

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Tactile Echoes: Multisensory Augmented Reality for the Hand

Permalink

<https://escholarship.org/uc/item/88x432gk>

Author

Kawazoe, Anzu

Publication Date

2021-04-02

Data Availability

The data associated with this publication are available at:

<https://datadryad.org/stash/dataset/doi:10.25349/D9BS5G>

Peer reviewed

Tactile Echoes: Multisensory Augmented Reality for the Hand

Anzu Kawazoe, Gregory Reardon, Erin Woo, Massimiliano Di Luca, *Member, IEEE*, and Yon Visell, *Member, IEEE*,

Abstract—Touch interactions are central to many human activities, but there are few technologies for computationally augmenting free-hand interactions with real environments. Here, we describe *Tactile Echoes*, a finger-wearable system for augmenting touch interactions with physical objects. This system captures and processes touch-elicited vibrations in real-time in order to enliven tactile experiences. We process these signals via a parametric signal processing network in order to generate responsive tactile and auditory feedback. Just as acoustic echoes are produced through the delayed replication and modification of sounds, so are Tactile Echoes produced through transformations of vibrotactile inputs in the skin. The echoes also reflect the contact interactions and touched objects involved. A transient tap produces discrete echoes, while a continuous slide yields sustained feedback. We also demonstrate computational and spatial tracking methods that allow these effects to be selectively assigned to different objects or actions. A large variety of distinct multisensory effects can be designed via ten processing parameters. We investigated how Tactile Echoes are perceived in several perceptual experiments using multidimensional scaling methods. This allowed us to deduce low-dimensional, semantically grounded perceptual descriptions. We present several virtual and augmented reality applications of Tactile Echoes. In a user study, we found that these effects made interactions more responsive and engaging. Our findings show how to endow a large variety of touch interactions with expressive multisensory effects.

Index Terms—Tactile augmented reality, wearable haptics, haptic rendering, multisensory feedback.

I. INTRODUCTION

INTERACTING with our environment frequently involves touching, exploring, or manipulating objects with the hands. Among the many haptic technologies that have been developed, few have been designed to augment naturally occurring touch interactions. Many existing haptic devices are based on controllers, instrumented surfaces, or hardware interfaces that must be operated by the hands. By occupying the hands, such interfaces often inhibit the great majority of manual interactions that support daily activities. We envisage new classes of electronic haptic interfaces that accommodate manual interactions involving direct skin contact with any

Anzu Kawazoe and Gregory Reardon are with the Media Arts and Technology Program, University of California Santa Barbara, CA 93106, USA (e-mail: anzu@ucsb.edu, reardon@ucsb.edu).

Erin Woo is with the Department of Electrical and Computer Engineering, University of California Santa Barbara, CA 93106, USA (e-mail: erinpwoo@ucsb.edu).

Massimiliano Di Luca is with the University of Birmingham, School of Psychology, Birmingham, UK (e-mail: m.diluca@bham.ac.uk).

Yon Visell is with the Department of Electrical and Computer Engineering and Media Arts and Technology Program, University of California Santa Barbara, CA 93106, USA (e-mail: yonvisell@ece.ucsb.edu)




Fig. 1: (A) Tactile Echoes system and concept. The wearable device captures vibrations in the fingertip that are produced during touch interactions, and processes them, and returns them to the finger as “echoes” of touch. (B) In this application, different Tactile Echoes are assigned to each region of the projected surface. The finger is tracked via camera, allowing different echoes to be assigned to different mapped regions on the surface. (C) In a musical controller application, a user controls a performance system by interacting with haptically augmented buttons, sliders, and knob interfaces rendered via a camera-projector system. (D) A VR experience involving a physical proxy object in which users are free to interact with miniature 3D printed objects upon which Tactile Echo feedback is superimposed.

object or surface in the surroundings. Only a few wearable devices for the hand have been designed to provide touch feedback without occluding skin-object contact [1], [2]. Addressing this gap could enable a wider range of human activities to be augmented with useful haptic information or evocative effects.

Here, we present a system for rendering effects that augment naturally occurring tactile sensations during manual interactions with objects and surfaces (Fig. 1). The system senses naturally occurring vibrations in the skin that are produced by contact with touched objects [3], [4] and transmitted throughout the skin [5], [6]. It processes the vibrations in real-time using a parametric signal network before returning them to the hand and ear as multisensory “echoes” of tactile interactions. Just as acoustic echoes are continuously produced in response to sound, Tactile Echoes can be continuously produced in

response to touch interactions. A hard tap produces a higher-amplitude response than a light tap, and a continuous slide produces feedback that is extended through time. The system can produce a wide array of responsive and evocative effects that can be parametrically designed using ten signal processing parameters.

Different Tactile Echoes can be assigned to different touched objects or interactions (Fig. 1B) by tracking movements of the hand in a mapped environment via optical, electromagnetic, ultrasound, or other technologies. This can enable a variety of applications in virtual, augmented, and mixed reality or human-computer interaction. Such applications can integrate informative surface-specific tactile feedback, introducing palpable digital information layers into physical environments, or can involve responsive, playful augmentations of ordinary touch interactions, among other possibilities.

In this paper, we first contextualize our work in the literature. We then describe the hardware and software, and the considerations that informed them. We next present experiments investigating how these unique haptic experiences are perceived. We describe three behavioral experiments and a multidimensional scaling (MDS) analysis based on user-provided descriptions and ratings. Analyses of the results shed considerable light on the perceptual dimensions underlying the unique experiences provided by our system. We next present several applications in which different objects, creative interfaces, or games are realized using Tactile Echoes. We then present a study evaluating how users appraised the form of tactile feedback provided by our system in one of these applications. We conclude with a discussion of these findings, opportunities for future work, and potential future implications for haptic engineering, augmented reality, and human-computer interaction design. This article is a revised and extended version of a paper we presented at the 2019 IEEE World Haptics Conference in Tokyo, Japan [7]. The present paper includes further content reviewing prior research projects and literature, additional detail about the signal processing network used to generate Tactile Echoes, additional vibrometry measurements, a more detailed analysis and discussion contrasting multisensory and haptic versions of the Tactile Echoes, and additional discussion about several facets of this wearable system. This paper also presents several new application demonstrations, with an accompanying user evaluation.

A. Background

The Tactile Echoes haptic feedback method shares similarities with other haptic feedback methods that are based on modulating the perceived properties of real objects by imposing forces felt via a haptic interface [8], [9] or with vibrations presented from a stylus [10], [11]. Such systems rely on generating signals to be reproduced via a device in response to performed motions or forces, but do not provide feedback during direct manual contact with touched objects. Closer to the approach taken in our work is the tactile magnification system of Yao and Hayward [12], which amplifies the sensations felt via a surgical tool.

Many other approaches to providing haptic feedback have been based on electronic gloves or exoskeletons [13]–[15], finger-mounted haptic devices [16]–[20], or grasped controls [21]. Few of these systems have integrated feedback from both real and virtual objects during free-hand interactions (in which the motion of the hand is essentially unrestricted). The great majority also introduce a surface or material between the hand and touched object, and thus restrict natural tactile sensation felt by the hand. Overcoming these limitations, as in our system, could pave the way for more effective and engaging haptic augmented and virtual reality systems.

In contrast, several methods have been proposed for superimposing touch-dependent haptic feedback on a tactile surface explored with the skin – typically a bare finger [22]–[27]. Similar to these methods, we compute tactile feedback via an algorithm that processes the sensed touch input. However, nearly all prior approaches of this kind provide feedback that is designed for a particular interaction type, such as textural sliding or tapping. The Tactile Echoes system generates feedback by processing the naturally occurring vibrations in the skin. The same algorithm can be applied to augment a wide variety of interactions – tapping, sliding, grasping, scratching with a finger, among other possibilities, all using the same system. One key difference between our work and the aforementioned examples is that our system augments real tactile interactions with unmistakably synthetic or “cartooned” haptic feedback that does not aim for realism, but rather at producing evocative effects. An analogy can be drawn to image distortion filters that are used for creative portraits, or to special effects in computer graphics, such as sparkles, glow effects, or explosions.

Recently, several researchers have described wearable electronic systems for capturing, amplifying, and reproducing natural tactile signals via skin-worn electronics. These include prior research in our lab [7], [28], and Makino et al. [29] as well as several collaborative works by Minamizawa, Maeda, Kakehi, Nakatani, Tsuchiya, Mihara, Peiris, and Tachi [30]–[32]. This research shows how it is possible to realize evocative experiences by concurrently sensing tactile signals elicited through skin-object contact and by amplifying the sensed signals to provide feedback on the same limb or another part of the body.

Such tactile amplification systems can yield interesting perceptual effects that are somewhat analogous to the auditory parchment skin or potato chip illusions [33], [34]. Several cross-modal effects of this type have also been uncovered. For example, in 1932, von Schiller reported tactile roughness perception to be influenced by the presence of concurrent auditory stimuli [35]. Other researchers have investigated the simultaneous use of vibrotactile and acoustic feedback associated with contact interactions. For example, Koehn and Kuchenbecker reported that users preferred haptic-auditory feedback from tool vibrations during robotic surgery [36]. Our system integrates haptic and auditory feedback in a way that is directed less at realism than at playfulness.

B. Summary of Contributions

Here, we show how both the sensing and feedback actuation

may be located on the same finger. Locating both sensing and actuation near the fingertip allows the physical and virtual sensations to better fuse into a single percept during touch interactions with physical objects. Crucial to our approach is our use of signal processing methods that minimize feedback instabilities, and that increase perceptual saliency by avoiding perceptual masking effects.

Prior examples of tactile amplification systems have provided for at least limited processing of the feedback that is supplied, including amplification. Maeda et al. went further by allowing for filtering, distortion, and other effects [31]. Here, we greatly expand on this approach by showing how a plurality of parametric processing stages can be used to yield a large continuum of haptic effects. We also use psychophysical methods to reveal several distinct underlying perceptual dimensions. The parameters in our system are addressable via UDP networked communication (as demonstrated in the applications presented in later sections of this paper).

