticians that need support for in-situ and fast prototyping for quicker iterations over a set of different vibrotactile feedback. More research is necessary, to fully understand how Weirding Haptics would benefit expert hapticians and what additional functionality is essential to them. Furthermore, we focused our research on designing haptic experiences inside a virtual environment, but our approach is not exclusive to this context; hapticians could also design experiences while manipulating physical objects, at the condition that the system could track and extract information from their actions. The vocalization layers displayed in-air in the virtual environment could then be displayed on a computer screen. Overall, we envision Weirding Haptics to be a broadly applicable concept for designing different types of haptic experiences in or outside VR.

ACKNOWLEDGMENTS

This project received funding from the European Research Council (ERC StG Interactive Skin 714797). We would like to thank the anonymous reviewers for their valuable feedback to improve the paper, and artists from the Noun Project⁷ for use of their artwork.

⁷thenounproject.com

REFERENCES

- [1] Kimi Akita and Mark Dingemans. 2019. Ideophones (Mimetics, Expressives). <https://doi.org/10.1093/acrefore/9780199384655.013.477>
- [2] Farah Arab, Sabrina Panéels, Margarita Anastassova, Stéphanie Coeugnet, Fanny Le Morellec, Aurélie Dommes, and Aline Chevalier. 2015. Haptic patterns and older adults: To repeat or not to repeat?. In *2015 IEEE World Haptics Conference (WHC)*. IEEE, Evanston, IL, USA, 248–253. <https://doi.org/10.1109/WHC.2015.7177721>
- [3] Sliman Bensmaïa and Mark Hollins. 2005. Pacinian representations of fine surface texture. *Perception & Psychophysics* 67, 5 (01 Jul 2005), 842–854. <https://doi.org/10.3758/BF03193537>
- [4] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. <https://doi.org/10.1191/1478088706qp0630a> arXiv:<https://www.tandfonline.com/doi/pdf/10.1191/1478088706qp0630a>
- [5] Lucie Brunet, Christine Megard, Sabrina Panéels, Gwenael Changeon, Jose Lozada, Marie P. Daniel, and Françoise Darses. 2013. “Invitation to the voyage”: The design of tactile metaphors to fulfill occasional travelers’ needs in transportation networks. In *2013 World Haptics Conference (WHC)*. 259–264. <https://doi.org/10.1109/WHC.2013.6548418>
- [6] Youngjun Cho, S Kim, M Joung, and J Lee. 2014. Haptic Cushion: Automatic Generation of Vibro-tactile Feedback Based on Audio Signal for Immersive Interaction with Multimedia. Messe Bremen.
- [7] Alain de Cheveigné and Hideki Kawahara. 2002. YIN, a fundamental frequency estimator for speech and music. *The Journal of the Acoustical Society of America* 111, 4 (April 2002), 1917–1930. <https://doi.org/10.1121/1.1458024>
- [8] Donald Degraen, Anna Reindl, Akhmadjon Maksadov, André Zenner, and Antonio Krüger. 2020. Envisioning Haptic Design for Immersive Virtual Environments. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference (Eindhoven, Netherlands) (DIS’ 20 Companion)*. Association for Computing Machinery, New York, NY, USA, 287–291. <https://doi.org/10.1145/3393914.3395870>
- [9] Donald Degraen, André Zenner, and Antonio Krüger. 2019. Enhancing Texture Perception in Virtual Reality Using 3D-Printed Hair Structures. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI ’19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300479>
- [10] Yves Guiard. 1987. Asymmetric Division of Labor in Human Skilled Bimanual Action. *Journal of Motor Behavior* 19, 4 (1987), 486–517. <https://doi.org/10.1080/00222895.1987.10735426>
- [11] Sean A Guynes. 2014. Four-Color Sound: A Peircean Semiotics of Comic Book Onomatopoeia. *Public Journal of Semiotics* 6, 1 (Dec. 2014), 58–72. <https://doi.org/10.37693/pjos.2014.6.11916>
