Short Paper

A Low-Parameter Rendering Algorithm for Fine Textures

David Arthur Burns, Roberta L. Klatzky, *Fellow, IEEE*, Michael A. Peshkin, *Senior Member, IEEE*, and J. Edward Colgate, *Fellow, IEEE*

***Abstract*—This paper introduces a novel rendering algorithm for virtual textures, specifically those with characteristic length scales below 1 mm. By leveraging the relatively lossy mode of human tactile perception at this length scale, a virtual texture with wide-band spectral characteristics can be reduced to a spatial sequence of single-frequency *texels*, where each frequency is pulled stochastically from a distribution. A psychophysical study was conducted to demonstrate that, below a limiting physical texel length, virtual textures defined by identical frequency distributions are perceptually indiscriminable. Additionally, an exploratory study mapped the distribution parameters of the texel-based rendering to spectral characteristics of perceptually similar multi-frequency virtual textures.**

Index Terms—Texture rendering, surface haptics, tactile data compression, haptic codec.

I. Introduction

Since David Katz's proposal of the *duplex theory of texture perception* [8] in 1925, substantial research has provided evidence that when textures are fabricated with features of sufficiently small length scale, spatial feature arrangement can be safely discarded without effect to the user [7]. The discarding of unneeded textural information is of great interest to researchers developing Haptic Codecs for the Tactile Internet: an envisioned communications network with latency low enough to enable perceptually "real-time" interactions (in practice, less than 1 ms latency) [2]. Such communication speed will require information to be in the most data-efficient format possible. For tactile information, this necessitates an algorithmic framework that selectively captures only those elements of tactile stimulation perceived during tactile exploration.

With the goal of compressing tactile signals, a study by Okamoto *et al.* leveraged human sensory detection thresholds and quantization of texture signal spectral components [14]. Hassen *et al.* have developed an algorithm that is capable of compressing a tactile signal using a linear prediction scheme [6]. Noll *et al.* have developed a compression strategy centered around discrete wavelet transforms [13].

The above compression schemes found success in reducing tactile signal data load, but each approached the problem of compression from a "top-down" perspective: taking a maximum-data signal and identifying perceptually redundant elements that can be safely removed. An alternate pathway is the reverse, a "bottom-up" strategy: building a tactile signal from basic elements, adding complexity until the signal is perceptually identical to some reference.

Manuscript received September 14, 2021; revised November 16, 2021; accepted December 18, 2021. Date of publication December 28, 2021; date of current version March 18, 2022. This work was supported in part by NSF under Grant IIS-1302422. This article was recommended for publication by Associate Editor Tommaso Lisini Baldi and Editor-in-Chief Domenico Prattichizzo upon evaluation of the reviewers' comments. *(Corresponding author: David Arthur Burns.)*

David Arthur Burns, Michael A. Peshkin, and J. Edward Colgate are with the Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208-0001 USA (e-mail: dburns@u.northwestern.edu; peshkin@northwestern.edu; colgate@northwestern.edu).

Roberta L. Klatzky is with the Department of Psychology, Carnegie Mellon, Pittsburgh, PA 15217 USA (e-mail: klatzky@cmu.edu).

Digital Object Identifier 10.1109/TOH.2021.3138839

II. Background

A. Tactile Frequency Sensitivity

A motivating factor for a "bottom-up" perspective lies in the relatively insensitive mode of human tactile perception. Martinez *et al.* demonstrated that single-frequency tactile signals, applied on physically separate points of the arm, cannot be reliably discriminated from a multi-frequency signal applied at a single point [10]. This result suggests that individual space-isolated tactile signals are not being perceived: rather, all signals within the receptive field are combined during perception. Friesen *et al.* further demonstrated that dual-frequency tactile signals on the fingertip are reliably matched with single-frequency tactile signals [4], indicating that a predictable single-frequency representation exists for multi-frequency tactile signals.