Another key contribution of our work is that we show how to realize programmable tactile augmented reality with direct skin-object contact. We achieve this aim by combining wearable sensing, processing, and amplification with spatial position tracking. This system allows us to selectively assign distinct haptic effects to different surface regions or objects in a spatially mapped environment. In some configurations, our approach is analogous to visual augmented reality techniques that use projection mapping or head-mounted displays. Our research expands on previous approaches to haptic augmented reality that are based on users of electronic haptic feedback to supplement what is felt during interactions with real objects and environments [37], [38]. Our approach contrasts with these tool-based approaches, and instead augments interactions involving direct skin contact, similar to the projects discussed in the foregoing. Another distinctive aspect of our approach is that we supply responsive haptic feedback that, while derived from measured natural tactile signals, is unmistakably synthetic or “cartooned”. Similar approaches have been used in gaming or other applications [39].

Various methods have been used to investigate the perception of haptic feedback or effects superimposed on physical surfaces [40]–[42]. However, the Tactile Echoes system provides augmented tactile feedback that could be compared to synthetically rendered graphic effects (e.g., explosions) superimposed on real visual scenes. Such feedback need not resemble any natural touch experience, and indeed is not intended to reproduce natural touch experiences. Informed by these observations, we studied how Tactile Echoes are perceived via behavioral experiments, using a multidimensional scaling (MDS) paradigm. Since it was unclear, a priori, what factors or descriptors would best match Tactile Echoes, we based our approach on a methodology in which we systematically collected labels from users themselves, rather than from descriptors that we judged to be appropriate. Similar MDS methods have been previously used to assess the perception of natural haptic materials [43]–[45] and mechanisms [46] and have also been used to characterize the perception of synthetic haptic effects [47], [48]. In addition to identifying the perceptual space that characterizes Tactile Echoes, we demonstrate opportunities for

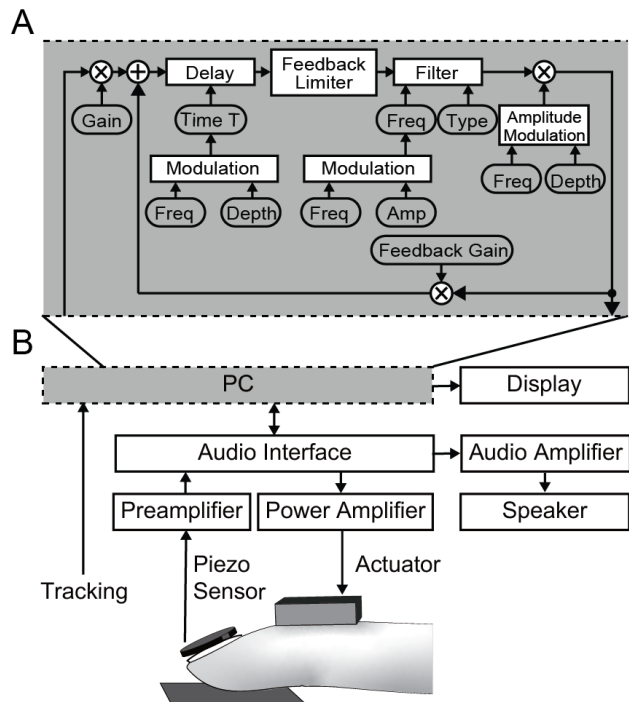


Fig. 2: (A) The Tactile Echoes are generated from the input via a parametric signal processing network. It includes modulated delay, nonlinear feedback limiting, amplitude modulation (tremolo), and modulated filtering. This architecture is sufficient to produce a wide variety of parametrized audio effects. (B) System Diagram: A piezoelectric sensor worn on the finger captures vibrations in the fingertip. The vibrations are amplified and concurrently processed by a computer. A signal processing network parametrically modifies the signals, which are amplified and returned to the finger via an inertial voice-coil actuator, and to the ear via a loudspeaker or headphone.

applying our system in several simple applications, including a VR application in which passive props [49], [50] acting as haptic proxies are augmented with dynamic, programmable tactile feedback.

II. SYSTEM DESIGN

The responsive and multisensory (haptic and auditory) feedback provided by Tactile Echoes is delivered by a system (Fig. 2) that captures and concurrently processes naturally occurring vibrotactile signals in the skin during manual interactions. The embodiment presented here senses vibrations in the finger as they are produced through touch interactions. It processes the sensed vibrations in real-time via a parametric signal network running on a computer, and continuously returns them to the finger and the ear (respectively using a vibrotactile and audio output device). The resulting tactile and auditory feedback augments what would normally be experienced during the touch interaction.

The wearable portion of the system consists of a fingernail-worn piezoelectric vibration sensor and a wide-bandwidth inertial voice coil actuator. The sensor, actuator, and cables are mounted in custom, ring-like brackets that were designed in CAD and fabricated in synthetic rubber via industrial

molding (Fig. 1A). The piezoelectric sensor signals are amplified (Puremini Amplifier, K&K Sound) and digitized in real-time using an audio analog-to-digital converter (Model 624, Mark of the Unicorn). The sampling rate is 44100 Hz. They are processed via a signal processing network running on a computer, and amplified (LP-2020A, Parts Express Inc.) after digital-to-analog conversion (Model 624, Mark of the Unicorn). The amplified signals drive the voice coil (Haptuator Mark II, Tactile Labs Inc.), returning the processed tactile signal to the finger with low latency (latency values are reported below). We use the same feedback signal in order to generate synchronous auditory feedback via a loudspeaker, headphone, or other device.

Tactile Echoes effects are produced by a parametric signal processing network (described below). The processing to be applied may be modified based on the proximity of the finger to different objects in the surroundings using standard tracking methods. For example, in Section 7 below, we demonstrate how to apply this technique when tracking the spatial position of the finger using the integrated camera of a smart projector system (Touch Xperia, Sony Inc.) for augmenting touch feedback on a projected touch surface, or via an optical hand tracking device (Leap Motion, Ultraleap Inc.) for augmenting tactile feedback on passive proxy objects in virtual reality. In such applications, the position tracking does not need to be precise enough to capture the contact event with high temporal accuracy. Instead, our system identifies the nearest mapped surface and selects the appropriate Tactile Echo before the surface is touched. Thus, many different motion tracking technologies could be used (for a recent review, see [51]). Our use of proximity to select the mode of feedback (i.e., the Tactile Echo settings) allows the tactile feedback to be responsively and automatically generated, synchronous with the touch event, because the Tactile Echo itself is driven by vibrations in the skin that are generated through finger-object contact.

While there are inevitable delays between the capture of input vibrations and the first feedback returned to the finger, our design leverages even longer delays (from 10 to 30 ms) than are imposed by system requirements, in order to enhance the effects themselves. During the course of designing our system, we observed that providing the aforementioned minimum delays greatly enhanced the perceptual saliency of the feedback. We hypothesize that this enhancement is due to a reduction in tactile forward masking effects. Kaaresoja et al. found that delayed tactile feedback increased the perceived mass of an electronic button [52]. The feedback delays in our system also reduce sensor-actuator feedback instabilities by allowing within-skin vibrations additional time to decay. Prior findings from our lab show that contact-like vibrations applied to the skin decay within a few tens of milliseconds [3].

A. Tactile Echoes – Signal Processing

Touch elicited vibrations in the finger are processed in real-time via software to yield a variety of parametrically-controlled effects. In our initial prototypes of this system, we used a guitar multi-effects box to explore the tactile feedback

generated by a set of 55 common audio effects during touch. These initial experiments revealed that some common audio effects, such as too long reverberation and distortion, seemed uninteresting, while others were highly evocative. Informed by this experience, we designed our system software to comprise a flexible, digital signal processing network of selected audio effects, with parametric controls over different processing stages (Fig. 2A). We use this signal processing network to generate a variety of Tactile Echoes by manipulating the values of the parameters. The network comprises a feedback delay structure with a variable gain, a resonant multimode filter, and nonlinear limiting integrated in the forward path of the delay structure. The limiting stage suppresses feedback loop instabilities and provides adaptive gain functions. Low frequency sinusoidal oscillators can optionally modulate each of the processing stages. In total, there are ten parameters that may be selected to specify the processing: output gain, feedback gain, delay time, filter corner frequency, filter type (highpass, lowpass, bandpass), filter resonance (Q factor), delay time modulation frequency and depth, and amplitude modulation (tremolo) frequency and depth. In other embodiments, a variety of other processing stages could also be used.

The amplitude modulation stage mitigates feedback instabilities that can arise due to the physical proximity of the sensor and actuator. We selected this feedback suppression method based on prior research in our lab, which evaluated several alternatives [53]. Feedback suppression is also aided by the imposed delays, as noted in the preceding section. While we have observed that such instabilities can occur for select settings within the large parameter space of our system, this only very occasionally arose during spontaneous use by hundreds of visitors to demonstrations we have given. For our experiments (described below), we selected the parameter settings of the stimuli to avoid feedback instabilities (and confirmed their absence through signal observation during the experiments).

Depending on the selected parameter values, the signal processing network can produce a large variety of effects. Some of these can resemble audio effects that are used in music production and performance, such as echo, slap-back, reverb, filtering, tremolo, filter delays, flange, or chorus effects, among others. Such effects have less often been used for the design of haptic or multisensory feedback. Through informal experimentation, we found delay time to produce the most appreciable qualitative changes. Delay times between 30 and 500 ms yielded especially interesting effects. The delay time also included a fixed feed-forward delay, due to input-output buffering in the digital audio hardware mentioned above. For our system settings, we measured this delay to be 20 ms. This value could be reduced significantly through software optimization, and could be reduced to below 1 ms using off the shelf hardware and software methods. However, we found that much stronger perceptual effects were produced if we ensured that a delay time of at least 30 ms elapsed between the sensor input produced by a touch interaction. We conjecture that this perceptual effect is due to tactile forward masking. We intend to explore this phenomenon in future work.




