

Perceptual Space of Algorithms for Three-to-One Dimensional Reduction of Realistic Vibrations

Hojin Lee¹, *Member, IEEE*, Güneş Işık Tombak², Gunhyuk Park³, *Member, IEEE*, and Katherine J. Kuchenbecker⁴, *Fellow, IEEE*

Abstract—Haptics researchers often endeavor to deliver realistic vibrotactile feedback through broad-bandwidth actuators; however, these actuators typically generate only single-axis vibrations, not 3D vibrations like those that occur in natural tool-mediated interactions. Several three-to-one (321) dimensional reduction algorithms have thus been developed to combine 3D vibrations into 1D vibrations. Surprisingly, the perceptual quality of 321-converted vibrations has never been comprehensively compared to rendering of the original 3D signals. In this study, we develop a multi-dimensional vibration rendering system using a magnetic levitation haptic interface. We verify the system’s ability to generate realistic 3D vibrations recorded in both tapping and dragging interactions with four surfaces. We then conduct a study with 15 participants to measure the perceived dissimilarities between five 321 algorithms (SAZ, SUM, VM, DFT, PCA) and the original recordings. The resulting perceptual space is investigated with multiple regression and Procrustes analysis to unveil the relationship between the physical and perceptual properties of 321-converted vibrations. Surprisingly, we found that participants perceptually discriminated the original 3D vibrations from all tested 1D versions. Overall, our results indicate that spectral, temporal, and directional attributes may all contribute to the perceived similarities of vibration signals.

Index Terms—Magnetic levitation haptic interface, vibrotactile rendering, perceptual space, dimensional reduction algorithm.

I. INTRODUCTION

PHYSICAL interactions between real objects generate vibrations that help humans perceive material properties

Manuscript received 3 August 2021; revised 11 January 2022 and 2 April 2022; accepted 29 April 2022. Date of publication 11 May 2022; date of current version 28 September 2022. This work was supported in part by the Max Planck Society in Germany. This article was recommended for publication by Associate Editor Editor Dr. Yoshihiro Tanaka and Editor-in-Chief Dr. Marcia K. O’Malley upon evaluation of the reviewers’ comments. (*Corresponding authors: Gunhyuk Park; Katherine J. Kuchenbecker.*)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Ethics Council of the Max Planck Society under protocol F010A of the Haptic Intelligence Department’s Framework Agreement, and performed in line with the Declaration of Helsinki.

Hojin Lee and Katherine J. Kuchenbecker are with the Haptic Intelligence Department, Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany (e-mail: hojinlee@is.mpg.de; kjk@is.mpg.de).

Güneş Işık Tombak is with the Haptic Intelligence Department, Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany, and also with the Department of Information Technology and Electrical Engineering, ETH Zürich, 8092 Zürich, Switzerland (e-mail: gtombak@student.ethz.ch).

Gunhyuk Park is with AI Graduate School, Gwangju Institute of Science and Technology (GIST), Gwangju 61005, South Korea (e-mail: maharaga@gist.ac.kr).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TOH.2022.3174229>, provided by the authors.

Digital Object Identifier 10.1109/TOH.2022.3174229

and textures. Tool-mediated contacts evoke particularly strong vibrations because a stiff object experiences many transients when making and breaking contact. By providing vibrotactile feedback in computer-mediated applications, recent haptic displays take advantage of the human sensitivity to vibrations and thereby enrich user interactions [1]. Software developers synthesize or pattern artificial vibration models to alert the user to important events or provide context-sensitive information [2]. Applications developed for virtual reality [3], [4] and robotic teleoperation [5], [6] often seek to deliver realistic vibrotactile cues similar to those felt in real interactions.

Natural vibrations from tool-mediated interactions are generally three-dimensional (3D) and contain energy at a broad range of frequencies [5], [7]. To mimic natural interactions in virtual environments, developers usually implement realistic vibrotactile feedback using either a physics-based model designed by hapticians or a data-driven model established from vibratory signals recorded in real environments [8]. However, in practice, rendering realistic vibrations via typical haptic devices and actuators is more challenging than computing the signal to be rendered. Many force-feedback devices do not excel at rendering high-frequency vibrations [9], [10] due to the friction, flexibility, and mechanical backlash that their mechanisms contain. These characteristics somewhat limit the controllability of the device [11]. Therefore, commercial vibration actuators that can output arbitrary broad-bandwidth waveforms are often used to display realistic vibrotactile cues [1]; many such devices are electromagnetic and function similar to small audio speakers. However, these actuators generate vibrations along only one fixed axis, so they cannot immediately output a 3D vibration that has been synthesized or recorded. Although three such actuators can be combined to create 3D vibrations, the system’s cost, complexity, weight, inertia, and volume all increase; a single actuator is more practical.

To render realistic vibrations using 1D vibration actuators, several haptics researchers have been looking for an effective method to convert 3D vibrations to 1D vibrations, called a three-to-one (321) dimensional reduction algorithm [12]. The belief that such a conversion exists somewhat stems from the fact that a Pacinian corpuscle (PC, the main mechanoreceptor that senses high-frequency mechanical vibrations in human skin) is less sensitive to the direction of stimuli than other mechanoreceptors [13], [14]. The ideal 321 algorithm would preserve all useful information when converting 3D vibrations into 1D so that the resulting signal feels indistinguishable

from the original version. Various 321 algorithms have been developed and tested, from the simplest approach that keeps only a single axis aligned with the main interaction direction, e.g., [15], to more computationally intensive options such as DFT321 [12], which leverages the discrete Fourier transform.

Previously, Park and Kuchenbecker investigated the effectiveness of 321 algorithms via performance metrics and a perceptual study [16]. Their users touched various surfaces with a pen-shaped sensing tool, while a one-axis actuating tool displayed a dimensionally reduced version of the measured three-axis vibration to the other hand in real time. The users rated the perceptual similarity between the vibrations felt on the two tools for each algorithm and surface. Although some of the tested 321 algorithms generated 1D vibrations that felt reasonably similar to the original 3D vibrations, most received comparable scores, and none matched exceptionally well.

For a more detailed performance examination, this paper aims to characterize the shared perceptual space of the original 3D and converted 1D vibrations produced by common 321 algorithms. Many studies have used the perceptual space approach [17] for haptics because this tool adeptly showcases the similarities and differences in human perception of complex stimuli. Such a comparison requires a haptic interface that can render multi-dimensional vibrotactile feedback effectively. One feasible way to create such vibrations uses magnetic levitation; the frictionless, stiff, and backlash-free transmission of such devices gives them the potential to produce high-frequency vibrations [18]. Recently, Zhang *et al.* used an untethered stylus-like magnetic levitation device constructed from large coils and a small permanent magnet to render 1D vibrations along/around each of the 6D translational and rotational axes [19]. They investigated 1D vibration detection thresholds for each axis as well as the effects of the physical parameters of the stylus for a precision grasping posture [20]. Furthermore, they used perceptual results gathered with this device to propose a six-to-one (621) dimensional reduction matrix that can effectively convert 6D vibrations into 1D at a single frequency (108 Hz) [21] or at a range of frequencies from 20 to 250 Hz [22].

Similar to the approach of Zhang *et al.* [21], this paper introduces a multi-dimensional vibration rendering system built around magnetic actuation. After careful characterization, this system can provide realistic vibration feedback, playing broad-bandwidth vibrations simultaneously along multiple axes. The paper then presents our findings from the user study, establishing the perceptual space of 321 algorithms for realistic vibrations rendered by the proposed system.

II. MULTI-DIMENSIONAL VIBRATION RENDERING SYSTEM

A. Hardware

We used a commercial magnetic levitation haptic interface (Maglev 200, Butterfly Haptics LLC) [23] to create a multi-dimensional vibration rendering system (Fig. 1). The interface can render continuous 6-DoF wrench feedback $W = [F_x F_y F_z \tau_x \tau_y \tau_z]^T$ [see coordinate frame in Fig. 1(a)] with zero friction and no mechanical backlash. This device's simple and joint-free mechanical design enables it to consistently output high-frequency vibration signals. The peak force and force




Fig. 1. Multi-dimensional vibration rendering system. (a) A custom stylus containing a high-bandwidth three-axis accelerometer is attached to the Maglev 200's handle. The top-right box shows the coordinate frame for outputting force and torque. (b) The precision-grip hand posture used in this study.

resolution of the device are 40 N and 20 mN, respectively. After testing various servo rates and wrench commanding rates, we set both to 1 kHz to keep the delay small and constant. As a result, according to the Nyquist-Shannon sampling theorem, the system can produce continuous vibrations in a frequency range up to 500 Hz. Although it does not fully cover the sensitivity of PCs (approximately 20–1000 Hz), it covers a broad range that impedance-type haptic devices with multi-joint structures may not easily produce [24].

We considered the common situation where a user holds a stylus with a precision grip and feels a cutaneous vibration caused by a surface interaction at the lower end of the stylus [Fig. 1(b)]. To this end, we custom-designed and attached a stylus-like rigid manipulandum (11.5 cm, $\varnothing = 6$ mm, assembled from a 10-cm aluminum cylinder-shaped rod and a copper base) to the top of the device end-effector. The weight of the manipulandum (450 g) is light enough to be compensated by the software. An accelerometer (ADXL354, Analog Devices) is rigidly mounted inside the manipulandum to record 3D acceleration signals aligned to the translational axes of x , y , and z . A data acquisition board (USB-6343, National Instruments) is used to record the stylus' three-axis acceleration data at a sampling rate of 10 kHz.

B. Dynamic Compensation

To render high-fidelity realistic vibration signals, one must characterize the transfer function (TF) of the rendering device from force/torque to acceleration and compensate for its dynamics [25]. We can approximately cancel the dynamics of the system by applying an inverse TF (ITF) filter to the desired acceleration signal to calculate the command [26]. When this model adequately captures the dominant dynamics, we can expect the device output to be similar to the desired acceleration waveform.