- [12] Seongkook Heo, Jaeyeon Lee, and Daniel Wigdor. 2019. PseudoBend: Producing Haptic Illusions of Stretching, Bending, and Twisting Using Grain Vibrations. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (New Orleans, LA, USA) (UIST ’19)*. Association for Computing Machinery, New York, NY, USA, 803–813. <https://doi.org/10.1145/3332165.3347941>
- [13] Hiroki Hojo, Nozomi Nomachi, Yutaro Tomoto, Tsuyoshi Nakamura, Masayoshi Kanoh, and Koji Yamada. 2014. Fundamental Study for Verbalization of Embodied Expertise based on Pattern Recognition. In *2014 IIAI 3rd International Conference on Advanced Applied Informatics*. 925–929. <https://doi.org/10.1109/IIAI-AAI.2014.182>
- [14] Brent Edward Insko. 2001. *Passive Haptics Significantly Enhances Virtual Environments*. Ph.D. Dissertation. University of North Carolina at Chapel Hill, USA. Advisor(s) Frederick P. Brooks Jr. <http://www.cs.unc.edu/techreports/01-017.pdf>
- [15] Interhaptics. 2021. Haptic Composer. <https://www.interhaptics.com/haptic-composer>
- [16] Roland S Johansson and J Randall Flanagan. 2009. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience* 10, 5 (2009), 345.
- [17] Johan Kildal. 2010. 3D-Press: Haptic Illusion of Compliance When Pressing on a Rigid Surface. In *International Conference on Multimodal Interfaces and the Workshop on Machine Learning for Multimodal Interaction (Beijing, China) (ICMI-MLMI ’10)*. Association for Computing Machinery, New York, NY, USA, Article 21, 8 pages. <https://doi.org/10.1145/1891903.1891931>
- [18] Johan Kildal. 2012. Kooboh: Variable Tangible Properties in a Handheld Haptic-Illusion Box. In *Proceedings of the 2012 International Conference on Haptics: Perception, Devices, Mobility, and Communication - Volume Part II (Tampere, Finland) (EuroHaptics’12)*. Springer-Verlag, Berlin, Heidelberg, 191–194. https://doi.org/10.1007/978-3-642-31404-9_33
- [19] Erin Kim and Oliver Schneider. 2020. *Defining Haptic Experience: Foundations for Understanding, Communicating, and Evaluating HX*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376280>
- [20] Robert Kovacs, Eyal Ofek, Mar Gonzalez Franco, Alexa Fay Siu, Sebastian Marwecki, Christian Holz, and Mike Sinclair. 2020. Haptic PIVOT: On-Demand Handhelds in VR. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (Virtual Event, USA) (UIST ’20)*. Association for Computing Machinery, New York, NY, USA, 1046–1059. <https://doi.org/10.1145/3379337.3415854>
- [21] Susan J Lederman and Roberta L Klatzky. 1987. Hand movements: A window into haptic object recognition. *Cognitive Psychology* 19, 3 (1987), 342–368. [https://doi.org/10.1016/0010-0285\(87\)90008-9](https://doi.org/10.1016/0010-0285(87)90008-9)
- [22] Jaebong Lee and Seungmoon Choi. 2013. Real-time perception-level translation from audio signals to vibrotactile effects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Paris France, 2567–2576. <https://doi.org/10.1145/2470654.2481354>
- [23] Jaeyeon Lee, Mike Sinclair, Mar Gonzalez-Franco, Eyal Ofek, and Christian Holz. 2019. TORC: A Virtual Reality Controller for In-Hand High-Dexterity Finger Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI ’19)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3290605.3300301>
- [24] Gwilym Lockwood and Mark Dingemans. 2015. Iconicity in the lab: a review of behavioral, developmental, and neuroimaging research into sound-symbolism. *Frontiers in Psychology* 6 (Aug. 2015). <https://doi.org/10.3389/fpsyg.2015.01246>
- [25] Lofelt. 2021. Lofelt Studio. <https://lofelt.com/design>