Fig. 3: (A) Waveforms produced by 35 designed effects in response to a single, pre-recorded finger tap captured by the piezo sensor (shown in red). The generated waveforms (in blue) are the “echo” signals furnished to the skin by the actuator. They were measured via Laser Doppler Vibrometer (LDV) along the axis of actuator vibration (top center). Each Echo is specified via values of ten processing parameters. In applications, different output waveforms are produced by each Echo in response to different touch inputs, much like different natural tactile signals are felt when touching an object differently. These Echoes each comprise one stimulus used in the perception experiments (Sec. III). (B) Computer voltage output (green trace) and actuator casing vibration (velocity, LDV measurements, blue trace) produced by a single echo, at four skin locations (i-iv, velocity, LDV measurements normal to skin surface, purple traces). Figure is adapted in part from the conference paper of which this article is a revised and extended version [7].

B. Tactile Echoes – Design and Mechanical Characterization

The large size of the 10-dimensional parameter space of our signal processing network precluded systematic evaluation of all parameter combinations. However, through manual search we identified regions of the parameter space that yielded palpable feedback and others that did not. Guided by these observations, we performed a heuristic search based on which we identified parameter settings for 88 varied Tactile Echoes that we judged to be interesting. We then selected 35 Echoes which we felt approximated the expressive range of effects that could be produced using our system and with our parametric signal processing network. We observed the differences between these 35 Echoes by measuring the vibrations produced by the actuator when attached to a participant’s finger (female, length of hand 16.5 cm, measured on right index finger). The hand from which measurements were captured arm was fixated to a vibration isolated table, with the measured finger left free. A non-contact Laser Doppler Vibrometer (Polytec PDV-100, Irvine, CA) measured the actuator velocity along the axis of vibration in response to the same, pre-recorded input from a discrete tap of the finger. The measured waveforms ranged in duration from 0.25 to 1 s, had varied densities of feedback, and different decay properties (Fig. 3A). While these sets of parameter values produced noticeably distinct waveforms, our heuristic selection process motivated the design of our experiments, which were based on user-supplied semantic

labels and ratings, and an MDS analysis.

When reproduced via the wearable hardware, Tactile Echoes yield mechanical vibrations of the skin that propagate as viscoelastic waves [3], [54]. From physics, such vibrations are expected to attenuate with distance d in a manner that depends on their frequency content [6]. For a vibration component of frequency f , a decrease in amplitude A with distance d is expected, with an approximately exponential relationship,

$$A(d) \sim \exp(-\alpha df), \quad (1)$$

where α is a damping coefficient. This damping contributes to the spatial localization of feedback in the finger, and reduces the influence of the actuator signal on the sensed signals. In our system, the combination of processing, feed-forward delay time, and damping in the skin reduce feedback instabilities, enabling larger gains to be used, and increasing dynamic range of the stimuli.

We empirically evaluated the vibrations imparted to the skin by the actuator when driven by Tactile Echoes waveforms using a non-contact scanning Laser Doppler Vibrometer (SLDV; Polytec PSV-500, Irvine, CA). The vibrometer measured the velocity of skin vibration in the direction normal to the volar skin surface at four locations (Fig. 3B). These measurements revealed that the Tactile Echoes system produced vibrations that were transmitted within the skin. The vibration waveforms at remote locations were similar to the those of the actuator signals. As expected from physics, the vibrations exhibit little

change in signal phase with distance (Fig. 3B), due to the relatively large (> 2 cm) wavelengths that occur at tactile frequencies ($f < 1000$ Hz). The vibrations attenuated with distance as expected from wave mechanics [6], [55].

III. PERCEPTION EXPERIMENTS

The goal of the experiments was to determine how touch interactions augmented by the Tactile Echoes were perceived and to identify a perceptual space that adequately described the perceptual similarity of different Tactile Echoes. The Tactile Echoes system can be applied to augment a wide variety of finger interactions such as sliding, grasping, tool-use, or scratching. We based our behavioral studies on a single gesture type, involving a discrete tap of the fingertip, which we judged to be an adequate proxy for transient contact events, such as initial skin-object contact, frequently occur during a large variety of manual interactions, such as pressing a switch, grasping an object, or touching a surface.

Our study design was informed by the fact that the Tactile Echoes stimuli are synthesized, and not intended to be realistic, and by our interest in avoiding biasing participant responses with our expectations about how the stimuli might be perceived. Our study is based on three perceptual experiments, a multidimensional scaling (MDS) procedure, and a regression analysis comparing the semantic ratings generated from the perceptual experiments with the MDS analysis. The three perceptual experiments consisted of a semantic labeling task, which employed a free verbalization method to elicit vocabulary which could be used to describe the sensations produced by the Tactile Echoes, a semantic sorting task, in which participants voted on the semantic labels to construct a unified set of 10 unipolar semantic labels to be used across participants, and a rating task, in which subjects rated the 35 Tactile Echoes based on semantic labels we determined via the preceding experiments. Our study was similar to those used in prior research [43]–[45]. This approach avoids difficulties that can arise if pre-determined adjective pairs are used [56]. Our system is also capable of producing multisensory feedback, by playing the Echoes as audio. To investigate the effect of this concurrent auditory feedback on how Tactile Echoes are perceived, we included both haptics-only and multisensory (audio-haptic) conditions.

A. Methods

1) *Participants*: In a first experiment, five native English speakers participated (ages 20 to 27, 3 male, 2 female). In a second experiment, a new set of seven native English speakers (ages 20 to 29, 4 male, 3 female) voted on the words that best described each stimulus. In a third experiment, fifteen new individuals (ages 20 to 50 years old, 10 male, 5 female) participated. All participants were right-hand dominant. Participants gave their written informed consent for the experiment, which was conducted according to the protocol approved by the UCSB institutional review board. Subjects were paid \$10 per hour for participating.

2) *Apparatus*: All experiments used the Tactile Echoes system. Participants were seated in a well-illuminated quiet room in front of a computer. Participants' hands were cleaned and sanitized in advance. The device was worn on the participant's dominant hand (right hand in all cases). In two conditions, haptic or multisensory, participants felt the Tactile Echoes with or without sound. In the multisensory condition, tactile and auditory feedback were produced concurrently via the same waveform used to drive the actuator. All experiments incorporated both conditions, haptic and multisensory. Every participant completed both conditions, one after the other, in random order per participant. Participants wore noise-cancelling headphones to prevent auditory cues, outside of those being presented in the multisensory condition. A curtain obstructed the view of the hand. We used a plastic-coated plywood sheet as the touch surface for all perceptual experiments. The surface was flat and uniform.

3) *Stimuli*: We used the set of 35 designed stimuli in all experiments (Fig. 3A). Each stimulus setting was presented once, individually, one per trial, in random order. During each trial, participants repeatedly tapped the surface at a rate of 0.67 Hz (guided by a visual metronome) while maintaining a tapping force between 1 and 1.5 N. We provided this guidance to ensure that participants experienced the stimuli in similar conditions. Software estimated the tapping force from the piezoelectric sensor signal, calibrated based on measurements from a laboratory force sensor. A visual indicator showed when participants tapped with appropriate or inappropriate force. Before the experiment, participants briefly practiced the procedure and practiced tapping with the requisite force levels.

B. Experiment 1: Descriptive Word Harvesting

In a first experiment, participants provided descriptive labels for the stimuli in each of the haptic and multisensory conditions. On each of the 35 trials, participants provided as many verbs and adjectives as they could to describe how the stimuli felt to them. Participants could experience each stimulus for as many times as they preferred and could enter responses as they proceeded. The duration of the first experiment was about 40 minutes.

C. Experiments 2: Word Voting

In a second experiment, a new set of participants voted on the words that best described each stimulus. We aggregated all of the words from the first experiment, after merging similar words using dictionary definitions and thesaurus associations. During each trial, participants were presented with one stimulus and a master list, in randomized order, of all words that had been collected for any stimulus via the first experiment. For each stimulus, participants selected any and as many words from the entire list that described what they felt. For each stimulus, participants could tap for as long as they preferred while they responded. The second experiment lasted about 30 minutes in total.

D. Experiment 3: Semantic Scaling

In a third experiment, a new set of participants rated each of the stimuli on a set of twelve semantic differential scales derived from the semantic labeling experiments. During each trial, participants rated one of the stimuli on 12 semantic differential continua. Responses were entered via computer. We used continuum scales rather than Likert scales to avoid introducing quantization (rounding) errors that could lose information. The semantic differential labels were chosen as the eleven most voted labels in Experiment 2. One further label “Real” was added by the experimenters, but yielded ambiguous results. Each of the 12 scales consisted of the label at the left extreme of the visual analog scale, and a second “not” label, indicating the literal converse, at the opposite side. Participants could experience each stimulus for as long as they preferred while they responded. We collected informal written comments and verbal reports from participants about their experience after the experiment. The duration of the third experiment was 1 hour, including a ten minute break.

E. Data Analysis

1) *Semantic Labeling*: The data from experiment 1 consisted of word sets that were aggregated to form the word list for voting in experiment 2. The word lists and votes were not further analyzed. The data from experiment 3 consisted of semantic differential scale ratings of each of the 35 stimuli in each condition (haptic, multisensory) by each participant. We analyzed the haptic and multisensory stimuli separately.