The authors of [4] further demonstrated that the parameters of tactile signal amplitude, central frequency, and spectral width represent potent parameters for rendering a wide array of fine textures [3]. This study found that an amplitude-weighted estimate of vibrational *pitch*, alongside the perceived *noisiness* of the signal, exist as salient points of tactile signal discrimination.

For the present study, we target technologies that stimulate the entire fingerpad simultaneously, such as variable friction displays and vibrotactile devices. This focus is appropriate due to the ubiquity of this type of device in both academia and industry, especially considering that the algorithm described here is intended for efficient texture rendering in devices with limited data capacity. We note, however, that recent studies have demonstrated that the spatiotemporal distribution of stimuli across the fingerpad is perceptually salient, even at fine texture frequencies [5]. While the algorithm described here does not make use of this sensitivity, the concepts that drive it have the potential for adaptation to a multi-point stimulus technology.

B. Textural Quanta

When considering compression of sensory stimuli, the visual *pixel* represents a hallmark. Essentially the smallest controllable element of a virtual image, the pixel's success lies in its alignment with the measured acuity of human vision: below a certain scale, the eye cannot discriminate the edges between pixels, and thus an image composed of a near-infinite gradient of colors can be compressed to a finite series of single-color blocks without perceptual effect.

Investigations into an analogous "tactile pixel" have been made in the past. Meyer *et al.* developed a framework built around such a *texel* [11], the term used in this case to represent a single finite-length

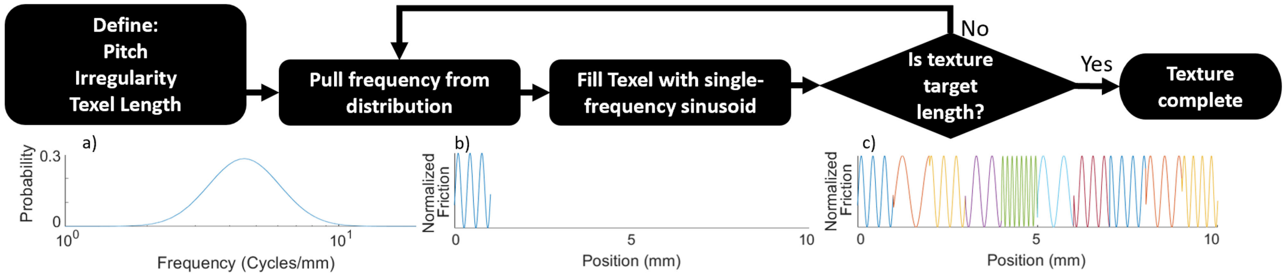

Fig. 1. Flowchart depicting rendering algorithm. a) Texel frequency distribution. b) First texel with stochastically-drawn frequency using a 1-mm texel length. c) Nine additional texels with stochastically-drawn frequencies.

texture signal composed of 101 spectral components. Briefly, the texture rendering algorithm generates the superposition of several overlapping 1-mm-long signals, each with a different specification of amplitudes for the frequency components. At the length-scale of features chosen (<0.25 mm), it was demonstrated that reliable spatial information was not perceived by the user [12], so the phase information of all spectral components could be safely discarded. Further, the friction force signals produced by a finger sliding over several real textures, when processed through the algorithm, exhibited spectral amplitude distributions that fit well within a Weibull model. This last point is particularly powerful: for the varied texture scans processed, a two-parameter statistical model could be used to reproduce the spectral signals in a given 1-mm-long patch.

III. Methods

A. *Single-Frequency Texel Rendering*

We present here a significant simplification of the *texel* described in [11]. Rather than allowing each spatial texel to embody several spectral components, we constrain each texel to a single sinusoidal frequency. Texture *irregularity* is achieved through variation of frequency texel-to-texel, governed by an underlying distribution where individual texel frequencies are pulled via a stochastic process. For the remainder of this paper, we will use "texel" to indicate a single-frequency texel.