The user grasps about 6 cm above the accelerometer located inside of the manipulandum. We observe that five different sources in the 6-DoF wrench (W) can independently contribute to the creation of each axial component of the 3D acceleration vector $a = [a_x a_y a_z]^T$ measured by the accelerometer. First, the 3-DoF force components (F_x , F_y , and F_z) mainly produce accelerations aligned with their respective axes (a_x ,

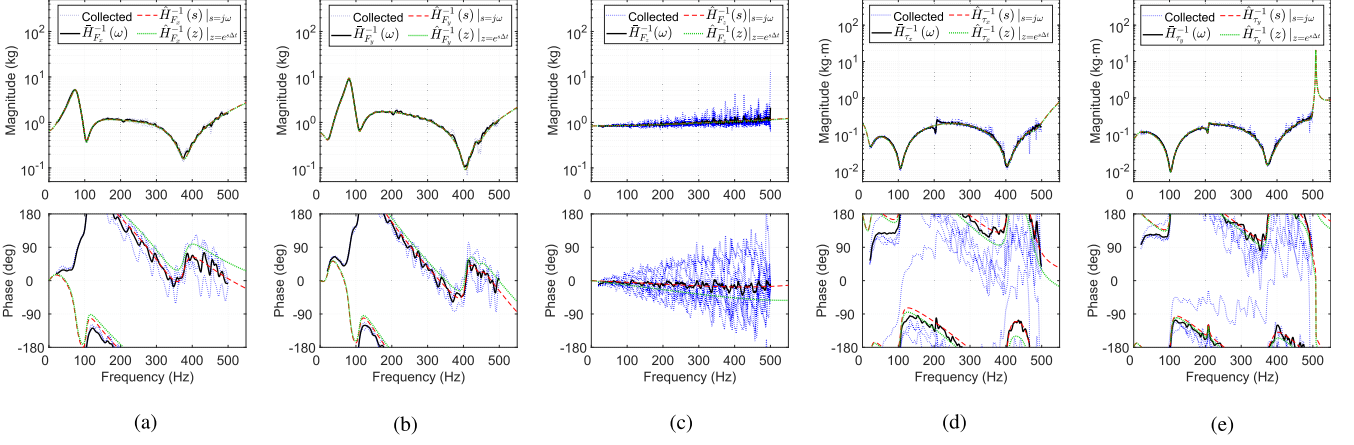


Fig. 2. ITF models for dynamic compensation. Blue dotted lines are trial-wise measurements. Black bold lines are smoothed nonparametric models, red dashed lines show continuous models, and green dotted lines show the final discrete models. (a) $H_{F_x}^{-1} : \frac{\mathcal{F}\{F_x\}}{\mathcal{F}\{a_x\}}$. (b) $H_{F_y}^{-1} : \frac{\mathcal{F}\{F_y\}}{\mathcal{F}\{a_y\}}$. (c) $H_{F_z}^{-1} : \frac{\mathcal{F}\{F_z\}}{\mathcal{F}\{a_z\}}$. (d) $H_{\tau_x}^{-1} : \frac{\mathcal{F}\{\tau_x\}}{\mathcal{F}\{a_y\}}$. (e) $H_{\tau_y}^{-1} : \frac{\mathcal{F}\{\tau_y\}}{\mathcal{F}\{a_x\}}$.

a_x , and a_x). Second, high-frequency torques around the x axis (τ_x) mainly generate stylus accelerations along y (a_y), just as torques around the y axis (τ_y) mainly cause accelerations along x (a_x). To limit the complexity of our model, we fit a single-input single-output (SISO) system with five independent TFs: H_{F_x} , H_{F_y} , H_{F_z} , H_{τ_x} (for a_y), and H_{τ_y} (for a_x), using the following data-driven procedure [27].

First, we command a 30-s, 10–500-Hz linear swept sinusoid signal on each force/torque source with a magnitude of 1 N or 0.1 Nm and measure the corresponding output acceleration. The identity of the user and how they grasp the stylus may have a measurable impact on the resulting transfer functions. However, changing the dynamic model according to each individual or hand posture would be inefficient in terms of data collection, model fitting, and device use. Instead, we chose the simple and practical solution of carefully fitting one dynamic model to a single individual’s varied dataset and employing it for all users. The selected participant conducted multiple measurement trials for each condition (three trials for F_x and F_y , ten for F_z , and five for τ_x and τ_y) to provide some variety in hand posture and to allow us to confirm that the data were repeatable. Then, based on the relationship between sweep input and output in each trial, we obtain a trial-wise frequency response of the system and its inverse form. For example, from a 1-N sweep input $F_{x,i}(t)$ and its measured output acceleration $a_{x,i}(t)$ in the i -th measurement trial, $H_i(\omega) = \frac{\mathcal{F}\{a_{x,i}(t)\}}{\mathcal{F}\{F_{x,i}(t)\}}$ and $H_i^{-1}(\omega) = \frac{\mathcal{F}\{F_{x,i}(t)\}}{\mathcal{F}\{a_{x,i}(t)\}}$ between $\omega_{\min} = 20\pi$ rad/s (10 Hz) and $\omega_{\max} = 1000\pi$ rad/s (500 Hz) are computed.

Next, a nonparametric ITF model $\bar{H}^{-1}(\omega)$ was obtained by smoothing all trial-wise frequency responses using an empirical TF estimate (ETFE) [28] after compensating for the time delay. Then, using the `tfest()` function in MATLAB, we fitted $\bar{H}^{-1}(\omega)$ into a continuous parametric ITF model $\hat{H}^{-1}(s)|_{s=j\omega}$ in a pole-zero representation such that

$$\hat{H}^{-1}(s) = e^{-\xi s} \frac{Z(s)}{P(s)}, \quad (1)$$

where the polynomials of $P(s)$ and $Z(s)$ are its poles and zeros, respectively. ξ is a time delay. $\hat{H}^{-1}(s)$ was established

by testing different numbers of poles and zeros to find the smallest error between $\bar{H}^{-1}(\omega)$ and $\hat{H}^{-1}(s)$. Then, we discretized $\hat{H}^{-1}(s)$ to get a discrete-time model $\hat{H}^{-1}(z)$ using the matched pole-zero (MPZ) transform, i.e., $z = e^{s\Delta t}$ [29]. $\Delta t = 0.1$ ms is a time interval of the accelerometer’s sampling rate (10 kHz). An identical ITF filter $\hat{H}^{-1}(\omega)$ can be implemented from the same poles and zeros of $\hat{H}^{-1}(z)$. By applying $\hat{H}^{-1}(\omega)$, we can transform each original acceleration command to a corresponding force or torque in the wrench command W .

C. Results

We conducted the characterization process via MATLAB (2019a, The MathWorks). The Bode plots of the five resulting ITF models in the frequency range of interest are presented in Fig. 2. We estimated the fitting accuracy using the root mean squared error (RMSE) and the normalized RMSE (NRMSE) between the empirical nonparametric model (\bar{H}^{-1}) and the resulting parameterized model (\hat{H}^{-1}), each calculated as follows:

$$\text{RMSE} = \sqrt{\sum_{\omega=\omega_{\min}}^{\omega_{\max}} (|\bar{H}^{-1}(\omega)| - |\hat{H}^{-1}(\omega)|)^2}, \quad (2)$$

$$\text{NRMSE} = \sqrt{\sum_{\omega=\omega_{\min}}^{\omega_{\max}} \left(\frac{|\bar{H}^{-1}(\omega)| - |\hat{H}^{-1}(\omega)|}{|\bar{H}^{-1}(\omega)|} \right)^2}. \quad (3)$$

Table I presents the order and fitting accuracy of each ITF model. All achieved a fitting accuracy over 90%, which is expected to sufficiently cancel out the dynamics of the system.

The planar force-based filters ($H_{F_x}^{-1}$ and $H_{F_y}^{-1}$) contain sharp peaks below 100 Hz [Fig. 2(a) and (b)]. When we apply these models to a vibration containing low-frequency components, the resulting force command may exceed the device limit. Therefore, to keep the commands within the device’s capabilities, we selected ITF filters as torque-based models for x and y ($H_{\tau_x}^{-1}$, $H_{\tau_y}^{-1}$) and a force-based model for z ($H_{F_z}^{-1}$).

TABLE I
THE NUMBER OF POLES AND ZEROS ($n(P)$ AND $n(Z)$) IN THE ITF MODELS AND THEIR FITTING ACCURACY

	$H_{F_x}^{-1}$	$H_{F_y}^{-1}$	$H_{F_z}^{-1}$	$H_{\tau_x}^{-1}$	$H_{\tau_y}^{-1}$
$n(P)$	6	9	1	11	7
$n(Z)$	11	12	5	13	12
RMSE	0.066 kg	0.086 kg	0.063 kg	0.009 kg·m	0.013 kg·m
NRMSE (%)	5.27	5.81	5.19	6.11	9.01

TABLE II

RELATIVE MAGNITUDE OF ACCELERATION OUTPUT ON THE ORTHOGONAL AXES (CROSSTALK) COMPARED TO THE CORRESPONDING DOMINANT AXIS FOR EACH OF OUR FIVE ACTUATION SOURCES

a_s	$R_{F_x}^{a_s}$	$R_{F_y}^{a_s}$	$R_{F_z}^{a_s}$	$R_{\tau_x}^{a_s}$	$R_{\tau_y}^{a_s}$
a_x	(100%)	10.50%	2.80%	6.96%	(100%)
a_y	15.49%	(100%)	3.37%	(100%)	6.07%
a_z	1.97%	3.17%	(100%)	3.06%	2.16%

D. Crosstalk Analysis

If the system does not follow our SISO assumption, its rendering performance will be poor due to interference across axes. We thus analyzed the undesired mechanical vibrations generated in the directions orthogonal to the desired axis (so-called crosstalk) for each considered force and torque input. Using the recorded sweep data and the same ETFE procedure, we obtained 15 different smoothed forward TFs, $\bar{H}_{W_s}^{a_s}(\omega)$, each describing the relationship from a single input force/torque command $W_s \in \{F_x, F_y, F_z, \tau_x, \tau_y\}$ in W to one of the single output accelerations $a_s \in \{a_x, a_y, a_z\}$ in 3D.

The relative crosstalk $R_{W_s}^{a_s}$ was then quantified as the ratio of the sum of the magnitude spectrum of $H_{W_s}^{a_s}$ on the target non-dominant axis to the sum of the magnitude spectrum of $H_{W_s}^{a_D}$ on its dominant axis, within our target frequency range from ω_{\min} to ω_{\max} , as follows:

$$R_{W_s}^{a_s} = \frac{\sum_{\omega=\omega_{\min}}^{\omega_{\max}} \left| \bar{H}_{W_s}^{a_s}(\omega) \right|}{\sum_{\omega=\omega_{\min}}^{\omega_{\max}} \left| \bar{H}_{W_s}^{a_D}(\omega) \right|}, \quad (4)$$

where a_D represents the acceleration along the dominant axis for W_s (e.g., $a_D = a_x$ for both F_x and τ_y). Table II shows the resulting values: the crosstalk generated by each actuation source is much smaller than the vibration the same command causes on its dominant axis (always less than 16%), thereby justifying our SISO assumption. The results further show that $R_{F_x}^{a_y}$ and $R_{F_y}^{a_x}$ are about two times larger than $R_{\tau_y}^{a_y}$ and $R_{\tau_x}^{a_x}$. This finding that torque commands cause less crosstalk than comparable force commands reinforces our decision to use torques to render lateral accelerations with this haptic device.

III. DATA ACQUISITION

We recorded a set of real three-axis vibration signals, $\mathbf{a}(t) = [a_x(t) \ a_y(t) \ a_z(t)]^T$. Fig. 3(a) demonstrates the data collection procedure. The accelerometer, the data acquisition board, the shape of the stylus, the hand posture, and the grasping position on the stylus were all the same as in Section II-A.