2) *MDS Analysis*: To assess the number of independent perceptual dimensions needed to describe the responses, and to derive a space that parametrized how the Tactile Echoes are perceived, we used the Classical Multidimensional Scaling (MDS) algorithm. It minimizes the mean residual error, called the strain between Euclidean distances (dissimilarities) among the original response vectors for each of the 35 Tactile Echoes gathered from the scaling experiment and the distances between their projection in a lower-dimensional embedding space. We computed MDS embeddings of dimensions 1 to 6, and computed the strain residuals for each. We selected embedding dimensionalities ($M = 2, 3$) based on the knee in the plot of strain residual vs. dimension (scree plot, Fig. 4), see discussion below. We computed the corresponding MDS embeddings for each value of the dimension, yielding four




Fig. 4: Scree plot showing the residual errors between the dissimilarity matrix and the MDS solutions as the number of dimensions increased.

spaces in total: two spaces for each condition and two spaces for each dimension. We computed mean response ratings for each stimulus and mapped each such value to one point in each MDS space.

3) *Regression between Scales and MDS Spaces*: We assessed the quality of the embeddings via Shephard diagrams – scatterplots of the dissimilarities vs. distances for each stimulus – and calculated R^2 values for each.

To further interpret the MDS mappings, and assess their quality, we used the entire dataset to fit the response data for each semantic differential scale as a function of the embedding coordinates. Regression of each scale yielded a line through the origin in each MDS space. We computed the R^2 values for each fit in order to assess the regression quality for each scale. This result allowed us to identify the semantic scales that were best predicted by the MDS coordinates, as those with the highest R^2 values. We identified orthogonal scales with high R^2 values (where $M = 2, 3$ is the embedding dimension) in order to interpret the MDS spaces in terms of participant-provided responses.

4) *Similarity of the Semantic Labels*: We used linear regression to map the perceptually-derived MDS spaces to each of the 12 semantic scales. Next, to investigate the perceptual dependence and independence of pairings of the response data for each semantic scale, we computed the relative angles between pairs of regression lines for each semantic scale in the MDS spaces (see Sec. 3.5.3) for both the haptic and multisensory conditions. For efficiency of presentation and to adhere to the length restrictions of this paper, we confined this further analysis to the 2D haptic and multisensory MDS spaces. We selected scales with $R^2 > 0.7$ for comparison. In brief, each regressed scale in each MDS space determined a vector with unit norm, \mathbf{u}_i , where $i = 1, 2, \dots, 12$. The geometric angle θ_{ij} between each pair of scales in each MDS space was computed as $\theta_{ij} = \arccos(\mathbf{u}_i \cdot \mathbf{u}_j)$. Angles close to 0 degrees are interpreted as the semantic scales describing identical sensations, while angles of 90 degrees are interpreted as the semantic scales describing independent perceptual dimensions. Angles of 180 degrees are interpreted as the semantic scales describing bipolar sensations.

5) *Comparison of Multisensory and Haptic Conditions*: To compare the perception of Tactile Echoes in the multisensory and haptic conditions, we computed the distributions of pairwise distances of the mean stimulus response values in each MDS space. These distributions describe the perceptual similarity between the stimuli in each condition. We compared the multisensory and haptic distributions for both the 2D and 3D MDS embeddings and used a Wilcoxon signed-rank test to ascertain whether the median perceptual distance between stimuli was different between conditions.

Next, we investigated differences in descriptor ratings between the haptic and multisensory conditions. In order to conduct this comparison between the multisensory and haptic conditions for each stimulus type, condition, and each semantic differential label, we computed a three-way ANOVA (conditions and stimuli, and descriptors as within-participant factors) and applied a Bonferroni multiple comparisons test. Before computing the ANOVA, we checked for normality

of residuals and homogeneity of variance. To check residual normality, the residuals from the model fit were analyzed graphically using Q-Q Plots; the residuals appeared normally distributed. Because there were 10,080 residuals, even small deviations from normality would be heavily penalized in conventional normality tests. To test for homogeneity of variance, we used multiple-sample tests for equal variances. After establishing significant main effects using an ANOVA, we used the Bonferroni multiple comparisons method to test for differences between the groupings of condition, descriptor and stimuli.

F. Results

1) *Semantic Scaling*: The results of Experiment 1 consisted of word sets that were aggregated to form a word list for voting in Experiment 2, which determined the semantic scales used in Experiment 3. We obtained 117 words in the haptic condition and 160 words in the multisensory condition. 46 words were common to both conditions.¹

2) *Perceptual Spaces*: Each of the four MDS analyses yielded a monotonically decreasing stress residual as dimensionality increased (Fig. 4), as expected. In both the multisensory and haptic conditions, the stress declined most as the dimension increased from 1 to 2 and from 2 to 3. The stress began to plateau as we increased the MDS solution space dimensions from 3 to 4. Thus, we focused our analysis on MDS spaces of dimension 2 and 3. Retaining both values of M for analysis allowed us to better understand how the MDS solution quality varied with dimensionality.

For each stimulus, we computed the mean value of all ratings across all presentations and mapped the resulting vector to the corresponding MDS space (Fig 5). The set of stimuli were widely distributed in all four spaces. The MDS optimization is invariant to orthogonal transformations – rotations and reflections of the data – so the orientation within these spaces is not informative.

Comparing the mean stimulus positions in the haptic and multisensory conditions, some Tactile Echoes that were near to one another in the haptic condition remained so when audio was added (examples in the 2D plot include 19 vs. 20, 29 vs. 9, 2 vs. 22, 29 vs. 9). Others that were near to one another in the haptic condition were farther apart in the multisensory condition (examples in the 2D plot include 10 vs. 34, 8 vs. 3, 2 vs. 25, 4 vs. 19). This is consistent with informal reports by participants that some Tactile Echoes features were perceived to be more prominent acoustically than haptically.

The linear regression analysis yielded a line representing each semantic scale in each MDS space (Fig 5). In the figure, line length is proportional to the R^2 value for the respective regression. The R^2 values ranged from 0.11 to 0.99. Several of the scales were nearly parallel, such as Wobbly and Echoing in the multisensory condition and Deep and Buzz in the haptic condition. These results suggest that

these scales were interpreted redundantly by participants in each condition. Others, including Hollow, remained nearly orthogonal to the other scales in all MDS cases, suggesting these ratings captured complementary perceptual ratings to the others. While there is no objective threshold for what constitutes a meaningful relationship, other researchers have relied on the judgement that scales with R^2 values greater than about 0.7 reflect substantial relationships [43], [44], [46]. In all four analyses, Deep, Rubbery, Rumble, and Wobbly yielded R^2 values greater than 0.7. It is often desired in such analyses to identify subsets of the scales of the same dimension as the space itself with high R^2 values. Such subsets can be used to interpret the MDS embedding coordinates of different stimuli. Suitable pairs in the 2D analyses include Deep-Wobbly in both the haptic and multisensory conditions, and Wobbly-Rumble or Wobbly-Deep (among other possibilities) in the haptic condition. In the 3D MDS analysis, one can point to triplets such as Wobbly-Rumble-Buzz in the haptic condition, or to Rubbery-Buzz-Wobbly in the multisensory condition.

3) *Perceptual Similarity Between Semantic Scales*: The relative angles between pairs of semantic scale regression lines (with $R^2 > 0.7$) in the MDS spaces reflected the perceptual similarity between the scales. Several pairs of scales yielded small, nearly parallel angles (angle magnitude < 15 degrees) reflecting high similarity, while several others were nearly orthogonal (90 ± 10 degrees) indicating high perceptual independence (Table I, shown in decreasing order of R^2 value).

4) *Comparison of Multisensory and Haptic Conditions*: The distribution of distances between mean stimuli ratings in the haptic and multisensory conditions was non-normal for both the 2D and 3D MDS embeddings. A Wilcoxon signed-rank test indicated a significant difference between the medians of the pairwise distances in the two conditions in 2D (median difference: 0.056, $Z = -3.7, p < 0.001$) and 3D (median difference: 0.057, $Z = -4.23, p < 0.0001$).

The three-way ANOVAs of the distributions of semantic ratings between conditions yielded residuals that were approximately normally distributed, with some light-tailed behavior, as determined graphically using Q-Q plots. Bartlett's multiple-sample tests showed that the variances in the semantic scale values across the stimuli, conditions, and descriptors were not significantly different ($p = 0.06$), supporting a constant variance analysis. The results of ANOVA test for all factors that were significant are shown in Table II. The Bonferroni-corrected comparisons of semantic scale descriptors revealed that several descriptors were significantly different between the haptic and multisensory conditions (Table III); ratings of Deep, Buzz, and Metallic were significantly higher, while ratings of Echoing, Bouncy, and Wobbly were significantly lower, in the haptic condition relative to the multisensory condition.

G. Perception Experiments: Discussion

1) *Perceptual Spaces for Tactile Echoes*: The descriptive word harvesting experiment revealed that participants employed a large variety of words to describe the effects.

¹The word lists are omitted for brevity. The lists, related results, and more details of the ten parameters in Tactile Echoes processing are summarized at this website: <http://spectrum2.mat.ucsb.edu/anzukawazoe/conf/TactileEchoes.html>.