- [26] Maud Marchal, Gabriel Cirio, Yon Visell, Federico Fontana, Stefania Serafin, Jeremy Cooperstock, and Anatole Lécuyer. 2013. *Multimodal Rendering of Walking Over Virtual Grounds*. Springer New York, New York, NY, 263–295. https://doi.org/10.1007/978-1-4419-8432-6_12
- [27] David Marino, Paul Bucci, Oliver S. Schneider, and Karon E. MacLean. 2017. Voodle: Vocal Doodling to Sketch Affective Robot Motion. In *Proceedings of the 2017 Conference on Designing Interactive Systems (Edinburgh, United Kingdom) (DIS ’17)*. Association for Computing Machinery, New York, NY, USA, 753–765. <https://doi.org/10.1145/3064663.3064668>
- [28] Michael Meehan, Brent Insko, Mary Whitton, and Frederick P. Brooks. 2002. Physiological Measures of Presence in Stressful Virtual Environments. *ACM Trans. Graph.* 21, 3 (July 2002), 645–652. <https://doi.org/10.1145/566654.566630>
- [29] Marianna Obrist, Sue Ann Seah, and Sriram Subramanian. 2013. *Talking about Tactile Experiences*. Association for Computing Machinery, New York, NY, USA, 1659–1668. <https://doi.org/10.1145/2470654.2466220>
- [30] Oculus Developers. [n.d.]. *Haptic Feedback | Oculus Developers*. <https://developer.oculus.com/documentation/native/pc/dg-input-touch-haptic/>
- [31] Breakfast Quay. 2021. *Rubber Band Library*. <http://breakfastquay.com/rubberband/>
- [32] Davide Rocchesso, Guillaume Lemaitre, Patrick Susini, Sten Ternström, and Patrick Boussard. 2015. Sketching Sound with Voice and Gesture. *Interactions* 22, 1 (Jan. 2015), 38–41. <https://doi.org/10.1145/2685501>
- [33] Joseph M. Romano and Katherine J. Kuchenbecker. 2012. Creating Realistic Virtual Textures from Contact Acceleration Data. *IEEE Transactions on Haptics* 5, 2 (April 2012), 109–119. <https://doi.org/10.1109/TOH.2011.38>
- [34] Eva-Lotta Sallnäs, Kirsten Rassmus-Gröhn, and Calle Sjöström. 2000. Supporting Presence in Collaborative Environments by Haptic Force Feedback. *ACM Trans. Comput.-Hum. Interact.* 7, 4 (Dec. 2000), 461–476. <https://doi.org/10.1145/365058.365086>
- [35] Ryoko Sasamoto. 2019. *Onomatopoeia and Relevance*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-26318-8>
- [36] Oliver Schneider, Karon MacLean, Colin Swindells, and Kellogg Booth. 2017. Haptic experience design: What hapticians do and where they need help. *International Journal of Human-Computer Studies* 107 (2017), 5–21. <https://doi.org/10.1016/j.ijhcs.2017.04.004> Multisensory Human-Computer Interaction.
- [37] Oliver S. Schneider, Ali Israr, and Karon E. MacLean. 2015. Tactile Animation by Direct Manipulation of Grid Displays. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology (Charlotte, NC, USA) (UIST ’15)*. Association for Computing Machinery, New York, NY, USA, 21–30. <https://doi.org/10.1145/2807442.2807470>
- [38] Oliver S. Schneider and Karon E. MacLean. 2014. Improvising design with a Haptic Instrument. In *2014 IEEE Haptics Symposium (HAPTICS)*. 327–332. <https://doi.org/10.1109/HAPTICS.2014.6775476>
- [39] Oliver S. Schneider and Karon E. MacLean. 2016. Studying design process and example use with Macaron, a web-based vibrotactile effect editor. In *2016 IEEE Haptics Symposium (HAPTICS)*. 52–58. <https://doi.org/10.1109/HAPTICS.2016.7463155>
- [40] Oliver S. Schneider, Hasti Seifi, Salma Kashani, Matthew Chun, and Karon E. MacLean. 2016. HapTurk: Crowdsourcing Affective Ratings of Vibrotactile Icons. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (San Jose, California, USA) (CHI ’16)*. Association for Computing Machinery, New York, NY, USA, 3248–3260. <https://doi.org/10.1145/2858036.2858279>
- [41] Hasti Seifi, Chamila Anthonypillai, and Karon E. MacLean. 2014. End-user customization of affective tactile messages: A qualitative examination of tool parameters. In *2014 IEEE Haptics Symposium (HAPTICS)*. 251–256. <https://doi.org/10.1109/HAPTICS.2014.6775476>