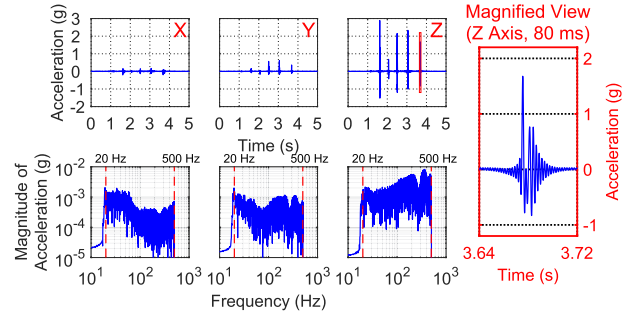
Fig. 3. Data acquisition of realistic vibrations using a combination of two movements and four surface materials. (a) The recording device assembled from a stylus and an accelerometer. (b) The four surfaces. (c) The tapping movement on acrylic. (d) The dragging movement on sandpaper.

Similar to [16], we selected four surfaces (100 mm \times 100 mm) to provide different hardness and roughness sensations during a stylus-based interaction. The tiles are made of sandpaper, acrylic, rough paper, and rubber, and each material represents the tactile sensation of hard/rough (HR), hard/smooth (HS), soft/rough (SR), and soft/smooth (SS), respectively [Fig. 3(b)].

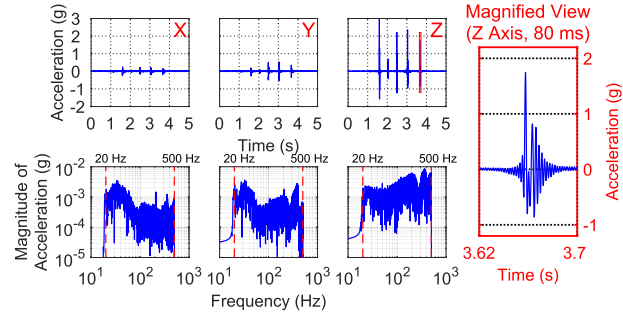
In [16], the experimenter interacted with each surface with a composite movement, tapping its surface three times and then dragging the stylus in a circle three times. In contrast, we separated tapping and dragging because each move generates accelerations mainly in a different direction. In our tapping recordings (denoted as T), the experimenter hits perpendicular to the surface five times with the instrumented stylus [Fig. 3(c)]. Since tapping is a 1D movement along the surface normal, this interaction results in a series of impulse-like vibrations mainly along the z direction. In dragging (denoted as D), the experimenter moves the stylus along the surface in an inverted-L-shaped stroke, once along the x -axis followed by once along the y -axis [Fig. 3(d)]. Dragging the stylus generates a vivid 3D vibration that depends on the geometry and material properties of both the tool tip and the surface; most of the signal energy is in the direction of travel and perpendicular to the surface.

We collected a total of eight interactions with the two different movements and four surfaces: THR, THS, TSR, TSS, DHR, DHS, DSR, and DSS. By repeating each measurement four times, we obtained a collection of 32 acceleration-based vibrations. During each interaction, a video was recorded simultaneously with sound using a standard webcam; see the examples of these recordings in part 1 (tapping) and part 2 (dragging) of the supplementary video. All of the collected vibration data and the corresponding videos can be found on IEEE DataPort [30].

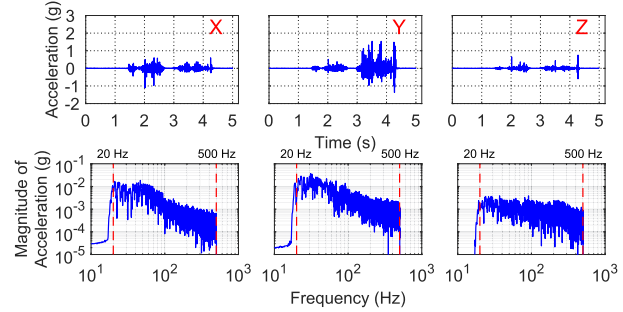
To output vibrations, we command a wrench $W_{in} = [0 \ 0 \ F_z \ \tau_y \ \tau_x \ 0]^T$ to the system, where W_{in} is obtained by applying the selected ITF filters to the corresponding input acceleration \mathbf{a}_{in} after application of a 20–500 Hz band-pass filter. The experimenter held the stylus and recorded the corresponding output signal \mathbf{a}_{out} during each execution. Fig. 4 illustrates input-output examples in situations that are somewhat challenging to render: tapping on a hard/smooth surface (THS) and dragging on a soft/rough surface (DSR). These examples



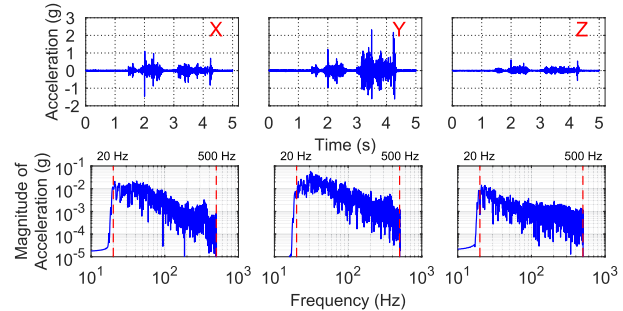
(a) \mathbf{a}_{in} : Input commanded to the system for tapping on acrylic, the hard smooth surface (THS). The red box at right shows a magnified view of the fifth tap in the time domain.



(b) \mathbf{a}_{out} : Output rendered by the system for tapping on acrylic, the hard smooth surface (THS). The red box at right shows a magnified view of the fifth tap in the time domain.



(c) \mathbf{a}_{in} : Input commanded to the system for dragging on rough paper, the soft rough surface (DSR).



(d) \mathbf{a}_{out} : Output rendered by the system for dragging on rough paper, the soft rough surface (DSR).

Fig. 4. Examples of input and output vibrations for two sample challenging interactions. (a) and (b): Tapping on acrylic (THS). (c) and (d): Dragging on rough paper (DSR). In each lettered subplot, the upper part visualizes time-domain data, and the lower part shows its corresponding magnitude spectrum.

showcase that \mathbf{a}_{in} and \mathbf{a}_{out} largely match in both time and frequency domains, verifying that dynamic compensation enables the system to generate realistic vibrations.

IV. ALGORITHMS AND OBJECTIVE PERFORMANCE METRICS

A. Dimensional Reduction Algorithms

This section introduces five computational algorithms commonly used to convert an original 3D acceleration signal $\mathbf{a}(t)$ into a reduced 1D signal $\hat{a}(t)$ for vibration rendering. The algorithms are described in more detail in [12], [16].

Single Axis Along Z (SAZ) is the simplest algorithm utilizing only one of the three measured acceleration directions as the reduced form, similar to a 1D accelerometer. We choose the single axis to be the surface normal, i.e., $\hat{a}(t) = a_z(t)$.

Sum of Components (SUM) simply adds all three acceleration axes into one, i.e., $\hat{a}(t) = \sum_{i \in \{x,y,z\}} a_i(t)$. This method is also relatively easy to compute and effective in transmitting all axes of $\mathbf{a}(t)$, but some information will be lost or distorted when the different axes cancel each other out.

Vector Magnitude (VM) computes the Euclidean norm of the three acceleration axes, i.e., $\hat{a}(t) = \sqrt{\sum_{i \in \{x,y,z\}} a_i^2(t)}$.

Although its computation is easy and fast, this algorithm has several limitations. First, the square and square-root computations in the time domain are nonlinear and distort the original frequency spectrum. Second, since the reduced signal $\hat{a}(t)$ consists of positive values only, intermediate processing is required before rendering (e.g., high-pass filtering).

Discrete Fourier Transform (DFT) utilizes the DFT to obtain frequency components $\{A_i(f)\}_{i \in \{x,y,z\}}$ from each component $\{a_i(t)\}_{i \in \{x,y,z\}}$ in $\mathbf{a}(t)$:

$$A_i(f) = \mathcal{F}\{a_i(t)\} = |A_i(f)|e^{j\Theta_i(f)}, \quad (5)$$

where $|A_i(f)|$ and $\Theta_i(f)$ denote the spectral magnitude and the phase information, respectively. The algorithm then synthesizes three frequency components into $\hat{A}(f) = |\hat{A}(f)|e^{j\hat{\Theta}(f)}$, where

$$|\hat{A}(f)| = \sqrt{\sum_{i \in \{x,y,z\}} |A_i(f)|^2}, \quad (6)$$

$$\hat{\Theta}(f) = \angle \left(\sum_{i \in \{x,y,z\}} A_i(f) \right), \quad (7)$$

respectively. The reduced 1D signal $\hat{a}(t)$ is transformed back from synthesized frequency components $\hat{A}(f)$ using the inverse DFT, i.e., $\hat{a}(t) = \mathcal{F}^{-1}\{\hat{A}(f)\}$. This algorithm mainly targets maintaining the spectral energy of $\mathbf{a}(t)$ in $\hat{a}(t)$. To preserve temporal signal properties and allow for real-time operation, instead of using the entire waveform $\mathbf{a}(t)$, a windowed time segment $\tilde{\mathbf{a}}(t)$ can be used. We extract $\tilde{\mathbf{a}}(t)$ from $\mathbf{a}(t)$ by applying a 10-ms Hann window with an overlapping ratio of 0.5. The processed segments $\tilde{\hat{a}}(t)$ after IDFT are synthesized back to reproduce $\hat{a}(t)$.

Principal Component Analysis (PCA) finds the principal axis of the 3D acceleration and projects the three-axis data onto that principal axis, i.e.,

$$\mathbf{w} = \text{PCA}(\mathbf{M}), \quad (8)$$

$$\hat{\mathbf{a}}(t) = \mathbf{w}^T \mathbf{a}(t), \quad (9)$$

where $\text{PCA}(\mathbf{M})$ finds a unit vector along the principal axis of the distribution in a data matrix \mathbf{M} , which is composed of the acceleration vector \mathbf{a} over time. This algorithm is similar to **SAZ**, except the chosen axis depends on the signal itself. Windowed time segments $\tilde{\mathbf{a}}(t)$ can be leveraged to respond to directional changes over short time intervals [16]. For consistency, the same window used in DFT was used for PCA.

B. Spectral and Temporal Matching Metrics

To quantitatively compare the performance of 321 algorithms, Landin *et al.* first introduced a spectral matching score and a time matching score, which compare the original 3D signal $\mathbf{a}(t)$ and the reduced 1D signal $\hat{\mathbf{a}}(t)$ in the frequency and time domains [12]. Adopting the same idea, several researchers have developed or improved similar metrics to assess the quality of vibrotactile signals [16], [31].

Spectral Matching (M_{sm}) computes the spectral similarity between the original 3D signal and the reduced 1D signal [12]. We use Park and Kuchenbecker's enhanced version of this metric, incorporating the perceptual characteristics of differential thresholds (DLs) of vibration frequencies [16]. First, the signals are normalized as

$${}^N \mathbf{a}(t) = \frac{\mathbf{a}(t)}{a_{\max}}, \quad (10)$$

$${}^N \hat{\mathbf{a}}(t) = \frac{\hat{\mathbf{a}}(t)}{a_{\max}}, \quad (11)$$

where $a_{\max} = \max(|\mathbf{a}(t)|)$. We then obtain the normalized frequency spectra $\{{}^N A_i(f) = \mathcal{F}\{a_i(t)\}\}_{i \in \{x,y,z\}}$ and ${}^N \hat{A}(f) = \mathcal{F}\{\hat{\mathbf{a}}(t)\}$ using DFT.