To render a texture, three parameters are first chosen: some measure of the location and spread of the spatial frequency distribution, referred to here as *Pitch* and *Irregularity*, respectively, and the physical length of one texel. From the distribution, the frequency of sinusoidal oscillation in controlled friction is drawn and applied to the leftmost texel in the display. This process is repeated for a series of non-overlapping texels proceeding rightward until a sufficient texture width is achieved, as demonstrated in Fig. 1. We note that, while the abrupt transition in frequency at the border between texels was not found to be perceptible in this study, future work will include provisions towards eliminating potentially perceivable transitions in anticipation of extending this algorithm to more sensitive tactile display devices, such as by ensuring continuity of oscillation phase at texel boundaries.

This rendering algorithm has two important threshold constraints to consider. First, the stochastic nature of the texture-drawing scheme is appropriate only for fine textures, where spatial information is not perceived. Such textures are defined by a maximum feature length-scale. Numerous studies have been performed to identify this threshold, providing typical values between 0.25 mm [12] and 1.0 mm [15]. As such, distribution frequencies are constrained to exceed 1 cycle/mm. (Although including this minimum poses the risk of creating coarse features, that only adds rigor to the test of our algorithm's elimination of salient spatial clues.) Second, the length of a texel must be carefully chosen: it must not be wide enough to generate a perceptible spatial pattern to the user, but any reduction in texel size necessarily increases the rendering time required. Note that this threshold is related, but not identical, to the fine texture boundary described earlier. We expect that very large texels may allow for recognition of spatial details by the user, but this may not necessarily occur at the length-scale thresholds provided in the literature.

B. *Experimental Design*

1) Texel Length Test: In our first experiment, we sought to determine the effectiveness of the rendering algorithm at displaying a fine texture without introducing spatial cues that can be used for discrimination. We asked users to differentiate multiple texel-based textures rendered from the same underlying frequency distribution while scaling the texel length. Virtual texture rendering was achieved using a TPad ultrasonic friction modulation display, which uses finger location sensing (sampled via infrared CCD array at 125 Hz) to produce a one-dimensional friction modulation spatial map 100 mm in length (with spatial resolution of 0.0053 mm) [16].

The subject was introduced to the TPad device through practice software. A GUI on a laptop running MATLAB r2020b allowed subjects to switch between three sample textures in each test trial. When one was selected, it could be explored via fingertip on the TPad device. After exploring the three textures, the subject selected which of the three textures did not match the other two. For those subjects given the *visual practice* condition, visual representations of the three friction maps were also provided on the GUI.

Following five minutes with the practice software, the subject began the experiment. The test procedure matched that of the practice software, except that each trial had a maximum duration of 30 seconds, at which point subjects were required to submit their response before continuing to the next trial.

During each trial, two unique textures were used, with one of the three choices an exact copy of another. These two textures were rendered using identical frequency distributions, but the stochastic sampling resulted in textures that were never spatially identical. Each texture contained the maximum number of texels for the trial's texel length that fit in the 100-mm display length. Texel length was scaled trial-to-trial using a *one-up three-down* (1U3D) adaptive staircase technique, decreasing texel length with correct responses. Starting at 26 mm, texel length changed by 2.5 mm steps, decreasing to 1.25 mm steps after seven reversals. The test concluded after ten reversals.

The underlying frequency distribution used was a lognormal distribution, selected on the basis of pilot studies that found it to perform well with both lower and higher Pitches. The Pitch (in this case, the lognormal mean) value was set to 4 cycles/mm and the Irregularity (the lognormal standard deviation) was 1.2 or 4.3 cycles/mm, depending on whether the subject was given the *narrow* or *wide* Irregularity condition for their test.