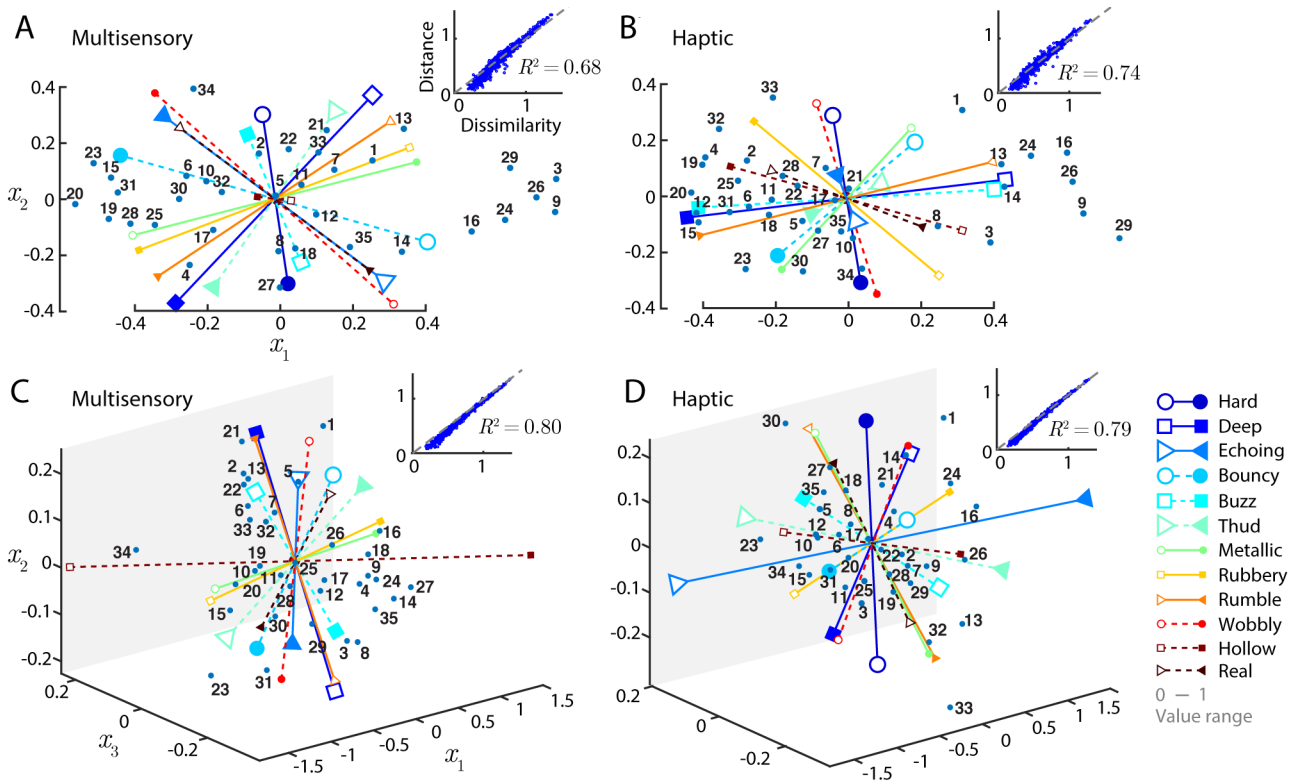


Fig. 5: The MDS analysis yielded embeddings of the Tactile Echoes stimuli in low dimensional spaces. The MDS embeddings were computed so that stimuli that are embedded near to each other received similar ratings in the semantic differential scaling experiments. For each condition, we computed MDS spaces in two dimensions (A: Multisensory, B: Haptic) and three dimensions (C: Multisensory, D: Haptic). The lines represent regression axes from MDS spaces to the semantic differential scale values; they ranged from 0 (hollow symbol) to 1 (filled symbol). The line length for each axis is proportional to the R^2 value of the regression, with longest lines denoting highest R^2 values. The Shepard plots (inset figures) show that the embedding quality increased for 3 vs 2 dimensions. Figure adapted from the conference paper of which this article is a revised and extended version [7].

TABLE I: Magnitudes of relative angles between semantic rating scale regression lines in the 2D MDS solution space (Left: Multisensory, Right: Haptic). Pairs with nearly orthogonal angles (90 ± 15 degrees) bold, in red cells. Pairs with small angles (0 ± 15 degrees) underlined in blue cells. Pairs with angles $180 - 15$ degrees are green.

	Wobbly	Deep	Bouncy	Echoing	Rumble		Deep	Buzz	Rumble	Rubbery	Wobbly	
Wobbly (0.944)	-						Deep (0.9568)	-				
Deep (0.8867)	101.6	-					Buzz (0.8851)	<u>4.1</u>	-			
Bouncy (0.879)	28.6	73.0	-				Rumble (0.8926)	<u>170.9</u>	<u>166.7</u>	-		
Echoing (0.8165)	<u>4.8</u>	96.7		-			Rubbery (0.7883)	55.1	51.0	115.7	-	
Rumble (0.7978)	88.5	<u>13.1</u>	59.8	83.6	-		Wobbly (0.7367)	84.2	80.1	93.4	29.1	
Rubbery (0.7901)	74.1	27.5	45.5	69.3	<u>14.4</u>		Hollow (0.72)	27.8	23.7	37.0	152.7	123.6

Examples included Wiggly, Thud, Twanging, Drop, Rattle, Thump, and Bouncy. In the haptic condition, words often evoked physical phenomena (Friction, Waves, Pulse, Thumping, Shock, and Reverberation). The multisensory experiment, which added auditory feedback, elicited a large number of descriptors that referred to material properties (Wood, Water, Marble, Glass, Liquid, Fluid, Woody, and Jelly) as well as words related to musical instruments (Drum, Banjo, and Guitar). The differences between the word lists in the haptic and multisensory conditions suggest that the presence of sound facilitated associations with material properties or objects, and that in the absence of sound, the effects evoked more abstract

phenomena.

In prior studies, it has been observed the perception of Roughness, Softness, and Temperature are involved in material recognition [57] and, in texture perception, that Roughness, Softness, and Sticky or Slippery are important perceptual dimensions [44]. In our experiments, participants provided words that are associated with roughness and softness (Rough, Gritty, Hollow, Soft, and Hard), and words that were found in haptic texture studies (Sticky, Smooth, and Slippery). We did not obtain words related to other dimensions, like temperature, which are not frequently associated with vibration signals. Thus, although some of the Tactile Echoes stimuli appeared

to evoke physical objects or processes, the association was limited in scope. In the voting experiment, the most commonly occurring word across both conditions was Bouncy. Others that were frequently selected included Echo, Short, Hard, Heavy, Rubbery, Rumble, and Light. Together, these comprised the most popular (top 15%) descriptors common to both conditions.

The MDS analysis revealed that despite the diversity of descriptors supplied by participants, and the ten different parameters used to design the stimuli, the perceptual similarity between the stimuli could be well-explained by just 2 or 3 dimensions. Several descriptors were highly correlated ($R^2 > 0.7$) with the MDS coordinates, including Deep, Buzz, Rumble, Rubbery, Wobbly, and Hollow in the haptic condition, and Wobbly, Deep, Bouncy, Echoing, Rumble, and Rubbery in the multisensory condition. In the further analysis of the 2D perceptual spaces, some pairs of the descriptors appeared to capture similar perceptual attributes, while others were complementary (Table I). In the multisensory condition, Echoing and Wobbly captured very similar perceptual information, as did Rumble and Deep, and Rumble and Rubbery. Thus the 2D perceptual space in the multisensory condition could best be parameterized via Deep-Wobbly dimensions, while the 2D perceptual space in the haptic condition could best be parameterized by the Wobbly-Rumble dimensions.

These results reflect differences between the Tactile Echoes stimuli in the conditions of the experiments, which involved tapping at approximately constant rates and forces on a relatively stiff surface. Further research is needed in order to clarify how these results might change if the tactile interactions were different. We hypothesize that a greater diversity of interaction types (e.g., continuous sliding on smooth or textured surfaces, tapping on soft surfaces) would increase the range of perceptual responses.

2) *Effects of the Sensory Conditions*: As indicated by the ANOVA, we found a significant interaction between “stimuli” and “conditions” (Table II). This significant interaction suggests that the presence of sound qualitatively altered how the stimuli were perceived. Further, the significant three-way interaction term between “conditions,” “descriptors,” and “stimuli” (Table II) implies that the qualitative change in how the stimuli were perceived in the presence of sound was dependent on the specific descriptor being rated. In another line of analysis, we found that the variation in responses, considered as the median pairwise MDS distances for the stimuli, was significantly smaller in the haptic than in the multisensory condition, indicating that the presence of sound increased the variation in responses. This is consistent with previous findings on multisensory perception [33], [58], [59].

The post-hoc Bonferroni multiple comparisons test indicated that there were significant differences in 6 of 12 descriptors ratings between the multisensory and haptic condition (Table III). Ratings of Deep, Buzz and Metallic were higher in the haptic condition, whereas ratings for Echoing, Bouncy, and Wobbly were higher in the multisensory condition. As for the reason for this difference in rating, it would appear that the auditory component made it possible to discriminate some stimuli that could not be distinguished from tactile information

TABLE II: Three-way ANOVA result in which conditions, and stimuli, and descriptors are within-participant factors. This table shows degrees of freedom (df), F-value, Significance (Sig.). Asterisks (*), (**), (****) indicate statistical significance at levels 0.05, 0.01, or 0.001.

Source	df	F-Value	Sig.
Conditions	1	9.364	**
Descriptors	11	28.826	****
Stimuli	34	16.029	****
Conditions*Descriptors	11	35.998	****
Conditions*Stimuli	34	1.642	*
Descriptors*Stimuli	374	2.639	****
Conditions*Descriptors*Stimuli	374	1.933	****

TABLE III: Multiple comparisons (Bonferroni’s test) of descriptors differing between conditions. The asterisks (****) indicate statistical significance at the 0.001 level.

Descriptor	Mean Rating Difference (Multisensory - Haptic)	p-Value
Deep	-0.124	$p < 0.0001$ ****
Metallic	-0.120	$p < 0.0001$ ****
Buzzing	-0.100	$p < 0.0001$ ****
Thud	-0.068	$p = 0.068$
Rumble	-0.017	$p = 1$
Bouncy	0.221	$p < 0.0001$ ****
Echoing	0.152	$p < 0.0001$ ****
Wobbly	0.112	$p < 0.0001$ ****
Rubbery	0.061	$p = 0.245$
Hollow	0.055	$p = 0.754$
Real	0.016	$p = 1$
Hard	0.007	$p = 1$

alone. Such difference could be due to the narrower tactile bandwidth limited below 100Hz by tactile actuator limitations above 700Hz by the rapidly decrease of tactile sensitivity. The results of the Bonferroni test, comparing the multisensory and haptic condition for each of the 35 stimuli, showed no significant difference between the mean semantic rating values for each stimulus. Thus, the presence of sound did not result in higher average ratings, although, as reported above, there was an effect when stimuli were grouped for each condition. The different results can be attributed to the conservative Bonferroni correction that is applied in the former case.