We then decompose the entire frequency range $f_{\min} = 20$ Hz to $f_{\max} = 500$ Hz into n different intervals $\{f : f_n \leq f < f_{n+1}\}$, updating f_n as

$$f_{n+1} = \begin{cases} 1.1f_n, & \text{if } 1.1f_n < f_{\max}, \\ f_{\max}, & \text{otherwise.} \end{cases} \quad (12)$$

until f_{n+1} reaches f_{\max} . Equation (12) is inspired by the perceptual property that a frequency increased by 10% from the reference frequency is still below the DL at moderate vibration magnitudes of 20 dB sensation level (dB SL) [32]. The dissimilarity value d_n is computed for each interval as

$$d_n = \frac{\sum_{f \in F_n} \left(|{}^N A(f)|^2 - |{}^N \hat{A}(f)|^2 \right)}{\sum_{f \in F_n} |{}^N A(f)|^2} \quad (13)$$

where $|{}^N A(f)|^2 = \sum_{i \in \{x,y,z\}} |{}^N A_i(f)|^2$. We then average all d_n 's to compute the mean spectral difference D between the original and reduced signals.

The decomposition process weights different human sensitivities according to different frequencies. However, when the

frequency resolution is not precise, such d_n 's depend on each individual value of the frequency components in the spectrum (especially at low frequencies); this causes an unexpected distortion of the actual values [16]. To prevent this side-effect, we acquire the spectral differences D_{20} and D_{21} , each of which starts from different frequencies, $f_1 = f_{\min} = 20$ Hz or $f_1 = 21$ Hz (purposefully separated by halving the initial DL), respectively. Finally, we compute M_{sm} as

$$M_{sm} = \begin{cases} 1 - \bar{D}, & \text{if } \bar{D} < 1, \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

where the smoothed value $\bar{D} = \frac{1}{2}(D_{20} + D_{21})$. Negative values are changed to zero to limit the score to the 0–1 range.

Temporal Matching (M_{tm}) computes the temporal similarity between the original and reduced signals. M_{tm} is computed as

$$M_{tm} = \frac{1}{3} \sum_{i \in \{x,y,z\}} \frac{|{}^N a_i(t) \star {}^N \hat{a}(t)|}{\sqrt{{}^N a_i(t) \star {}^N a_i(t)} \cdot \sqrt{{}^N \hat{a}(t) \star {}^N \hat{a}(t)}}, \quad (15)$$

where the operator \star denotes the cross-correlation of two signals evaluated at zero time delay. Equation (15) averages the absolute values of the normalized cross-correlations (in the 0–1 range) across the three axes.

C. Objective Evaluation

Before investigating the perceptual space of these 321 methods, we evaluated their objective performance with both matching metrics. Using our multi-dimensional rendering system, we rendered 160 vibrations (32 recorded 3D acceleration signals converted to 1D by the five algorithms) along the z axis and measured the acceleration while a user held the stylus.

Fig. 5 depicts the means of M_{sm} and M_{tm} between the input and output accelerations, averaged across the four different measurements for each interaction. The input and the rendered output had similar matching scores, again showcasing the high fidelity of our rendering approach. These objective scores are ratio-scaled data; therefore, to study the effects of *surface material* and *321 algorithm*, we performed a parametric test of two-way analysis of variance (ANOVA) on each of the four datasets (M_{sm} or M_{tm} for tapping or dragging). Every analysis reported that both main effects and their interaction were statistically significant at $\alpha = 0.05$ ($p < 0.05$).

For more detailed post-hoc analysis, we additionally conducted Tukey's HSD test (Fig. 5). We found that most of our results were consistent with the previous studies of 321 algorithms ([12] and [16]). The overall ranking of M_{sm} in tapping, which mainly occurs along the z axis, was $\text{DFT} \geq \text{PCA} \geq \text{SAZ} > \text{SUM} > \text{VM}$, where the highest is best. VM consistently showed the lowest M_{sm} , followed by SUM. DFT showed the highest score, which indicates that the normalized frequency spectrum of DFT was the most similar to the original 3D signal. In dragging interactions, the overall ranking of M_{sm} was $(\text{DFT} \approx \text{PCA}) > \text{SUM} > \text{VM} > \text{SAZ}$. SAZ, which is limited to display only the vibrations orthogonal to the surface, showed a considerable limitation in the presentation of spectral information in dragging. DFT and PCA were similarly rated as the highest in

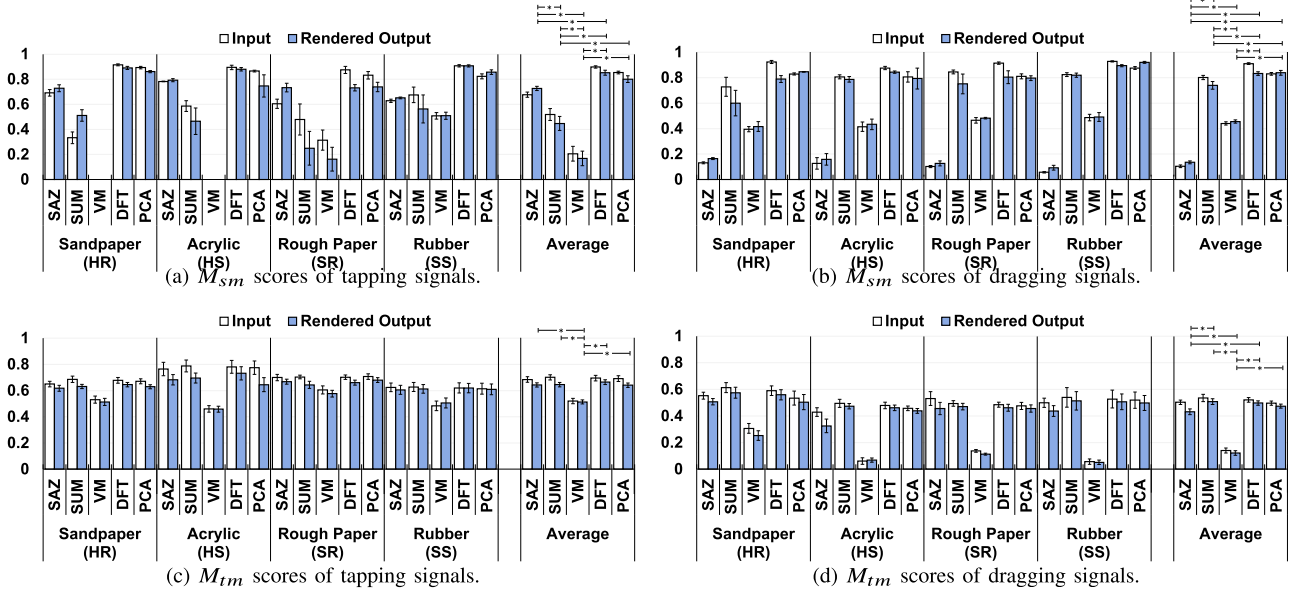


Fig. 5. Means of the spectral (M_{sm} ; (a) and (b)) and temporal (M_{tm} ; (c) and (d)) matching scores across four repeated trials for each surface interaction and each 321 algorithm. The rightmost graphs show the mean score across the four surfaces for each interaction. Asterisks (*) indicate significant differences in Tukey's test on the rendered output data ($\alpha = 0.05$).

M_{sm} , contending with each other to be the best matching algorithm over all types of surface material. In both tapping and dragging, the overall ranking of M_{tm} scores varied significantly depending on the surface texture type, because most methods showed similar scores. VM consistently achieved the lowest score for M_{tm} .

In summary, the spectral similarity between dimensionally reduced acceleration signals and the original signal is different from method to method. DFT or PCA showed the best score in both interactions; thus, these methods can be expected to be the best dimensional reduction methods to keep the spectral characteristics of the original 3D signal. VM was found to be the worst method for spectral match in tapping and temporal match in dragging. All other evaluated dimensional reduction algorithms show a similar capability to maintain the temporal properties of the original signal. Their scores are lower in dragging interactions than tapping interactions, meaning that the algorithms may lose more temporal information when applied to rich 3D signals compared to mainly 1D signals.

V. PERCEPTUAL EXPERIMENT

We conducted an experiment to compare the cutaneous perception of vibrotactile stimuli created by different 321 algorithms. All vibration signals were rendered by the multi-dimensional system using dynamic compensation.

A. Participants

A total of fifteen participants (six males and nine females; 11 right-handed and four left-handed; 24–41 years old, M: 31.3, SD: 5.6) participated in this experiment. The location of the haptic device was rearranged before the experiment so each participant could use their dominant hand comfortably. All participants reported no visual, auditory, or sensory-motor

disabilities. This experiment was approved by the Ethics Council of the Max Planck Society under the Haptic Intelligence Department's framework agreement (F010A). All participants signed a consent document after being informed about the procedure. Participants not employed by the Max Planck Society were compensated 20 EUR (8 EUR per hour).

B. Stimuli

Our recorded dataset of real vibrations ([30]; see Section III) includes four examples for each surface interaction. From each set, we selected the example whose renderings best showed the trend of the average M_{sm} across the five algorithms in Section IV-C. Each representative recording was converted to six different experimental stimuli, including the original 3D vibration (ORI) and 1D vibrations dimensionally reduced by the five 321 algorithms (SAZ, SUM, VM, DFT, PCA). All 1D vibrations were displayed using only the z -axis, which is aligned with the stylus. The resulting set of 48 wrench commands (4 surfaces \times 2 interactions \times 6 algorithms) was used in the study.

C. Procedure

Participants were seated in front of the system in a comfortable posture and instructed to stretch their dominant hand to grasp the stylus with a precision grip (Fig. 6). The view of their hand and the haptic device was blocked by a curtain. They also wore noise-isolating headphones during the experiment to block auditory cues from the device and other noise sources.

During the experiment, the experimenter monitored the participant and corrected their hand posture if it changed in order to minimize the trial-wise variations in the system dynamics. Each block of the study featured one interaction and consisted of a pair of consecutive sessions. The first session in the pair




Fig. 6. A right-handed participant performing the perceptual experiment. The participant’s hand behind the curtain is presented in translucency to show the hand posture.

was a practice session to let participants accustom themselves to all of the stimuli. Thus, only the responses in the second session were used for the analysis. The order of the eight featured interactions was randomized. Each session contained 15 comparison trials to compare each pair of the six algorithms. The order of the stimuli was also randomized.

In each trial, the participant was presented with two 5-s vibration stimuli and was instructed to rate their pairwise dissimilarity. Each vibration was mapped onto either the **A** or **B** buttons on a graphical interface. After a participant clicked either **A** or **B** using the mouse with the non-dominant hand, the corresponding vibration was played on the device. In synchrony with the rendered vibration, the associated video recording was displayed in a separate pop-up window, and the sound recording was simultaneously played on the headphones. The **A** and **B** buttons were disabled after being clicked so that the participant could feel each vibration only once.