TABLE I
REFERENCE TEXTURE PARAMETERS (SECOND EXPERIMENT)

Central Frequency (cycles/mm)	ζ
1	0.05
4	0.5
20	5

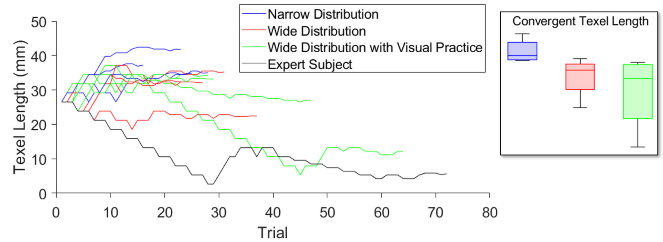

Fig. 2. Individual subject performance in experiment 1. Expert subject performance (using "Narrow Distribution") shown for reference. Box plots collect convergent texel lengths for each group.

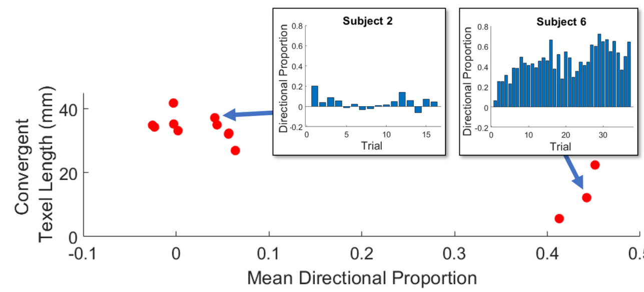

Fig. 3. Average values of directional proportion for each subject across all trials. Inset: sample plots of directional proportion for each trial for two sample subjects. (No data points are hidden by the inset).

2) Texel Parameter Mapping: Our second experiment was an exploratory study investigating perception of the underlying frequency distribution parameters of the rendering algorithm. We asked subjects to match a texel rendering to a multi-frequency texture signal by freely editing the Pitch and Irregularity parameters.

A GUI on a laptop running MATLAB r2020b allowed subjects to adjust two sliding controls, labeled "Pitch" and "Irregularity," between values of 0 and 100. When a new value was selected, the TPad display would present the newly-rendered texture, allowing the subject to explore freely with the finger. The active surface was divided in half, the left half always displaying the reference texture and the right half displaying the texture generated through the subject's Pitch and Irregularity choices. When the subject was satisfied with the match, the parameters could be submitted and the next trial would begin. This continued for nine trials, each with a unique reference texture.

Following these nine trials, the subject was asked to rate the similarity between one reference texture and one texel-based signal based on subject responses, rendered side-by-side on the TPad surface. The user was asked to rate similarity on a scale from 0 (completely different) to 10 (identical). To minimize test fatigue, only 25 of the possible 81 comparisons were investigated, selected on the basis of building a set of comparisons that spanned both distribution parameters at each level tested. Note that a subset of these comparisons involved a reference signal and its texel-based best match.

Reference textures were generated via the process established in [3]: white noise filtered using a spectral band-pass with a defined central frequency and Q-factor, where a smaller Q-factor results in a wider spectral width. For clarity, we define $\zeta = \frac{1}{2Q}$, so that increasing ζ corresponds with increasing signal irregularity. The nine reference textures were generated from the permutation of three values each of central frequency and ζ given in Table 1.

Editable textures were generated using the subject-defined distribution parameters. To allow for nonlinear scaling of distribution parameters, subject-defined Pitch P and lognormal mean μ were related by $\mu = (1.04^P)0.75$ (in units of cycles/mm), while subject-defined Irregularity I and lognormal standard deviation σ were related by $\sigma = \frac{3I}{10}$.

Based on the results of the first experiment, texel length was chosen for the second experiment to remain at a constant value of 1 mm. This value was selected because it was smaller by a sufficient safety margin (2x) than the best-performing subject could use to reliably discriminate textures drawn from the same frequency distribution parameters.

IV. RESULTS

A. Texel Length Test

There were 12 subjects for the first experiment, 5 female, all students of Northwestern University. Subjects were instructed to use the index finger of their dominant hand for texture exploration, but any finger on any hand to control the GUI. The protocol was approved by the Northwestern University Institutional Review Board, all subjects gave informed consent, and all subjects were paid for participation.