IV. DEMONSTRATING APPLICATIONS

We explored demonstrations of our system, informed in part by approaches adopted in previous research projects that have used wearable systems to haptically supplement naturally-occurring sensations felt during touch contact with real, physical objects associated with digital objects in virtual or augmented reality environments [60], [61].

We implemented three demonstration applications to illustrate how Tactile Echoes can be applied in virtual and augmented reality, human-computer interaction, and gaming. Our applications highlight the practical ways in which Tactile Echoes can be used to augment touch interactions with tactile feedback that is highly responsive, is parametrically and perceptually varied (as our experiments show), and can be assigned to different real or virtual objects, surfaces, or controls. The feedback is very responsive to the physics of the interaction because it is generated from vibrations

in the skin that are produced when touching real objects. The applications also illustrate how different low-complexity tracking methods are sufficient for enabling distinct Tactile Echoes to be selectively assigned to different objects, surface regions, or actions.

A. Multisensory Memory Game in VR with Augmented Passive Tangible Proxy Objects

In one application, we created a Virtual Reality memory game, modeled after the classic electronic game “Simon” (Fig. 6A). In it, users wear a head-mounted virtual reality headset. In a virtual game environment, they experience four, three-dimensional colored blocks. The blocks must be tapped in a specified sequence, matching a pattern that is first shown by the computer. After a user successfully reproduces a given sequence, the computer demonstrates a longer sequence. This proceeds until the user makes an error. The goal is to reproduce the longest sequence, yielding a high score. In the demonstration, the virtual blocks are co-located with physical blocks, which serve as passive haptic proxy objects [50], at corresponding registered positions in the physical environment. Each block is assigned a different multisensory Tactile Echo which is felt and heard by the user when activating one of the blocks. When the computer demonstrates a sequence, a representative, pre-recorded signal is used to generate the Tactile Echo associated with each block, which is heard, but not felt, by the user. Thus, the challenge can be regarded as that of memorizing the sequence as determined by the color, position, and Tactile Echoes assigned to the block.

In the application, Tactile Echoes are elicited when users touch the objects with a finger wearing our device. A Leap Motion hand-tracking camera provides relatively coarse information about the position of the finger relative to the block. This information allows the system to pre-activate the Tactile Echo corresponding to a block well before it is touched. This process is transparent to the user, since the Tactile Echoes feedback is driven responsively by the real contact between the finger and the object. Thus, a user experiences a seamless association of each block with the corresponding Tactile Echoes.

B. Augmenting a 2D Tactile Drawing Application

In a second demonstration, we created a 2D finger drawing application in which drawing actions are interactively augmented with different Tactile Echoes (Fig. 6B). The application is presented via an augmented reality surface generated by a smart projector system (Touch Xperia, Sony Inc.) running the Android operating system. An integrated camera in the projector tracks the user’s touch gestures. The user selects one of a large variety of sprite shapes and colors from a palette for fingerpainting on the projected display. Each color and shape is associated with a different Tactile Echo, which, when interacting in the specified drawing region on the interface, evokes an artificial, texture-like effect. The interface allows for drawing with continuous strokes or discrete taps, yielding discrete or sustained Tactile Echoes feedback. When we demonstrated this application in an exhibition at our university, we observed a

wide range of users, ranging from children to older adults, enjoy interacting with this multisensory creative experience. This application demonstrates how Tactile Echoes makes it possible to augment ordinary surfaces in the environment with continuously interactive projected interfaces that provide responsive tactile feedback.

C. Augmenting 2D Tactile Control Surfaces

In another demonstration, we used the same smart projector system to create a projected control surface. We mapped different tactile echoes to each of six colored regions (Fig. 1B). The dark areas of the interface, where no control button exists, are assigned to produce no Tactile Echoes feedback. In a separate application, we used the same approach and hardware to augment a projected touch screen based music controller (TouchOSC, Hexler, Ltd. [62]) with tactile feedback. The application provides a reconfigurable array of control surface elements, including sliders, dials, and buttons, for musical performance (see supplementary media and Fig. 1C). Tactile Echoes in this application are activated using control data transmitted via the Open Sound Control streaming network protocol [63]. Such augmented control surfaces can enable responsive, playful interfaces for creative applications. These applications demonstrate how it is possible to selectively assign tactile effects to different designer-specified control elements associated with a projected surface in a real environment.

D. Augmenting a 2D Video Game with Tactile Feedback

We created another simple demonstration in which we used Tactile Echoes to augment a controller for a side-scrolling video game based on a touch screen (Fig. 6C). The game runs on the smart projector system described above. In it, the player character, a rabbit, continuously travels to the right. The user is tasked with catching as many floating carrots as possible, which have been spawned at different heights, in the allotted time. In order to catch the carrots, the user taps on virtual buttons which make the rabbit jump. By tapping on the virtual buttons with more or less force, the user is able to control the height of the rabbit’s jump; the harder the user taps, the higher the rabbit jumps. In order to estimate tapping force, we used the piezoelectric sensor in the Tactile Echoes wearable device (Figs. 1A, 2B) in the manner described in the perception experiment. The tapping force was also automatically reflected in the Tactile Echoes feedback. This demonstration shows how Tactile Echoes make it possible to augment playful touch screen interactions with tactile feedback.

E. User Study

In order to evaluate whether users found Tactile Echoes to provide a more engaging and immersive experience in an application setting when compared to that of traditional vibrotactile feedback, we performed a user study based on the rabbit game demonstration. In the experiment, two virtual buttons were placed in the lower third of the screen (Fig. 6C). Each button (which we denoted “A” and “B”) was randomly




Fig. 6: Applications of Tactile Echoes with audio in virtual and augmented reality and human-computer interaction. A) A memory game in virtual reality using four passive haptic proxy objects augmented with different Tactile Echoes. B) A drawing application augmenting planar stroking, tapping, or scratching interactions with Tactile Echoes that depend on the selected color. C) A side-scrolling game in which a user controls a hopping rabbit (capturing carrots) via tapping with haptic feedback.

assigned to provide either Tactile Echoes feedback or a simple vibrotactile notification, consisting of a 200 Hz vibration cue with fixed amplitude lasting 250 ms. 12 participants volunteered for the experiment (7 male, 5 female). All subjects gave their written informed consent. Before the experiment began, each subject underwent a short, three-minute training phase in which they were free to press both buttons (i.e., simple haptic feedback and Tactile Echoes feedback) and learn the mechanics of the game. After training, subjects played the game twice. In the first trial, the Tactile Echoes feedback and the simple notification feedback was randomly assigned to either button “A” or “B” (e.g. “A” provided Tactile Echoes feedback and “B” provided a simple notification). In the second trial, the feedback assigned to each button was swapped (e.g. “A” was assigned to provide a simple notification, while “B” provided Tactile Echoes feedback). Each trial lasted three minutes. Participants were naive with respect to the purpose of the experiment, and were not informed about the different feedback modes. After each trial, participants answered three questions for each of the two buttons that were based on standard presence questionnaires:

- How responsive was button A/B to motion?
- How engaging was button A/B?
- How much agency or control do you feel when using the A/B?

Subjects answered using 7-point Likert scales (1 = Not at all; 4 = somewhat; 7 = completely). Subject responses were averaged across trials, resulting in 6 ratings per subject, 3 ratings describing how the Tactile Echoes feedback was perceived and 3 ratings describing how the simple notification feedback was perceived. We used a Wilcoxon signed-rank test to analyze the difference in median ratings between the two different types of feedback for question.

The median ratings for all questions were higher in the Tactile Echoes condition than in the control feedback condition. Participants judged the Tactile Echoes feedback to be more responsive to motion (median rating 6.5 vs. 4.0; $W=65$, $Z=2.9$, $p=0.002$, $r=0.59$), more engaging (6.25 vs. 3.75; $W=66$, $Z=3.04$, $p=0.001$, $r=0.62$), and more agency or control to

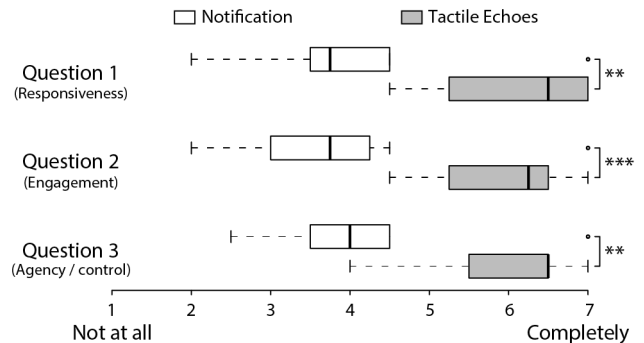


Fig. 7: Results of the user study of Tactile Echoes in video gaming. Boxes, whiskers, and points present the medians and IQRs, the $1.5 \times \text{IQR}$, and the outlier of evaluation value, respectively. The asterisks (**) and (***) indicate statistical significance at the 0.01 and 0.001 level, respectively.

facilitate (6.5 vs. 3.75; $W=55$, $Z=3.0$, $p=0.002$, $r=0.605$)². This result suggests that Tactile Echoes could enhance user experiences in many applications that currently rely on simpler haptic notifications.

V. CONCLUSION

This paper presents a wearable method and system for multisensory augmentation of manual touch interactions with objects and surfaces. This enables responsive haptic effects to be rendered during manual interactions involving direct contact with the skin. Our method allows tactile feedback to be introduced into naturally occurring interactions without requiring the touched object to be engineered and without imposing any device, such as a handheld controller or instrumented surface, between the skin and touched object. Thus, it can be used in a great variety of environments and interactions. This system represents a promising design approach for tactile augmented or mixed reality. It could be compared emerging visual augmented reality methods like those based

² W , Z , p , r are test statistics, critical z-value, p-value, and effect size, respectively.

on head mounted displays or projection systems. Our work also demonstrates how tactile feedback can be programmably assigned to objects or surface regions (Figure 6).