We followed the general protocol of absolute magnitude estimation (on a free scale without a standard stimulus or a modulus) [33]. After feeling both stimuli, the participant was asked to enter a non-negative dissimilarity score for the two experienced stimuli, using the number keypad with the non-dominant hand. The participant was instructed to enter zero if the two stimuli felt identical. The comparison procedure is shown in part 3 of the supplementary video. A participant could proceed to the next trial or session by clicking the **Next** button. Previous responses could be changed by clicking the **Previous** button, but the other buttons remained disabled to prevent the participant from feeling the previous stimuli again.

After completing two consecutive sessions for the same condition, a participant was asked to take a break for at least one minute to prevent fatigue and sensory adaptation. A participant could take additional breaks between sessions. It took about 7.5 minutes to finish each session. Thus, the whole experiment was completed in two and a half hours.

D. Data Analysis

Since humans tend to use absolute scales rather than ratio scales for judging stimuli, this absolute magnitude estimation protocol avoids the contextual effects of natural human

biases [34]. We first normalized the dissimilarity scores of each participant to a 0–100 scale using his or her maximum dissimilarity rating across all test trials to minimize individual differences that might exist. Then, the normalized dissimilarity scores in each interaction type were averaged across the participants and formatted into a dissimilarity matrix. These absolute dissimilarity values ensure ratio-scaled psychophysical data [35], rather than the ordinal or interval scales used in ordinary human response tests [36]. Therefore, we used each dissimilarity matrix as the input to a metric multi-dimensional scaling (MDS) to establish a corresponding perceptual space. During the computation of MDS, we assessed the goodness-of-fit of every dimension using Kruskal’s stress values. The stress value ranges from 0–1, and smaller indicates a better fit. A detailed procedure for MDS can be found in [37].

Although perceptual spaces implicitly visualize the relationship between stimuli, the axes generally do not coincide with any dimensions that are meaningful for interpretation [38]. Therefore, further investigation is usually conducted by analyzing additional ratings using subjective measures [39]. The most common approach is discovering semantic information related to perception by using adjectives as descriptors [40], [41]. In contrast, some studies have applied physical properties directly as descriptors to find systematic relationships between physical and perceptual properties because the responses to construct perceptual spaces are based on the apparent physical properties in tactile stimuli [36], [42].

Similarly, we aim to determine which characteristics of dimensional reduction algorithms can lead to perceived discrepancies. Therefore, we adopt several ratings related to the physical quality of vibrations to clarify the relationship between perceptual and physical spaces. First, we computed the matching scores M_{sm} and M_{tm} , which show the spectral and temporal similarity to the original vibrations for all stimuli used in the experiment. $M_{sm} = 1$ and $M_{tm} = 1$ were assumed for ORI. Second, we consider the energy (or power) of the vibration signal transmitted to the skin mechanoreceptors as another major factor in recognition of perceived differences [36]. We calculated the signal energy E for each vibration:

$$E = \sum_{f \in F} \sum_{i \in \{x,y,z\}} |A_i(f)|^2, \quad (16)$$

where F is the frequency range that can be rendered and transmitted to the participant by our rendering system (20–500 Hz). Then, we applied multiple linear regression analysis with all 2D positions in each perceptual space as independent variables, using M_{sm} , M_{tm} , or E as dependent variables.

Lastly, to generalize and illustrate the perceptual relationship between 321 algorithms more clearly, we applied Procrustes analysis to bridge perceptual and physical properties. The Procrustes analysis primarily aims to map two representations onto each other as closely as possible by using only linear transformations (translation, scaling, reflection, and rotation, which are valid operations because the perceptual space yields only relative positions of percepts) [43], [44]. Therefore, the optimal transformation matrix to fit the points in a particular shape can

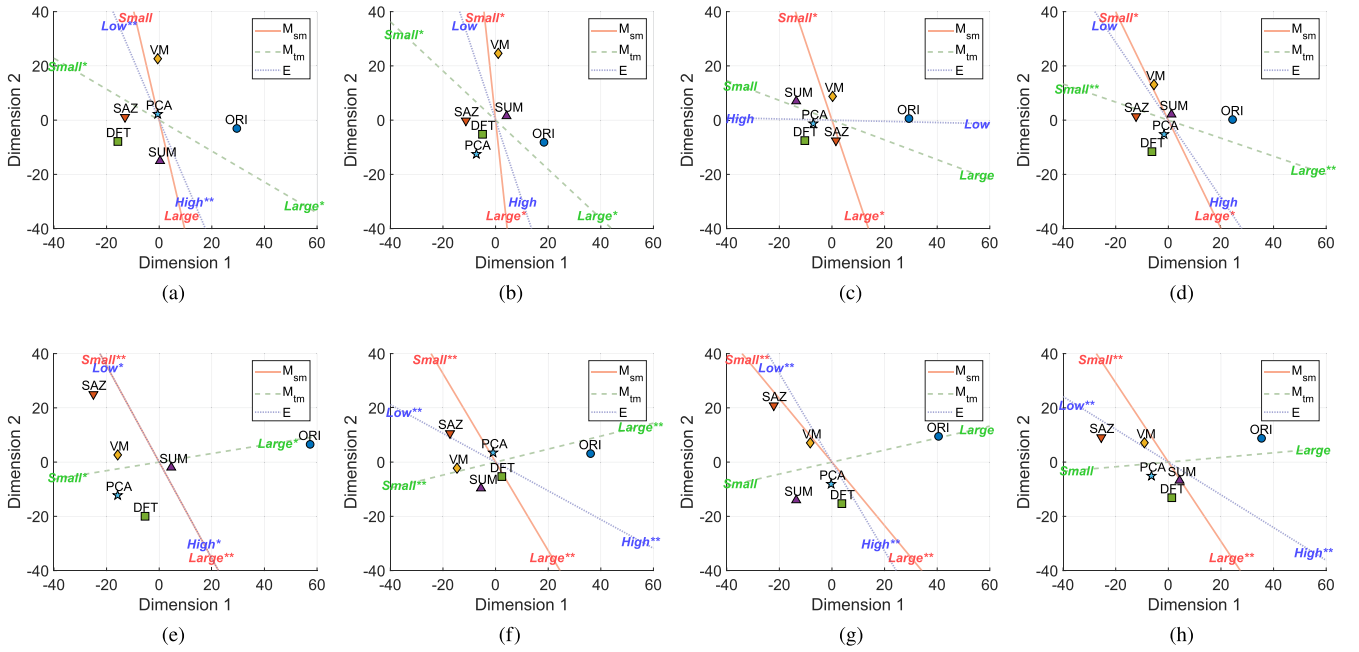


Fig. 7. Perceptual spaces corresponding to the eight types of surface interaction. Red solid and green dashed lines illustrate the axes regressed by M_{sm} and M_{tm} of each element, respectively. Blue dotted lines show the axis regressed by the signal energy E . All lines are drawn to cross the center of mass of the six percepts. Asterisks (*) and double asterisks (**) indicate statistical significance acquired from the regression analysis ($\alpha = 0.10$ and $\alpha = 0.05$, respectively). (a) P_{THR} : Tapping on Sandpaper. (b) P_{THS} : Tapping on Acrylic. (c) P_{TSR} : Tapping on Rough Paper. (d) P_{TSS} : Tapping on Rubber. (e) P_{DHR} : Dragging on Sandpaper. (f) P_{DHS} : Dragging on Acrylic. (g) P_{DSR} : Dragging on Rough Paper. (h) P_{DSS} : Dragging on Rubber.

be suggested by the analysis. As mentioned before, most perceptual spaces show a relationship close to orthogonal between the two M_{sm} and M_{tm} axes. Hence, we fitted each perceptual MDS output to each physical space of M_{sm} and M_{tm} scores by linearly transforming them using a Procrustes function in MATLAB. This Procrustes transformation into the combined physical-perceptual space strongly preserves the relative locations of points in the perceptual space and provides additional information about their tendencies along the M_{sm} axis and the M_{tm} axis. Because this transformation is a linear approximation, the stimuli may be shown at locations that are differently from their original two-dimensional locations in the M_{sm} - M_{tm} plane. The Procrustes function also outputs a goodness-of-fit value that calculates the sum of squared errors between the original points and the fitted points, where lower values are better.

E. Results and Discussion

There exists a trade-off between the difficulty of the visualization in perceptual spaces and the accuracy of the MDS result. Therefore, stress values lower than 0.15 are commonly recommended for the goodness-of-fit [45]. For a 2D representation, we achieved stress values of 0.125 (THR), 0.109 (THS), 0.101 (TSR), 0.148 (TSS), 0.073 (DHR), 0.058 (DHS), 0.088 (DHS), and 0.070 (DSS). These low stress values indicate that two dimensions sufficiently capture the relationships among all stimuli; 2D visualizations are also beneficially easy to understand. We plot a total of eight different 2D perceptual spaces (denoted as P_X for the surface interaction X) individually in Fig. 7.

The Euclidean distance between two stimuli shows the average perceived difference between that pair of corresponding vibrations, i.e., the perceived difference between those two 321 algorithms. Thus, one can discover some meaningful insights by analyzing pairwise distances.

First, we computed and summarized the distances from every 321 algorithm to ORI in Table III. This metric simply reveals the perceived difference from each stimulus to the original 3D vibration. Most of the perceptual spaces showed a distance between ORI and all 1D vibrations that is larger than the distances between the 1D vibrations, particularly in the dragging conditions, which involve 3D signals. The average distance also appears larger in the interactions with rough textures than with smooth ones. In the tapping condition, the overall ranking of mean distances from ORI to each dimensional reduction algorithm across four different textures was $DFT > VM > SAZ > PCA > SUM$. However, this ranking hardly matches with the ranking for each surface type. While SUM or PCA was located closest to ORI in three perceptual spaces (P_{THR} , P_{THS} , and P_{TSS}), SAZ was the nearest in P_{TSR} . In the dragging condition, the overall ranking of mean distances between ORI and each algorithm was $SAZ > VM > PCA > SUM > DFT$. Unlike the tapping conditions, this result was consistent with the individual rankings; DFT or SUM was located closest or second closest to ORI. DFT was usually located slightly closer to ORI than PCA (except P_{DSR}), whereas SUM showed a similar distance to ORI with DFT and PCA but had more variation across surfaces. In all dragging conditions, SAZ and VM had the most and the second most different perception from ORI, respectively.