Subjects were divided evenly into three groups: the first used the *narrow* frequency distribution, the second used the *wide* frequency distribution, and the third used the *wide* frequency distribution and also received the *visual practice* software. For comparison, the performance of an expert subject (the primary author) is shown. Individual subject staircases, as well as summarized box plots of convergent texel lengths, are shown in Fig. 2.

To investigate discrimination strategies, the swipe speeds of all subjects were recorded. Both average swipe speed as well as standard deviation of swipe speed for each subject were found to be poorly correlated with performance (as measured by convergent texel length), $R = 0.2419$ and $R = 0.3171$, respectively. A metric of swipe directionality was defined by the following: for each trial, a running tally was incremented by 1 for every sample where the finger was swiping right, and decremented by 1 for leftward swipe. The tally was divided by the total number of samples in that trial and termed *directional proportion*, effectively describing the proportion of a trial when a subject's swipes were biased in one direction. This metric was found to have a stronger correlation with performance, $R = -0.867$. Each subject's average directional proportion is plotted against convergent texel length in 3, with two sample subjects' directional proportion values displayed for each trial.

B. Texel Parameter Mapping

There were 5 subjects for the second experiment, 2 female.

In Fig. 4(a), editable Pitch values are plotted against matching reference central frequency values for all subjects, for each ζ value. The diagonal line shows equal values in both axes. In Fig. 4(b), editable Irregularity values are plotted against matching reference ζ values for all subjects, for each reference central frequency value. Note that editable Irregularity values have been normalized across each subject's total set of matching tasks.

Similarity rating values were investigated for trends. In Fig. 5(a), box plots depict the set of similarity ratings between textures with the

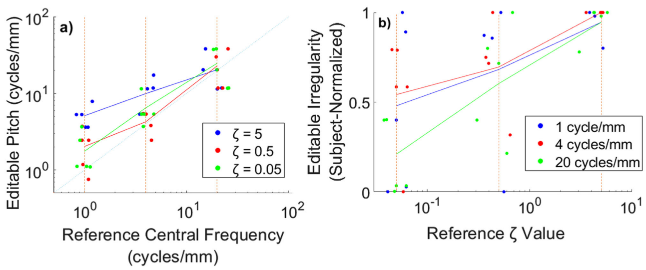

Fig. 4. a) Subject-matched values for Pitch. b) Subject-matched values for Irregularity. Solid lines indicate average values across all subjects. Points offset horizintally for clarity. Vertical dashed lines indicate actual reference values.

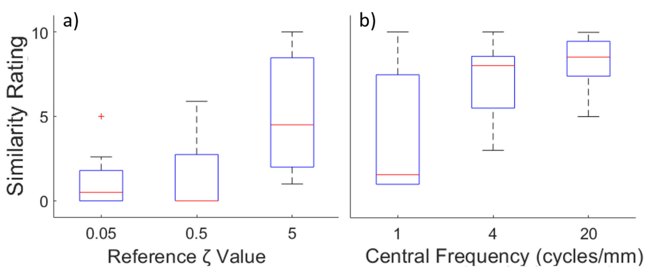

Fig. 5. a) Box plot depicting similarity ratings when comparing the lowest- and highest-central-frequency textures for a given value of ζ. b) Box plot depicting similarity ratings when comparing the lowest- and highest-ζ textures for a given value of central frequency.

highest and lowest reference central frequency, with ζ equal for both, for each value of ζ. Similarly, in Fig. 5(b), box plots depict the set of similarity ratings comparing textures with the highest and lowest reference ζ, with central frequency equal for both, for each value of central frequency. Compare this to the average similarity rating between each reference signal and its texel-based best match: 8.11 (standard deviation: 2.76).

V. Discussion

The performance of all subjects on the first experiment, including the expert subject, suggests the potential of this rendering algorithm to produce what we call textural *surrogates:* textures that are perceptually identical despite differing physical structures. Below a threshold texel length, texel-based fine textures drawn from the same underlying statistical parameters (Pitch and Irregularity) cannot be discriminated reliably. This threshold differed greatly between subjects (see Fig. 2), but no subject was capable of performing this discrimination near the 1-mm mark.