- //doi.org/10.1109/HAPTICS.2014.6775463
- [42] Hasti Seifi, Matthew Chun, Colin Gallacher, Oliver Schneider, and Karon E. MacLean. 2020. How Do Novice Hapticians Design? A Case Study in Creating Haptic Learning Environments. *IEEE Transactions on Haptics* 13, 4 (2020), 791–805. <https://doi.org/10.1109/TOH.2020.2968903>
 - [43] Hasti Seifi and Karon E. MacLean. 2017. Exploiting haptic facets: Users' sense-making schemas as a path to design and personalization of experience. *International Journal of Human-Computer Studies* 107 (2017), 38–61. <https://doi.org/10.1016/j.ijhcs.2017.04.003>
 - [44] David M. Sidhu and Penny M. Pexman. 2017. Five mechanisms of sound symbolic association. *Psychonomic Bulletin & Review* 25, 5 (Aug. 2017), 1619–1643. <https://doi.org/10.3758/s13423-017-1361-1>
 - [45] Mandayam A. Srinivasan and Cagatay Basdogan. 1997. Haptics in virtual environments: Taxonomy, research status, and challenges. *Computers & Graphics* 21, 4 (1997), 393–404. [https://doi.org/10.1016/S0097-8493\(97\)00030-7](https://doi.org/10.1016/S0097-8493(97)00030-7)
 - [46] Paul Strohmeier, Sebastian Boring, and Kasper Hornbæk. 2018. From Pulse Trains to "Coloring with Vibrations": Motion Mappings for Mid-Air Haptic Textures. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173639>
 - [47] Paul Strohmeier, Seref Güngör, Luis Herres, Dennis Gudea, Bruno Fruchard, and Jürgen Steimle. 2020. BARefoot: Generating Virtual Materials Using Motion Coupled Vibration in Shoes. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 579–593. <https://doi.org/10.1145/3379337.3415828>
 - [48] Paul Strohmeier and Kasper Hornbæk. 2017. *Generating Haptic Textures with a Vibrotactile Actuator*. Association for Computing Machinery, New York, NY, USA, 4994–5005. <https://doi.org/10.1145/3025453.3025812>
 - [49] C. Swindells, E. Maksakov, K.E. MacLean, and V. Chung. [n.d.]. The Role of Prototyping Tools for Haptic Behavior Design. In *2006 14th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE. <https://doi.org/10.1109/haptic.2006.1627084>
 - [50] Colin Swindells, Seppo Pietarinen, and Arto Viitanen. 2014. Medium fidelity rapid prototyping of vibrotactile haptic, audio and video effects. In *2014 IEEE Haptics Symposium (HAPTICS)*, IEEE, Houston, TX, USA, 515–521. <https://doi.org/10.1109/HAPTICS.2014.6775509>
 - [51] Valve Corporation. [n.d.]. *SteamVR Unity Plugin*. https://valvesoftware.github.io/steamvr_unity_plugin/index.html
 - [52] Ronald T Verrillo. 1966. Vibrotactile sensitivity and the frequency response of the Pacinian corpuscle. *Psychonomic Science* 4, 1 (1966), 135–136.
 - [53] Raquel Viciania-Abad, Arcadio Reyes Lecuona, and Matthieu Poyade. 2010. The Influence of Passive Haptic Feedback and Difference Interaction Metaphors on Presence and Task Performance. *Presence* 19, 3 (2010), 197–212. <https://doi.org/10.1162/pres.19.3.197>
 - [54] Yon Visell, Keerthi Adithya Duraikkannan, and Vincent Hayward. 2014. A Device and Method for Multimodal Haptic Rendering of Volumetric Stiffness. In *Haptics: Neuroscience, Devices, Modeling, and Applications*, Malika Auvray and Christian Duriez (Eds.). Vol. 8618. Springer Berlin Heidelberg, Berlin, Heidelberg, 478–486. https://doi.org/10.1007/978-3-662-44193-0_60
 - [55] Junji Watanabe and Maki Sakamoto. 2012. Comparison between onomatopoeias and adjectives for evaluating tactile sensations. In *Proceedings of the 6th International Conference of Soft Computing and Intelligent Systems and the 13th International Symposium on Advanced Intelligent Systems (SCIS-ISIS 2012)*. 2346–2348.
 - [56] Dennis Wittchen, Bruno Fruchard, Paul Strohmeier, and Georg Freitag. 2021. TactJam: A Collaborative Playground for Composing Spatial Tactons. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Salzburg, Austria) (TEI '21). Association for Computing Machinery, New York, NY, USA, Article 64, 4 pages. <https://doi.org/10.1145/3430524.3442699>
 - [57] Junya Yoshino, Akira Yakata, Yuichiro Shimizu, Masaharu Haginoya, and Maki Sakamoto. 2013. Method of evaluating metal textures by the sound symbolism of onomatopoeia. In *The Second Asian Conference on Information Systems*. 618–624.
 - [58] André Zenner and Antonio Krüger. 2019. Drag:On: A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300441>
 - [59] André Zenner and Antonio Krüger. 2017. Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 23, 4 (2017), 1285–1294. <https://doi.org/10.1109/TVCG.2017.2656978>