Second, we computed and summarized the distances between SAZ and all other stimuli in Table IV because the displacement

TABLE III

EUCLEIDEAN DISTANCES FROM ORI TO OTHER ALGORITHMS IN THE PERCEPTUAL SPACES. A DARKER BACKGROUND COLOR CORRESPONDS TO A LONGER DISTANCE

Perceptual Space	Distance to ORI					Average
	SAZ	SUM	VM	DFT	PCA	
P _{THR}	44.74	32.55	41.03	47.29	32.21	39.56
P _{THS}	31.49	17.98	38.18	23.68	27.05	27.68
P _{TSR}	29.79	45.17	31.54	41.61	37.95	37.21
P _{TSS}	37.79	24.18	33.62	33.98	26.96	31.31
Average	35.95	29.97	36.09	36.64	31.04	
P _{DHR}	88.81	56.28	77.07	71.51	78.75	74.49
P _{DHS}	55.58	45.33	52.75	36.05	38.28	45.60
P _{DSR}	66.21	61.60	49.88	46.44	46.39	54.10
P _{DSS}	63.26	36.49	46.16	42.67	46.35	46.99
Average	68.47	49.93	56.46	49.17	52.45	

TABLE IV

EUCLEIDEAN DISTANCES FROM SAZ TO OTHER ALGORITHMS IN THE PERCEPTUAL SPACES. A DARKER BACKGROUND COLOR CORRESPONDS TO A LONGER DISTANCE

Perceptual Space	Distance to SAZ				Average
	SUM	VM	DFT	PCA	
P _{THR}	21.81	26.33	9.50	13.02	17.66
P _{THS}	15.90	28.00	8.82	13.40	16.53
P _{TSR}	21.77	16.95	12.05	11.18	15.49
P _{TSS}	13.73	13.89	14.45	13.35	13.85
Average	18.30	21.29	11.20	12.74	
P _{DHR}	41.55	25.20	51.40	40.54	39.67
P _{DHS}	24.21	13.53	26.19	18.41	20.59
P _{DSR}	36.36	20.58	46.25	37.63	35.21
P _{DSS}	35.08	17.20	35.88	24.03	28.05
Average	34.30	19.13	39.93	30.15	

of SAZ from the other 1D vibrations is one observable factor shown differently in the perceptual spaces between tapping and dragging. For tapping, the overall ranking of average distances from SAZ was VM > SUM > PCA > DFT, and this ranking is more prominently shown in the conditions involving hard surfaces (THR and THS). In dragging interactions, the perceptual differences between SAZ and other algorithms were greater than in tapping interactions, and these differences became even larger on rougher surfaces (DHR and DSR). The overall ranking of the means of this distance metric was consistent as DFT > (PCA or SUM) > VM across surfaces. Therefore, DFT was perceived as most different from SAZ, whereas VM was perceived the most similar to SAZ.

We depicted potential descriptive axes acquired by the multiple linear regression analysis in every perceptual space in Fig. 7. The standardized regression coefficients were used to represent the slope of each axis. Table V summarizes the detailed statistical data. To provide more freedom of interpretation, we present the results using two different levels of significance, one at a typical level at $\alpha = 0.05$ and the other at the less conservative level of significance at $\alpha = 0.10$; this second level indicates only marginal significance and should be interpreted with caution.

The results showed that the relative 2D positions of stimuli arranged in the perceptual spaces are closely related to the M_{sm} score or the signal energy E , with high R^2 values and significant p -values. In three perceptual spaces for tapping (P_{THS}, P_{TSR}, and P_{TSS}), M_{sm} was the most descriptive with the highest correlation with the spatial arrangement of the stimuli. E was the

TABLE V

STATISTICS OF MULTIPLE LINEAR REGRESSION ANALYSIS FOR M_{sm} , M_{tm} , AND E

Perceptual Space	M_{sm}		M_{tm}		E	
	R^2	p -value	R^2	p -value	R^2	p -value
P _{THR}	0.448	0.410	0.794	0.094*	0.974	0.004**
P _{THS}	0.832	0.069*	0.860	0.052*	0.416	0.447
P _{TSR}	0.855	0.055*	0.772	0.109	0.145	0.791
P _{TSS}	0.856	0.055*	0.872	0.046**	0.440	0.419
P _{DHR}	0.907	0.028**	0.806	0.085*	0.846	0.061*
P _{DHS}	0.875	0.044**	0.886	0.039**	0.975	0.004**
P _{DSR}	0.965	0.007**	0.622	0.232	0.902	0.031**
P _{DSS}	0.883	0.040**	0.587	0.265	0.938	0.016**

*: significance $\alpha = 0.10$, **: significance $\alpha = 0.05$.

most significant metric only in P_{THR}. Both M_{sm} and E were also effective at explaining the distribution of positions in most of the dragging perceptual spaces, except P_{DHR}, which achieved marginal significance. M_{tm} also showed a correlation to the relative positions to explain some relationship in most of the perceptual spaces, particularly in the perceptual spaces involving hard surfaces (P_{THR}, P_{THS}, P_{DHR}, and P_{DHS}). P_{TSS} was the only perceptual space for a soft-surface interaction that was highly correlated to M_{tm} . The results indicate that the perceptual differences and the physical (spectral and temporal) differences between algorithms are correlated with each other to some extent.

The M_{sm} axis (spectral similarity) and the E axis (overall spectral energy) were often close to parallel, indicating a strong correlation. On the other hand, the two axes of M_{sm} and M_{tm} (temporal similarity) were closer to orthogonal in most perceptual spaces, which implies a complementary relationship. Except for P_{TSR}, most perceptual spaces showed similar trends, and the perceptual spaces associated with dragging showed them more clearly. We further investigated the exceptional condition TSR and found that this is the only condition where the signal energy E of SUM ($0.70 \text{ m}^2/\text{s}^3$) becomes unusually larger than that of ORI ($0.43 \text{ m}^2/\text{s}^3$). This case is a particular example that reveals the weakness of SUM in maintaining the total energy, which means that generalization is more difficult than with other methods.

The perceptual spaces mapped to the M_{sm} - M_{tm} -physical matching parameters (later called physical-perceptual spaces) and their Procrustes values (low enough to show a reasonable fit) are presented in Fig. 8. Each group of four physical-perceptual spaces for the two featured interactions (tapping and dragging) shows a similar horizontal-vertical alignment of elements, regardless of the surface type. Almost every physical-perceptual space shows that the perceptual difference between 1D vibration and 3D vibration is greater than the differences between various 321-converted vibrations. This finding means that the participants were able to sense a discrepancy between the original 3D vibration and all 321-converted 1D vibrations regardless of which algorithm had been applied and which surface or interaction was featured.

VI. GENERAL DISCUSSION

A. Comparison With Prior Perceptual Results

In [16], Park and Kuchenbecker assessed the same five algorithms by measuring the subjective similarity between the

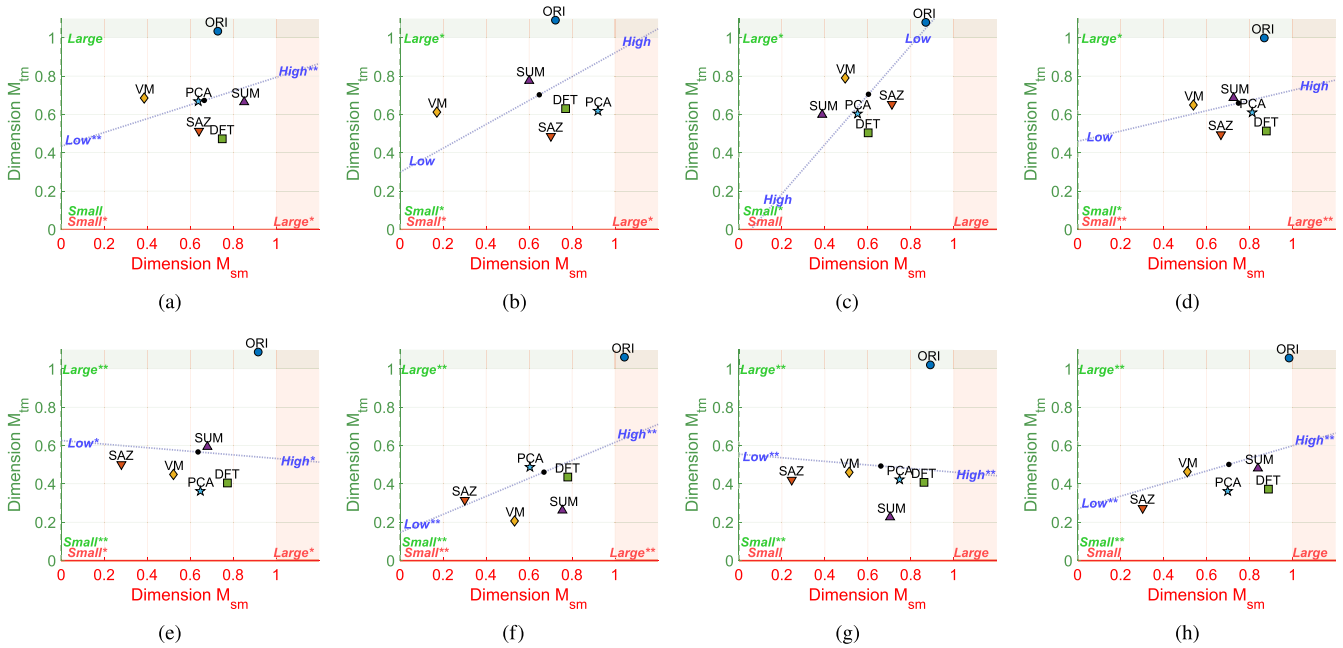


Fig. 8. M_{sm} - M_{tm} physical-perceptual spaces linearly transformed from the original perceptual spaces from Fig. 7. Blue dotted lines show the axis regressed by the signal energy E crossing the center of mass of the six stimuli (represented as a blue dot). The red and green colored areas show $M_{sm} > 1$ and $M_{tm} > 1$, respectively (regions that are presumably hard to reach for 321-converted vibrations). The corresponding Procrustes values are (from left) 0.607, 0.317, 0.435, and 0.273 in tapping, and (from left) 0.214, 0.153, 0.263, and 0.326 in dragging. (a) P_{THR} : Tapping on Sandpaper. (b) Tapping on Acrylic. (c) Tapping on Rough Paper. (d) Tapping on Rubber. (e) Dragging on Sandpaper. (f) Dragging on Acrylic. (g) Dragging on Rough Paper. (h) Dragging on Rubber.

original haptic interaction and its real-time-converted 1D vibration from participants' responses on a Likert scale from 1 to 7. Several interpretations can be inferred from this previous study: first, the similarity scores of dimensionally reduced vibrations were about 4.8, implying that participants could tell the difference between the original 3D interaction (which included low-frequency contact forces) and the reduced 1D vibrations. Second, the overall variance of the similarity scores was relatively smaller than that of the physical similarity scores, showing that physical dissimilarities and perceptual dissimilarities are not in simple proportionality. Third, although every algorithm was rated well, DFT seemed to provide a slightly better feel, i.e., a more similar perception to the original 3D vibration, than others, particularly during interactions with rough surfaces. Interestingly, SUM often achieved a perceptual similarity comparable to DFT. Since this method has lower computational cost and induces no time delay, it may be a good option for real-time implementation.