Discrimination performance, measured by convergent texel length, showed improvement with increasing Irregularity. This is a logical effect: as Irregularity increased, a wider range of texel frequencies could be sampled. This, in turn, increased the mean texel-to-texel frequency difference, allowing more dramatic changes in texel frequency to be detected as spatial features by the subject.

Adding a visual practice round before testing improved discrimination performance further. By receiving a visual model of the task, these subjects were more likely to continue successfully identifying spatial features at smaller texel lengths. This behavior seems to relate to the "Eureka" effect described by Ahissar *et al.*: a clearly perceived encounter that accelerates learning in the task at hand [1].

Among all groups, neither swipe speed nor variation in swipe speed correlated with performance. Directionality of swipe, however, had a powerful effect. Although all subjects were given freedom to explore textures in any manner, those that constrained themselves to a single direction of swipe (picking up their finger at the end of a swipe to return it to its original position) universally performed better than the other subjects (see Fig. 3). This strategy seems to have allowed for better discrimination of spatial features. It is likely that bi-directional swiping does not easily allow for a mental (spatial) map of the texture to be drawn as easily as for a single-direction swipe.

The behavior of the subjects in the second experiment provides confidence that the underlying statistical parameters could be used by texture designers in a predictable manner. While the average chosen values for Pitch and Irregularity follow the expected trend given the central frequency and ζ of the reference textures, some notable deviations were observed.

As shown in Fig. 4(a), at very high ζ, subjects reliably overestimated Pitch. This is likely an artifact of the nature of the lognormal distribution: as the distribution is incapable of crossing the origin of the frequency axis, a lowering of the mean necessarily compresses the left (low-frequency) tail of the distribution without affecting the right (high-frequency) tail, leaving more high-frequency content intact. This result suggests that a more optimal frequency distribution shape could be sought in order to better retain Pitch information without crossing the origin.

As shown in Fig. 4(b), while average chosen values of Irregularity increase monotonically as reference ζ increases, some subjects chose high values of Irregularity for the lowest ζ value when the chosen Pitch was low. In follow-up interviews, these subjects reported that the "bumpiness" or "waviness" of low-central-frequency reference textures could be achieved by either choosing a low Pitch value or a high Irregularity value. This behavior may be due to the additional low-frequency content generated when Irregularity is increased, but does not explain why the additional high-frequency content also generated did not counteract the effect. While this behavior was observed in only some of the subjects, future studies are required to better illuminate the interacting effects of Pitch and Irregularity as perceptual dimensions.

The similarity ratings demonstrate some of this interaction. Irregularity is a more potent factor for discrimination at low Pitch values than for high Pitch values. This is evidenced in Fig. 5(b) by the similarity ratings between minimum- and maximum-Irregularity versions of a texture of constant Pitch: Irregularity differences for low-Pitch comparisons result in significantly lower similarity ratings than those for high-Pitch comparisons. This is likely due to the fact that human tactile frequency discrimination decreases as frequency increases: a greater difference in frequency is required to discriminate higher frequencies than lower frequencies [9]. This suggests the need for an Irregularity parameter with gain that increases with increasing Pitch.

Additionally, at high Irregularity, the effectiveness of Pitch as a factor for discrimination decreases. This is evidenced in Fig. 5(a) by the high mean similarity rating between minimum- and maximum-Pitch versions of a texture at the highest Irregularity value, indicating Pitch discrimination was difficult. At lower Irregularity, similarity ratings are lower, as Pitch differences are more obvious to subjects. This can be explained by the fact that wider distributions tend to overlap more and share more frequency components than narrow distributions.

VI. Conclusion

We have introduced a low-parameter fine texture rendering algorithm and demonstrated its use in drawing virtual fine textures. Given the findings that no subject could reliably discriminate textures drawn from the same frequency distribution below a threshold texel length, and that subjects repeatably rate textures more similar when the frequency distribution parameters are closer to equal, we assert that textures with identical frequency distribution parameters represent fine-texture surrogates. This algorithm, then, can be a powerful tool in repeatably defining fine textures using an ultra-low amount of required data.