The Tactile Echoes system captures naturally occurring vibrations in the skin that are elicited via touch contact during manual interactions. It processes the vibrations and returns them to the hand as echoes of touch and to the ear as sound. The feedback automatically reflects the attributes of the contact event or touched object. Our system provides ten parameters to design these effects via a signal processing network. The same processing can be used to generate either only tactile feedback or concurrent tactile and auditory feedback, yielding multisensory experiences. Many other signal processing architectures and parameters can be used to realize such effects.

In perceptual experiments, we characterized how Tactile Echoes are perceived using semantic labels that were provided by participants. MDS analyses yielded low-dimensional, semantically grounded descriptions of the underlying perceptual spaces. While these results reflect design choices we adopted, and many other such choices are possible. The labels were often related to familiar physical processes or objects. We hypothesize that aspects of the perceptual mapping revealed here would be preserved in other embodiments of our approach, but further research is needed to clarify this hypothesis.

The promising nature of these results suggests several avenues for further investigation. First, the effects that we designed proved to be evocative and diverse, but not necessarily natural. Nonetheless, participants frequently described them using terms that referred to physical processes. Further research on how these effects might be designed to match natural touch sensations, or to modify the perceived properties of surfaces, is warranted. Second, individual differences in perception could arise from variations in the size, stiffness, and shape of the finger, as would be appropriate for further study. Third, as our applications demonstrate, this design can be used to generate responsive haptic effects in response to a variety of touch interactions, including tapping, textural sliding, and scratching, among others. Our perception experiments focused mainly on touch contact via tapping, while the applications also demonstrate sliding contact. Further research is warranted to investigate the perception of Tactile Echoes accompanying more general interactions. This research deduced perceptual spaces grounded in user-supplied semantic descriptors. It would be interesting to leverage these low-dimensional representations to simplify the design of Tactile Echoes effects. We plan to explore this design simplification in future work. Fourth, while we have presented several different demonstrations, the majority involve the haptic augmentation of nearly flat extended surfaces of objects. We have explored an array of potential interactive scenarios (including several not described here), and have informally found scenarios involving the augmentation of low-curvature surfaces to produce more interesting results than are typically obtained using three-dimensional objects. This could be due, in part, to the single-finger nature of the interactions involved, but other considerations may also be at play. We plan to investigate these issues further in future work.

We designed the physical implementations presented here based upon a piezoelectric vibration sensor, inertial voice coil actuator, motion sensing and display systems that were efficient to implement and appropriate for the experiments and demonstrations. However, many other variations on this system and these components are also possible. The implementations in our system are all tethered through physical wires, but this system can be made wireless and battery-powered, with wireless data transmission link to a remote desktop computer. We prototyped such a configuration in an earlier project in our lab [53]. The computing and motion sensing portions of the system could also be made wearable, leveraging contemporary head-mounted augmented reality glasses, goggles, and computer vision sensing, as we plan to explore in future work.

APPENDIX

The supplemental media of applications and experiments can be retrieved here: <https://youtu.be/HrR5WuPiMmU>

ACKNOWLEDGMENT

This work was supported by the US National Science Foundation (NSF-1628831, NSF-1623459, and NSF-1751348). We thank S. Fyke, B. Cunningham, and K. Sparackin of Snap Pea Design for their contributions to an earlier version of this system, and thank Y. Shao and T. Hachisu for comments.

REFERENCES

- [1] C. Bermejo and P. Hui, "A survey on haptic technologies for mobile augmented reality," *CoRR*, vol. abs/1709.00698, 2017.
- [2] A. Matsubayashi, Y. Makino, and H. Shinoda, "Direct finger manipulation of 3d object image with ultrasound haptic feedback," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. New York, NY, USA: ACM, 2019, pp. 87:1–87:11.
- [3] Y. Shao, V. Hayward, and Y. Visell, "Spatial patterns of cutaneous vibration during whole-hand haptic interactions," *Proceedings of the National Academy of Sciences*, 2016.
- [4] —, "Compression of dynamic tactile information in the human hand," *Science Advances*, (In Press) 2020.
- [5] H. Shaefer, Z. Wells, Y. Shao, and Y. Visell, "Transfer properties of touch elicited waves vary with hand posture and contact conditions," in *Proc. IEEE World Haptics Conference*, 2017.
- [6] B. Dandu, Y. Shao, A. Stanley, and Y. Visell, "Spatiotemporal haptic effects from a single actuator via spectral control of cutaneous wave propagation," in *Proc. IEEE World Haptics Conference*, 2019.
- [7] A. Kawazoe, M. Di Luca, and Y. Visell, "Tactile echoes: A wearable system for tactile augmentation of objects," in *Proceedings of the 2019 IEEE World Haptics Conference*.
- [8] K. J. Kuchenbecker, J. Fiene, and G. Niemeyer, "Improving contact realism through event-based haptic feedback," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 2, pp. 219–230, March 2006.
- [9] A. M. Okamura, M. R. Cutkosky, and J. T. Dennerlein, "Reality-based models for vibration feedback in virtual environments," *IEEE/ASME Transactions on Mechatronics*, vol. 6, no. 3, pp. 245–252, Sep 2001.
- [10] T. Hachisu and H. Kajimoto, "Vibration feedback latency affects material perception during rod tapping interactions," *IEEE Transactions on Haptics*, vol. 10, no. 2, pp. 288–295, April-June 2017.
- [11] —, "Hachistack: dual-layer photo touch sensing for haptic and auditory tapping interaction," in *2013 ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '13, Paris, France, April 27 - May 2, 2013*, 2013, pp. 1411–1420.
- [12] H.-Y. Yao, V. Hayward, and R. E. Ellis, "A tactile magnification instrument for minimally invasive surgery," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004*, C. Barillot, D. R. Haynor, and P. Hellier, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 89–96.