Our study presents the same implications as [16] with additional interpretation in more detail. In most of the spaces in Fig. 8, the 321-converted vibrations were distributed along the horizontal direction, seemingly parallel to the M_{sm} axis. Moreover, SAZ and DFT were placed farthest apart as the leftmost and rightmost algorithms in those spaces for dragging, following their low and high M_{sm} scores, respectively. The vibration in tapping mainly occurs along the z -axis, which is the principal axis of the movement. Conversely, the vibration in dragging always occurs both along the movement direction (in 2D, both x and y) and along the z -axis, resulting in a relatively even distribution to all three axes. In these regards, the low M_{sm} score of SAZ in dragging reflects that its spectrum includes only one

axis. Furthermore, DFT can be regarded as the best algorithm for dragging because this method has the highest M_{sm} score and results in the closest (rightmost) perception to the original 3D vibration. These results imply that matching the spectral information contributes to the perceptual similarity, especially for rich signals with broad spectral components. On the other hand, our physical-perceptual spaces for both tapping and dragging consistently show a vertical gap between the 1D vibrations and the 3D vibration, similar to Park and Kuchenbecker's finding that none of the tested algorithms earned an average similarity score higher than 5 out of 7 [16].

B. Factors Contributing to Perceptual Dissimilarity

What information allowed the participants to differentiate the reduced 1D vibrations from the original 3D vibrations? Since it is aligned with the M_{tm} axis in our analyses and the 1D vibrations all earned relatively low scores, it is tempting to attribute this difference to poor temporal matching. However, this gap could also originate from other differences that exist between these stimulus sets. It is not yet known what information is most important and how humans use it when making perceptual judgments. Therefore, we discuss some clues that participants might have used in addition to the established effect of spectral similarity.

First, we notice that the phase of the frequency components in each axial signal is not considered by current 321 algorithms. Accordingly, participants could potentially use distortion in the relative phase information when judging similarity; phase changes shift signals in time and therefore might affect temporal matching. Several decades ago, researchers suggested that two

groups of PCs might exist. Even though one PC may not easily distinguish the direction of vibrations [13], the groups could respond to different phases of stimuli [46]–[48] and could thus work together to obtain some phase-related information that is distorted by all of the tested 321 algorithms. Later psychophysical research, however, has reported that the effect of phase-related differences has no significant influence on perception. According to Bensaïa and Hollins, humans cannot perceptually discriminate superimposed sinusoidal vibrations at two frequencies felt via the PC channel (100 and 300 Hz) with different phase shifts [49]. As a result, researchers now generally agree that phase-related information is perceptually unimportant and that perception of such high-frequency vibrations is dominated by magnitude-related information [50]–[52]. We thus believe that phase information could have played a role in our study only for lower frequencies and not for the PC channel.

In this regard, our second hypothesis for the vertical gaps between ORI and the other stimuli in Fig. 8 arises from the different directionality of 1D and 3D vibrations. In our implementation, even 321 algorithms with high spectral matching scores create only single-axis vibrations along the z -axis, while the original recordings output vibrations across all 3D axes. In this regard, Brisben *et al.* found that the direction in which a grasped physical object vibrates does slightly but consistently affect the absolute threshold for vibration detection [53]. They conducted several psychophysical experiments using a custom cylindrical motor-operated probe with vibration stimuli at either 40 Hz or 300 Hz, finding that participants were slightly more sensitive to vibration movements parallel to the skin than to vibration movements perpendicular to the skin. Thus, we believe that the 3D nature of ORI stimuli might have caused them to feel fundamentally different from all of the 1D stimuli. Indeed, recent work by Serhat and Kuchenbecker used a detailed 3D finite element model of the human fingertip to show that changing the direction of an oscillatory force applied at the center of the fingerpad can greatly affect the resulting tissue deformations both locally and at a distance [54]. PCs or other mechanoreceptors distributed over the large contact area of the precision grip posture (such as Ruffini endings through the SA II channel for sensing skin stretch at 0.5–400 Hz or Meissner’s corpuscles through the FA I channel for fluttering sensations from the shearing force at 3–40 Hz; see [14] for details) may decode relevant directional information from the rendered vibration.

Our third hypothesis is that the dimensionally reduced vibration stimuli may have been perceived as different from the original signals due to technical artifacts in our approach to rendering 1D and/or 3D vibrations, such as rotational motion, vibrational crosstalk between the axes, or audible cues. We tried to limit these artifacts during the study design process, but they cannot be completely eliminated from consideration.

C. Future Work

Importantly, all three of our possible explanations may have contributed to our finding that none of the tested 321 algorithms could produce 1D vibrations that feel the same as the

original 3D recordings. In-depth studies with a wider variety of stimuli would be needed to test our hypotheses and determine their relative impacts on the perceptual performance of both existing and new 321 algorithms. For example, we suggest also presenting the original 3D vibration with axes swapped or with other reflections and rotations to determine whether such directional distortions are perceivable even when the spectral match and temporal match metrics are high.

One interesting finding is the relative positions of SUM in the physical-perceptual spaces. In our results, the vertical displacement of SUM from ORI varied depending on the type of interaction and the type of surface, more arbitrarily than other algorithms. We speculate that due to the summation of multiple axes in SUM, the phase-related information of the original signal can be randomly dispersed, e.g., some of the phase data may be canceled or even doubled. Yet, no existing algorithm takes phase-related information into account, and even DFT is biased to focus on the amplitude-related information. Finding a way to preserve the phase information of the original signal as well as matching the frequency spectrum would be a good direction for future research on 321 algorithms. We believe it should be possible to update DFT to improve its performance in this regard.

Except for SAZ and VM, the horizontal differences between DFT and other advanced 321 algorithms (SUM and PCA) were smaller than the vertical difference to ORI, suggesting that humans have a poorer ability to perceptually distinguish differences in the spectral magnitude than the differences induced by present approaches to dimensional reduction. As an alternative explanation to the three hypotheses already presented, the lower temporal match scores we observed may partially stem from the design of the metric itself; we assigned $M_{tm} = 1$ for ORI, but a 1D signal may not be able to reach this value for some original 3D vibrations. Other ways of quantifying temporal match should thus also be explored. Thinking more broadly, compact and lightweight hardware for generating high-quality 3D vibrations should be investigated as a way of circumventing the need for dimensional reduction in the first place.

VII. CONCLUSION

This study investigated the perceptual differences between three-to-one dimensional reduction algorithms. We developed a multi-dimensional vibration rendering system using a magnetic levitation haptic interface to output realistic 1D and 3D vibrations. We then collected a dataset of real vibrations from two representative actions (tapping and dragging) with four surfaces having different hardness and roughness. Using the developed system, we rendered the original 3D and the 1D vibrations converted by five different 321 algorithms and obtained their spectral and temporal matching scores. We proceeded to a user study with multiple pairwise comparisons of a subset of the same vibration stimuli. Then, we analyzed the results with MDS to obtain 2D perceptual spaces visualizing the perceptual differences between the original and 321-converted vibrations.

We found strong evidence of a sensory-physical relationship in multi-dimensional vibrations; the perceived dissimilarities are related to the physical dissimilarities based on features of the vibrations. Therefore, algorithms that seek to the spectral difference (like DFT) can still be considered the best 321 approach when perceptual similarity is prioritized, particularly for offline processing. Nevertheless, we also found that even the best available algorithms have a perception gap between the original vibration and the 321-converted vibrations. It remains to be seen whether any 1D vibration can completely imitate the feel of original, realistic 3D vibrations as would be desired for many applications in virtual reality and teleoperation.

ACKNOWLEDGMENT

The authors thank Ralph L. Hollis, Bernard Javot, Joey Burns, Christian Wallraven, and Lijuan Wang for their valuable technical support on this project. The authors also express gratitude to Gökhan Serhat and Eric Young for providing advice about building the system, Ifat Gertler, Yijie Gong, and Saekwang Nam for assisting with the creation of images and videos, and Ravali Gourshetti for proofreading.