Future work will be conducted to fully understand the interacting effects of the Pitch and Irregularity tactile dimensions, as well as identifying a distribution shape that better represents tactile perception.

Acknowledgment

I am deeply grateful to the technical assistance of Arielle Menasce Traub, without whom completing this work would have been impossible.

References

[1] M. Ahissar and S. Hochstein, "Task difficulty and the specificity of perceptual learning," *Nature*, vol. 387, pp. 401–406, 1997.

[2] G. P. Fettweis, "The tactile internet: Applications and challenges," *IEEE Veh. Technol. Mag.*, vol. 9, no. 1, pp. 64–70, Mar. 2014.

[3] R. F. Friesen, R. L. Klatzky, M. A. Peshkin, and E. Colgate, "Building a navigable fine texture design space," *IEEE Trans. Haptics*, vol. 14, no. 4, pp. 897–906, Oct.–Dec. 2021.

[4] 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.

[5] R. V. Grigorii, "Capture, playback, and enhancement of tactile texture - Merging physics and perception," Ph.D. dissertation, Northwestern Univ., Evanston, Illinois, Jun. 2021.

[6] R. Hassen, B. Gülecyüz, and E. Steinbach, "PVC-SLP: Perceptual vibrotactile-signal compression based-on sparse linear prediction," in *IEEE Trans. Multimedia*, vol. 23, pp. 4455–4468, 2021, doi: 10.1109/TMM.2020.3042674.

[7] M. Hollins, S. Bensmaia, and E. Roy, "Vibrotaction and texture perception," *Behav. Brain Res.*, vol. 135, no. 1, pp. 51–56, 2002.

[8] D. Katz and L. Krueger, *The World of Touch*. Psychology Press, Hove, East Sussex, England, U.K., 1989.

[9] D. A. Mahns, N. M. Perkins, V. Sahai, L. Robinson, and M. J. Rowe, "Vibrotactile frequency discrimination in human hairy skin," *J. Neuriophysiol.*, vol. 95, no. 3, pp. 1442–1450, 2006.

[10] J. S. Martinez, H. Z. Tan, and R. W. Cholewiak, "Psychophysical validation of interleaving narrowband tactile stimuli to achieve broadband effects," in *Proc. IEEE World Haptics Conf.*, Jul. 2021, pp. 709–714.

[11] D. Meyer, M. Peshkin, and J. Colgate, "Tactile paintbrush: A procedural method for generating spatial haptic texture," in *Proc. IEEE Haptics Symp.*, 2016, pp. 259–264, Apr. 2016.

[12] D. J. Meyer, M. A. Peshkin, and J. E. Colgate, "Modeling and synthesis of tactile texture with spatial spectrograms for display on variable friction surfaces," in *Proc. IEEE World Haptics Conf.*, Jun. 2015, pp. 125–130.

[13] A. Noll, L. Nockenberg, B. Gülecyüz, and E. Steinbach, "VC-PWQ: Vibrotactile signal compression based on perceptual wavelet quantization," in *Proc. IEEE World Haptics Conf.*, 2021, pp. 427–432.

[14] S. Okamoto and Y. Yamada, "Lossy data compression of vibrotactile material-like textures," *IEEE Trans. Haptics*, vol. 6, no. 1, pp. 69–80, Jan.–Mar. 2013.

[15] A. Weber *et al.*, "Spatial and temporal codes mediate the tactile perception of natural textures," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 42, pp. 17107–17112, 2013.

[16] L. Winfield, J. Glassmire, J. E. Colgate, and M. Peshkin, "T-pad: Tactile pattern display through variable friction reduction," in *Proc. 2nd Joint EuroHaptics Conf. Symp. Haptic Interfaces Virtual Environ. Teleoperator Syst.*, 2007, pp. 421–426.