- [13] M. Bouzit, G. Burdea, G. Popescu, and R. Boian, "The Rutgers master ii-new design force-feedback glove," *IEEE/ASME Transactions on Mechatronics*, vol. 7, no. 2, pp. 256–263, Jun 2002.
- [14] "CyberGrasp force feedback system and CyberTouch tactile feedback system. Online , Last accessed on July, 2018."
- [15] J. Hummel, J. Dodiya, G. A. Center, L. Eckardt, R. Wolff, A. Gerndt, and T. W. Kuhlén, "A lightweight electro-tactile feedback device for grasp improvement in immersive virtual environments," in *2016 IEEE Virtual Reality (VR)*, March 2016, pp. 39–48.
- [16] F. Chinello, M. Malvezzi, C. Pacchierotti, and D. Prattichizzo, "Design and development of a 3rrs wearable fingertip cutaneous device," in *2015 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, July 2015, pp. 293–298.
- [17] M. Solazzi, A. Frisoli, and M. Bergamasco, "Design of a novel finger haptic interface for contact and orientation display," in *2010 IEEE Haptics Symposium*, March 2010, pp. 129–132.
- [18] D. Prattichizzo, F. Chinello, C. Pacchierotti, and M. Malvezzi, "Towards wearability in fingertip haptics: A 3-dof wearable device for cutaneous force feedback," *IEEE Trans. Haptics*, vol. 6, no. 4, pp. 506–516, oct 2013.
- [19] C. Pacchierotti, G. Salvietti, I. Hussain, L. Meli, and D. Prattichizzo, "The hring: A wearable haptic device to avoid occlusions in hand tracking," in *2016 IEEE Haptics Symposium (HAPTICS)*, April 2016, pp. 134–139.
- [20] S. B. Schorr and A. M. Okamura, "Fingertip tactile devices for virtual object manipulation and exploration," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, May 2018, pp. 3115–3119.
- [21] E. Whitmire, H. Benko, C. Holz, E. Ofek, and M. Sinclair, "Haptic revolver: Touch, shear, texture, and shape rendering on a reconfigurable virtual reality controller," in *CHI 2018*. ACM, April 2018.
- [22] J. Fleureau, Y. Lefevre, F. Danieau, P. Guillotel, and A. Costes, "Texture rendering on a tactile surface using extended elastic images and example-based audio cues," vol. 9774, 07 2016, pp. 350–359.
- [23] M. Wiertelowski, D. Leonardis, D. J. Meyer, M. A. Peshkin, and J. E. Colgate, "A high-fidelity surface-haptic device for texture rendering on bare finger," in *Haptics: Neuroscience, Devices, Modeling, and Applications*, M. Auvray and C. Duriez, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 241–248.
- [24] M. Niwa, T. Nozaki, T. Maeda, and H. Ando, "Fingernail-mounted display of attraction force and texture," in *Haptics: Generating and Perceiving Tangible Sensations*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 3–8.
- [25] K. Sakai, T. Hachisu, and Y. Hashimoto, "Characteristics of illusory vibration in the toe pad induced by vibration applied to toenail and toe force to surface," 06 2017, pp. 281–286.
- [26] S. Asano, S. Okamoto, and Y. Yamada, "Vibrotactile stimulation to increase and decrease texture roughness," *IEEE Transactions on Human-Machine Systems*, vol. 45, pp. 393–398, 2015.
- [27] —, "Toward augmented reality of textures: Vibrotactile high-frequency stimuli mask texture perception to be rougher or smoother?" *2013 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 510–515, 2013.
- [28] M. Chartier, N. Thomas, Y. Shao, and Y. Visell, "Toward a wearable tactile sensory amplification device: Transfer characteristics and optimization," in *Work in Progress of the 2016 Haptic Symposium Conference*. IEEE., 2016, pp. 546–551.
- [29] Y. Makino, M. Murao, and T. Maeno, "Life log system based on tactile sound," in *Haptics: Generating and Perceiving Tangible Sensations*, A. M. L. Kappers, J. B. F. van Erp, W. M. Bergmann Tiest, and F. C. T. van der Helm, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 292–297.
- [30] K. Minamizawa, Y. Kakehi, M. Nakatani, S. Mihara, and S. Tachi, "Tactile toolkit: a prototyping tool for design and education of haptic media," in *Proceedings of the 2012 Virtual Reality International Conference*. ACM, 2012, p. 26.
- [31] T. Maeda, R. Peiris, M. Nakatani, Y. Tanaka, and K. Minamizawa, "Wearable haptic augmentation system using skin vibration sensor," in *Proceedings of the 2016 Virtual Reality International Conference*. ACM, 2016, pp. 25:1–25:4.
- [32] T. Maeda, K. Tsuchiya, R. Peiris, Y. Tanaka, and K. Minamizawa, "Haptic-aid: Haptic experiences system using mobile platform," in *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 2017, pp. 397–402.
- [33] V. Jousmäki and R. Hari, "Parchment-skin illusion: sound-biased touch," *Current Biology*, vol. 8, no. 6, pp. R190 – R191, 1998.
- [34] M. Zampini and C. Spence, "The role of auditory cues in modulating the perceived crispness and staleness of potato chips," *Journal of Sensory Studies*, vol. 19, no. 5, pp. 347–363, 2004.
- [35] P. von Schiller, "Die rauhhigkeit als intermodale erscheinung (roughness as an intermodal phenomenon)," *Zeitschrift für Psychologie Bildung*, vol. 127, pp. 265–289, 1932.
- [36] J. K. Koehn and K. J. Kuchenbecker, "Surgeons and non-surgeons prefer haptic feedback of instrument vibrations during robotic surgery," *Surgical Endoscopy*, vol. 29, pp. 2970–2983, 2014.
- [37] S. Jeon and S. Choi, "Haptic augmented reality: Taxonomy and an example of stiffness modulation," *PRESENCE: Teleoperators and Virtual Environments*, vol. 18, pp. 387–408, 2009.
- [38] H. Culbertson and K. J. Kuchenbecker, "Ungrounded haptic augmented reality system for displaying roughness and friction," *IEEE/ASME Transactions on Mechatronics*, vol. 22, pp. 1839–1849, 2017.
- [39] B. T. Gleeson and D. E. Johnson, "Expressive haptic rendering with cartoon-inspired effects," *2010 IEEE Haptics Symposium*, pp. 191–194, 2010.
- [40] S. Okamoto, M. Konyo, and S. Tadokoro, "Vibrotactile stimuli applied to finger pads as biases for perceived inertial and viscous loads," *IEEE Transactions on Haptics*, vol. 4, no. 4, pp. 307–315, October 2011.
- [41] S. Yoshimoto, Y. Kuroda, M. Imura, and O. Oshiro, "Material roughness modulation via electro-tactile augmentation," *IEEE Transactions on Haptics*, vol. 8, no. 2, pp. 199–208, April 2015.
- [42] S. Jeon and S. Choi, "Haptic augmented reality: Taxonomy and an example of stiffness modulation," *Presence*, vol. 18, no. 5, pp. 387–408, Oct 2009.
- [43] M. Hollins, R. Faldowski, S. Rao, and F. Young, "Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis," *Perception & Psychophysics*, vol. 54, no. 6, pp. 697–705, Nov 1993.
- [44] M. Hollins, S. Bensmaïa, K. Karlof, and F. Young, "Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling," *Perception & Psychophysics*, vol. 62, no. 8, pp. 1534–1544, Dec 2000.
- [45] W. M. B. Tiest and A. M. Kappers, "Analysis of haptic perception of materials by multidimensional scaling and physical measurements of roughness and compressibility," *Acta Psychologica*, vol. 121, no. 1, pp. 1 – 20, 2006.
- [46] Q. Liu, H. Z. Tan, L. Jiang, and Y. Zhang, "Perceptual dimensionality of manual key clicks," in *2018 IEEE Haptics Symposium (HAPTICS)*, March 2018, pp. 112–118.
- [47] D. Ternes and K. E. Maclean, "Designing large sets of haptic icons with rhythm," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 2008, pp. 199–208.
- [48] J. Luk, J. Pasquero, S. Little, K. MacLean, V. Levesque, and V. Hayward, "A role for haptics in mobile interaction: initial design using a handheld tactile display prototype," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 2006, pp. 171–180.
- [49] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell, "Passive real-world interface props for neurosurgical visualization," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '94. New York, NY, USA: ACM, 1994, pp. 452–458.
- [50] L.-P. Cheng, E. Ofek, C. Holz, H. Benko, and A. D. Wilson, "Sparse haptic proxy: Touch feedback in virtual environments using a general passive prop," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ser. CHI '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 3718?3728.
- [51] M. Yahya, J. A. Shah, K. Kadir, Z. M. Yusof, S. Khan, and A. Warsi, "Motion capture sensing techniques used in human upper limb motion: a review," *Sensor Review*, vol. 39, pp. 504–511, 2019.
- [52] T. Kaaresoja, E. E. Hoggan, and E. Anttila, "Playing with tactile feedback latency in touchscreen interaction: Two approaches," in *INTERACT*, 2011.
- [53] H. Schafer, "Touch amplification for human computer interaction," Master's thesis, UC Santa Barbara, 2017.
- [54] R. R. Z. H. Chunhui Li, Guangying Guan and R. K. Wang, "Determining elastic properties of skin by measuring surface waves from an impulse mechanical stimulus using phase-sensitive optical coherence tomography," *J R Soc Interface*, vol. 9, no. 70, pp. 831–41., 2011.
- [55] L. R. Manfredi, A. T. Baker, D. O. Elias, J. F. Dammann III, M. C. Zielinski, V. S. Polashock, and S. J. Bensmaïa, "The effect of surface wave propagation on neural responses to vibration in primate glabrous skin," *PLoS one*, vol. 7, no. 2, 2012.
- [56] S. Okamoto, H. Nagano, and Y. Yamada, "Psychophysical dimensions of tactile perception of textures," *IEEE Transactions on Haptics*, vol. 6, no. 1, pp. 81–93, First 2013.

- [57] L. Jones and S. Lederman, *Human Hand Function*. Oxford University Press, 2006.
- [58] S. Guest, C. Catmur, D. Lloyd, and C. Spence, "Audiotactile interactions in roughness perception," *Experimental Brain Research*, vol. 146, no. 2, pp. 161–171, Sep 2002.
- [59] M. Zampini and C. Spence, "The role of auditory cues in modulating the perceived crispness and staleness of potato chips," *Journal of Sensory Studies*, vol. 19, no. 5, pp. 347–363, 2004.
- [60] O. Bau and I. Poupyrev, "Revel: tactile feedback technology for augmented reality," *ACM Trans. Graph.*, vol. 31, pp. 89:1–89:11, 2012.
- [61] S. V. Salazar, C. Pacchierotti, X. de Tinguy, A. Maciel, and M. Marchal, "Altering the stiffness, friction, and shape perception of tangible objects in virtual reality using wearable haptics," *IEEE Transactions on Haptics*, vol. 13, pp. 167–174, 2020.
- [62] A. Gräf, "A touchosc midi bridge for linux," 2014.
- [63] A. Freed and A. Schmeder, "Features and future of open sound control version 1.1 for nime," in *In Proceedings of the 9th Conference on New Interfaces for Musical Expression*, 2009, pp. 116–120.



Yon Visell received the PhD Degree in Electrical and Computer Engineering from McGill University (2011), and MA degree in Physics from Univ. Texas-Austin and the BA degree in Physics from Wesleyan Univ. He is Associate Professor at the University of California, Santa Barbara – in the Media Arts and Technology Program, Department of Electrical and Computer Engineering, and Department of Mechanical Engineering (by courtesy) – where he directs the RE Touch Lab. Assistant Professor (2013–2015) of Electrical and Computer Engineering at Drexel University. Post-Doctoral Fellow (2011–2012) at the Inst. of Intelligent Systems and Robotics, Sorbonne University. He has been employed in industrial R&D for sonar, speech recognition, and music DSP at several technology companies. His research interests include haptics, robotics, and soft electronics.



Anzu Kawazoe received the BS and MS degrees in Information Science from Soka University, Tokyo, Japan, in 2014 and 2016 respectively. She was assistant from April to August in 2017 at Keio University in Kanagawa, Japan. In 2017, she started working at the RE Touch Laboratory, University of California, Santa Barbara, where she is a Ph.D. student in Media arts and technology. Her research interests include the design and fabrication of novel haptic device and tactile psychophysics with applications in haptics, thermal devices and tactile illusion.



Gregory Reardon received the BS Degree in Applied Mathematics-Economics and the MM Degree in Music Technology from Brown University (2016) and New York University (2018), respectively. Since 2018, he has been a PhD student in the Media Arts and Technology Program at the University of California, Santa Barbara. His research interests include haptic and spatial audio perception and acoustics.



Erin Woo Erin Woo is an undergraduate student studying computing within the College of Creative Studies at the University of California, Santa Barbara. She was a EUREKA! Scholars undergraduate research intern at the RE-Touch Lab from June-August 2017. She spent the summer developing interactive virtual reality applications for finger-worn haptic devices and is currently pursuing a career in industry as a software engineer.



Massimiliano Di Luca received the PhD Degree in Cognitive Science from Brown University (2006) and the Laurea in Psychology from the Università di Trieste in 2000. He has been scientist at the Max Planck Institute for Biological Cybernetics in Tübingen, visiting scientist at Oculus, and research scientist at Facebook Reality Labs. He is Senior Lecturer at the University of Birmingham. His research interests include psychophysics and computational modeling of multisensory perception and action.