REFERENCES

- [1] S. Choi and K. J. Kuchenbecker, "Vibrotactile display: Perception, technology, and applications," *Proc. IEEE*, vol. 101, no. 9, pp. 2093–2104, Sep. 2013.
- [2] V. Hayward, O. R. Astley, M. Cruz-Hernandez, D. Grant, and G. Robles-De-La-Torre, "Haptic interfaces and devices," *Sens. Rev.*, vol. 24, no. 1, pp. 16–29, 2004.
- [3] A. Ferreira and C. Mavroidis, "Virtual reality and haptics for nanorobotics," *IEEE Robot. Autom. Mag.*, vol. 13, no. 3, pp. 78–92, Sep. 2006.
- [4] C. Park, J. Park, S. Oh, and S. Choi, "Realistic haptic rendering of collision effects using multimodal vibrotactile and impact feedback," in *Proc. IEEE World Haptics Conf.*, 2019, pp. 449–454.
- [5] W. McMahan *et al.*, "Tool contact acceleration feedback for telerobotic surgery," *IEEE Trans. Haptics*, vol. 4, no. 3, pp. 210–220, Jul.–Sep. 2011.
- [6] R. P. Khurshid, N. T. Fitter, E. A. Fedalei, and K. J. Kuchenbecker, "Effects of grip-force, contact, and acceleration feedback on a teleoperated pick-and-place task," *IEEE Trans. Haptics*, vol. 10, no. 1, pp. 40–53, Jan.–Mar. 2017.
- [7] K. J. Kuchenbecker, J. M. Romano, and W. McMahan, "Haptography: Capturing and recreating the rich feel of real surfaces," in *Proc. Int. Symp. Robot. Res.*, 2011, pp. 245–260.
- [8] J. M. Romano and K. J. Kuchenbecker, "Creating realistic virtual textures from contact acceleration data," *IEEE Trans. Haptics*, vol. 5, no. 2, pp. 109–119, Apr.–Jun. 2012.
- [9] C. Salisbury, R. B. Gillespie, H. Tan, F. Barbagli, and J. K. Salisbury, "Effects of haptic device attributes on vibration detection thresholds," in *Proc. IEEE World Haptics Conf.*, 2009, pp. 115–120.
- [10] W. McMahan, J. M. Romano, A. M. A. Rahuman, and K. J. Kuchenbecker, "High frequency acceleration feedback significantly increases the realism of haptically rendered textured surfaces," in *Proc. IEEE Haptics Symp.*, 2010, pp. 141–148.
- [11] S. Choi and H. Z. Tan, "Toward realistic haptic rendering of surface textures," *IEEE Comput. Graph. Appl.*, vol. 24, no. 2, pp. 40–47, Mar./Apr. 2004.
- [12] N. Landin, J. M. Romano, W. McMahan, and K. J. Kuchenbecker, "Dimensional reduction of high-frequency accelerations for haptic rendering," in *Proc. Lecture Notes Comput. Sci.*, 2010, pp. 79–86.
- [13] G. Westling and R. S. Johansson, "Responses in glabrous skin mechanoreceptors during precision grip in humans," *Exp. Brain Res.*, vol. 66, no. 1, pp. 128–140, 1987.
- [14] L. A. Jones and S. J. Lederman, *Human Hand Function*. New York, NY, USA: Oxford Univ. Press, 2006.
- [15] H.-Y. Yao, V. Hayward, and R. E. Ellis, "A tactile enhancement instrument for minimally invasive surgery," *Comput. Aided Surg.*, vol. 10, no. 4, pp. 233–239, 2005.
- [16] G. Park and K. J. Kuchenbecker, "Objective and subjective assessment of algorithms for reducing three-axis vibrations to one-axis vibrations," in *Proc. IEEE World Haptics Conf.*, 2019, pp. 467–472.
- [17] I. Hwang, J. Seo, and S. Choi, "Perceptual space of superimposed dual-frequency vibrations in the hands," *PLoS ONE*, vol. 12, no. 1, 2017, Art. no. e0169570.
- [18] J. B. Brink, A. J. Petruska, D. E. Johnson, and J. J. Abbott, "Factors affecting the design of untethered magnetic haptic interfaces," in *Proc. IEEE Haptics Symp.*, 2014, pp. 107–114.
- [19] R. Zhang, A. J. Boyles, and J. J. Abbott, "Six principal modes of vibrotactile display via stylus," in *Proc. IEEE Haptics Symp.*, 2018, pp. 313–318.
- [20] R. Zhang and J. J. Abbott, "Characterizing detection thresholds for six orthogonal modes of vibrotactile display via stylus with precision grasp," *IEEE Trans. Haptics*, vol. 15, no. 1, pp. 164–177, Jan.–Mar. 2022.
- [21] R. Zhang, T. J. Schwehr, and J. J. Abbott, "Dimensional reduction for 6D vibrotactile display," *IEEE Trans. Haptics*, vol. 13, no. 1, pp. 102–108, Jan.–Mar. 2020.
- [22] R. Zhang, T. J. Schwehr, and J. J. Abbott, "Addendum to "dimensional reduction for 6D vibrotactile display,"" *IEEE Trans. Haptics*, vol. 14, no. 2, pp. 445–448, Apr.–Jun. 2021.
- [23] R. L. Hollis, "Magnetic levitation haptic interface system," U.S. Patent 8 497767 Jul. 2013.
- [24] A. M. Okamura, M. R. Cutkosky, and J. T. Dennerlein, "Reality-based models for vibration feedback in virtual environments," *IEEE/ASME Trans. Mechatronics*, vol. 6, no. 3, pp. 245–252, Sep. 2001.
- [25] W. McMahan and K. J. Kuchenbecker, "Dynamic modeling and control of voice-coil actuators for high-fidelity display of haptic vibrations," in *Proc. IEEE Haptics Symp.*, 2014, pp. 115–122.
- [26] Y. Jung and M. Enqvist, "Estimating models of inverse systems," in *Proc. IEEE Conf. Decis. Control*, 2013, pp. 7143–7148.
- [27] F. Boeren, T. Oomen, and M. Steinbuch, "Iterative motion feedforward tuning: A data-driven approach based on instrumental variable identification," *Control Eng. Pract.*, vol. 37, pp. 11–19, 2015.
- [28] L. Lennart, *System Identification: Theory for the User*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 1999.
- [29] M. A. Al-Alaoui, "Novel approach to analog-to-digital transforms," *IEEE Trans. Circuits Syst. I*, vol. 54, no. 2, pp. 338–350, Feb. 2007.
- [30] H. Lee, G. I. Tombak, G. Park, and K. J. Kuchenbecker, "Vibrations in tool-mediated interactions: Three-dimensional accelerations and corresponding videos during human-surface haptic interactions," 2021. [Online]. Available: <https://dx.doi.org/10.21227/8s0d-ca94>
- [31] R. Hassen and E. Steinbach, "Subjective evaluation of the spectral temporal SIMilarity (ST-SIM) measure for vibrotactile quality assessment," *IEEE Trans. Haptics*, vol. 13, no. 1, pp. 25–31, Jan.–Mar. 2020.
- [32] A. Israr, H. Z. Tan, and C. M. Reed, "Frequency and amplitude discrimination along the kinesthetic-cutaneous continuum in the presence of masking stimuli," *J. Acoust. Soc. Amer.*, vol. 120, no. 5, pp. 2789–2800, 2006.
- [33] G. A. Gescheider, *Psychophysics: The Fundamentals*, 3rd ed. Mahwah, NJ, USA: Lawrence Erlbaum Associate, 1997.
- [34] J. J. Zwislocki and D. Goodman, "Absolute scaling of sensory magnitudes: A validation," *Percept. Psychophys.*, vol. 28, no. 1, pp. 28–38, 1980.
- [35] G. Young and A. Householder, "A note on multidimensional psychophysical analysis," *Psychometrika*, vol. 6, no. 5, pp. 331–333, 1941.
- [36] Y. Vardar, C. Wallraven, and K. J. Kuchenbecker, "Fingertip interaction metrics correlate with visual and haptic perception of real surfaces," in *Proc. IEEE World Haptics Conf.*, 2019, pp. 395–400.
- [37] G. Park and S. Choi, "Perceptual space of amplitude-modulated vibrotactile stimuli," in *Proc. IEEE World Haptics Conf.*, 2011, pp. 59–64.
- [38] S. Okamoto, H. Nagano, and Y. Yamada, "Psychophysical dimensions of tactile perception of textures," *IEEE Trans. Haptics*, vol. 6, no. 1, pp. 81–93, Jan.–Mar. 2013.
- [39] D. Picard, C. Dacremont, D. Valentin, and A. Giboreau, "Perceptual dimensions of tactile textures," *Acta Psychol.*, vol. 114, no. 2, pp. 165–184, 2003.
- [40] I. Hwang and S. Choi, "Perceptual space and adjective rating of sinusoidal vibrations perceived via mobile device," in *Proc. IEEE Haptics Symp.*, 2010, pp. 1–8.
- [41] S. Mun, H. Lee, and S. Choi, "Perceptual space of regular homogeneous haptic textures rendered using electrovibration," in *Proc. IEEE World Haptics Conf.*, 2019, pp. 7–12.
- [42] R. Rosenkranz and M. E. Altinsoy, "Mapping the sensory-perceptual space of vibration for user-centered intuitive tactile design," *IEEE Trans. Haptics*, vol. 14, no. 1, pp. 95–108, Jan.–Mar. 2021.


- [43] N. Gaißert, C. Wallraven, and H. H. Bühlhoff, "Visual and haptic perceptual spaces show high similarity in humans," *J. Vis.*, vol. 10, no. 11, pp. 1–20, 2010.
- [44] E. Baumgartner, C. B. Wiebel, and K. R. Gegenfurtner, "Visual and haptic representations of material properties," *Multisens. Res.*, vol. 26, no. 5, pp. 429–455, 2013.
- [45] P. Dugard, J. Todman, and H. Staines, *Approaching Multivariate Analysis: A Practical Introduction*, 2nd ed. London, U.K.: Routledge, 2010.
- [46] S. J. Bolanowski and J. Zwislocki, "Intensity and frequency characteristics of Pacinian corpuscles. II. Receptor potentials," *J. Neurophysiol.*, vol. 51, no. 4, pp. 812–830, 1984.
- [47] S. J. Bolanowski, "Transduction mechanisms in Pacinian corpuscles," in *Mechanoreceptors: Development, Structure, Function*. Berlin, Germany: Springer, 1988, pp. 201–207.
- [48] J. Bell, S. Bolanowski, and M. H. Holmes, "The structure and function of Pacinian corpuscles: A review," *Prog. Neurobiol.*, vol. 42, no. 1, pp. 79–128, 1994.
- [49] S. J. Bensmaïa and M. Hollins, "Complex tactile waveform discrimination," *J. Acoust. Soc. Amer.*, vol. 108, no. 3, pp. 1236–1245, 2000.
- [50] R. F. Friesen, R. L. Klatzky, M. A. Peshkin, and J. E. Colgate, "Single pitch perception of multi-frequency textures," in *Proc. IEEE Haptics Symp.*, 2018, pp. 290–295.
- [51] R. Zhang and J. J. Abbott, "Vibrotactile display of patterned surface textures with kinesthetic haptic devices using balanced impulses," *IEEE Trans. Haptics*, vol. 14, no. 4, pp. 776–791, Oct.–Dec. 2021.
- [52] Y. Yoo, I. Hwang, and S. Choi, "Perceived intensity model of dual-frequency superimposed vibration: Pythagorean sum," *IEEE Trans. Haptics*, to be published, doi: [10.1109/TOH.2022.3144290](https://doi.org/10.1109/TOH.2022.3144290).
- [53] A. Brisben, S. Hsiao, and K. Johnson, "Detection of vibration transmitted through an object grasped in the hand," *J. Neurophysiol.*, vol. 81, no. 4, pp. 1548–1558, 1999.
- [54] G. Serhat and K. J. Kuchenbecker, "Free and forced vibration modes of the human fingertip," *Appl. Sci.*, vol. 11, no. 12, 2021, Art. no. 5709.



Hojin Lee (Member, IEEE) received the B.S. and Ph.D. degrees in computer science and engineering from the Pohang University of Science and Technology, Pohang, South Korea, in 2010 and 2019, respectively. He is currently a Postdoctoral Researcher with the Haptic Intelligence Department, the Max Planck Institute for Intelligent Systems, Stuttgart, Germany. His research interests include haptic interfaces, virtual reality, human-computer interaction, motor learning, cognitive psychology with human factors, and other human-machine technologies.



Güney Işık Tombak received the B.S. degrees in electrical and electronics engineering and physics from Boğaziçi University, Istanbul, Turkey, in 2020. He is currently working toward the master's degree in information technology and electrical engineering with ETH Zürich, Switzerland. He completed internships with the Physical Intelligence and Haptic Intelligence Departments of the Max Planck Institute for Intelligent Systems, Stuttgart, in 2018 and 2019, respectively. His primary research interests include signal analysis, computer vision, graph networks and statistical learning.



Gunhyuk Park (Member, IEEE) received the B.S. degrees in computer science and engineering and in electronic and electrical engineering in 2007, and the Ph.D. degree in computer science and engineering in 2017 from the Pohang University of Science and Technology, Pohang, South Korea. He is currently an Assistant Professor with the AI Graduate School, Gwangju Institute of Science and Technology, Gwangju, South Korea. He was a Postdoctoral Researcher with the Max Planck Institute for Intelligent Systems from 2017 to 2019. His research interests include haptic rendering and perception, mainly on tactile systems. He has worked on haptics for vibrotactile feedback perception, illusory vibrotactile feedback, automobiles, and mobile phones.



Katherine J. Kuchenbecker (Fellow, IEEE) received the B.S., M.S., and Ph.D. degrees in mechanical engineering from Stanford University, Stanford, CA, USA, in 2000, 2002, and 2006, respectively. After Postdoctoral Research with the Johns Hopkins University, Baltimore, MD, USA, she was a Faculty Member with the School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, PA, USA, for nine and half years. She moved to the Max Planck Society in 2017 where she directs the Haptic Intelligence Department with the Max Planck Institute for Intelligent Systems, Stuttgart, Germany. Her research interests include haptic interfaces, haptic sensing systems, and human-robot interaction.