A PILOT STUDY

Category	Actions	Object	#Objects	Category	Actions	Object	#Objects
Push	flick	paper football	5	Tool Slide	slide	small coin	4
		marble				table	
	push	small box			slide	boxcutter knife	
		medium box				camera trolley	
		large					
Slide	sliding finger over	smooth	6	Tool Rotate	rotate	(smooth) potentiometer	5
		rough textured surface				servomotor	
		rougher				stepper motor	
	sliding hand over	smooth				dial (boxcutter)	
		rough textured surface			turn with both hands	steering wheel	
		rougher					
Rotate	mix	water	3	Tool Pull/Push	pull	bike pump	4
		sand			open	door	
		in cup with finger			open	drawer	
	waving	fan			open	scissors	
Pull	stretch with fingers	soft	6	Tool Press	push	small	4
		strong elastic band				medium	
		textured hairtie				large	
	stretch with hands	soft			press	pressure pen	
		strong exercise band					
	pull	feet of small tripod					
Passive Feel	touch	base	3	Press	press finger	soft	5
		tip toothbrush				stiff	
					squeeze	soft	
	pour water	bottle				stiff	
					pressing two hands	pregnancy ball	

Figure 11: List of actions and objects used in the pilot study.

B PARTICIPANT DESIGN EXAMPLES

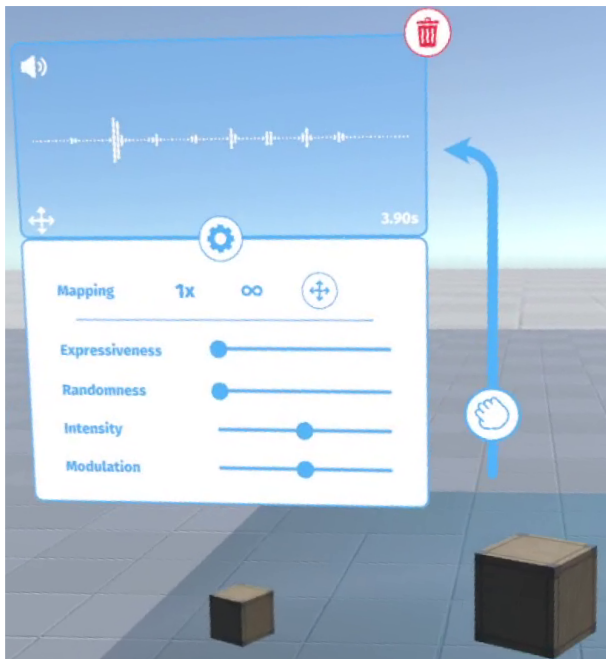


(a) Design of vibrotactile feedback for a slider with textures on top. Here, the difference in vocalization strength (see waveform) is balanced out using the expressiveness and intensity modifiers




(b) Design of vibrotactile feedback for a virtual lightsaber. Here, a layer with a continuous mapping provides active background vibrations, while a layer using a motion mapping emphasizes the swing of the lightsaber.

Figure 12: Example designs using multi-layered approaches.




(a) Illusion of pebbles inside a box. Here, the initial vocalization provided the intended result as no modifiers were changed.



(b) Illusion of a rock inside a box. Here, the modulation modifier was used to provide a sharper sensations.

Figure 13: Example designs for simulating illusions inside boxes.



(a) A smooth continuous sensation ends with a “boom” upon opening the drawer and reaching the end.



(b) Here, the participant aimed to create a sensation when dropping the sword, indicating a limitation of the current framework.

Figure 14: Example designs for a metallic drawer and a metal sword.




Figure 15: Design of vibrotactile feedback for touching a brick wall. Here, the modifiers were extensively changed to enhance the experience. Note the high amount of randomness aimed at providing the sensation of touching stone features.