

NORTHWESTERN UNIVERSITY

A Multi-Scale, Low-Parameter Rendering Algorithm for Virtual Textures

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Mechanical Engineering

By

David Arthur Burns

EVANSTON, ILLINOIS

March 2023

© Copyright by David Arthur Burns 2023

All Rights Reserved

ABSTRACT

A Multi-Scale, Low-Parameter Rendering Algorithm for Virtual Textures

David Arthur Burns

As the technology enabling touch-sense rendering of virtual textures grows in efficacy and prevalence, so too grows the need for a standardized means of storing, transferring, and reconstituting textural signals. Furthermore, to achieve the ultra-low-latency requirements of the next generation of global communications networks, this digital texture representation must be data-efficient: retaining only the minimum set of textural features required to produce a rich virtual texture experience.

In this thesis, I introduce the *Texel rendering algorithm*. Designed to exploit the limitations of human vibrotactile discrimination, this algorithm reproduces a rich and varied set of fine textures using only three parameters as input. Through the use of a series of psychophysical studies, I demonstrate that these parameters are sufficient to define self-similar fine textures; that close perceptual matches can be made using this algorithm to a diverse group of fine texture families; that the parameters can be successfully navigated even by users unfamiliar with virtual texture rendering devices, especially when enabled

by a novel assistive algorithm; and that textures rendered with this algorithm are robust to additional compression via spectral quantization.

Additionally, I describe the capacity of this algorithm to produce coarse textural features. In a psychophysical study, I demonstrate that the appearance of simultaneous fine textural features does not affect the perception of coarse features, suggesting that the algorithm is successful in leveraging the *duplex theory of texture perception*: the simultaneous and separate modes by which coarse and fine textures are perceived. In this way, the Texel rendering algorithm can be used as an effective tool to design perceptually-rich virtual textures requiring a very small amount of input data.

Acknowledgements

While the number of people who lent me their support are too numerous to list here, the following individuals stand out in their generosity.

To Ed Colgate, I am truly indebted for the years of guidance, inspiration, and unflappable conviction that this whole thing was even possible.

I am greatly thankful for the mentoring by Bobby Klatzky and Michael Peshkin, on whom I relied heavily on matters psychophysical and electromechanical, respectively.

The other students in my lab group helped me more than I can express here, and I am especially thankful for the selfless assistance and caring friendship from Becca, Roman, and Sylvia.

To my parents, Marcia and Scott: thank you so much. This wouldn't have been possible without you both.

Most of all, to my wife, Brigid: thank you for the long years of tireless support, patient understanding, and willing participation in more subject tests than any other person (by a large margin). I owe everything here to you.

Dedication

To my sons, Thomas and Conor

Table of Contents

ABSTRACT	3
Acknowledgements	5
Dedication	6
Table of Contents	7
List of Figures	10
Chapter 1. Introduction	14
1.1. The Sense of Touch	14
1.2. Touch in Technology	15
1.3. Thesis Overview	16
Chapter 2. Background	18
2.1. Touch Sensing	18
2.2. Surface Haptics Technology	22
2.3. Compression of Tactile Signals	28
2.4. Statistical Signal Representations	32
Chapter 3. Spatial perception of textures depends on length-scale	38
3.1. Tactile Working Memory	38

	8
3.2. Experimental Design	40
3.3. Results	46
3.4. Discussion	48
Chapter 4. A low-parameter rendering algorithm for fine textures	51
4.1. Single-Frequency Texel Rendering	51
4.2. Experimental Design	55
4.3. Results	58
4.4. Discussion	62
Chapter 5. The Single-Pitch Texel: A flexible and practical texture rendering algorithm	65
5.1. Texel Rendering: Rigorous Testing	65
5.2. Optimization techniques	66
5.3. Methods	68
5.4. Results	77
5.5. Discussion	85
Chapter 6. Coarse Features in Texel Rendering	93
6.1. Coarse Feature Rendering via Amplitude Modulation	93
6.2. Experimental Design	94
6.3. Results	99
6.4. Discussion	99
Chapter 7. Conclusions	103
7.1. Texel Rendering Algorithm Summary	103

	9
7.2. Future Work	105
References	108
Appendix A. Parameter List	117

List of Figures

2.1	High-Performance TPaD surface haptics device	26
3.1	Sample texture maps and corresponding probability distributions for three values of average feature spacing: 2.5 mm (a), 4.6 mm (b), and 13.5 mm (c) and (d). 'A' represents normalized TPad friction reduction amplitude, and 'P' represents normalized probability amplitude. (All probability distribution widths are twice the given average feature spacing.) Note that, during a trial with the <i>different</i> condition, texture (a) and texture (b) would be compared for the narrow spacing group, and texture (c) and texture (d) would be compared for the wide spacing group.	42
3.2	Graphical User Interface (GUI) presented to subjects in the texture discrimination with distraction test.	44
3.3	Subject performance in Main Task (texture discrimination) and Distractor Task (matrix recall). Vertical line indicates performance at chance for Main Task.	46
3.4	Average texture discrimination performance for each of two texture scales, with and without distractor. Bar: average performance; whisker: standard error of the mean	48
4.1	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.	52
4.2	Individual subject performance in experiment 1. Expert subject performance (using "Narrow Distribution") shown for reference. Box plots collect convergent texel lengths for each group.	59
4.3	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.)	60
4.4	a) Subject-matched values for Pitch. b) Subject-matched values for Irregularity. Solid lines indicate average values across all subjects.	

		11
	Points offset horizontally for clarity. Vertical dashed lines indicate actual reference values.	61
4.5	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.	61
5.1	The nine reference textures used in this study represented by 6-mm-long samples of the space-domain TPad texture maps. TPad Amplitude refers to the friction reduction amplitude achieved at any point through ultrasonic vibration between the TPad surface and the test subject's finger, scaled by the maximum amplitude of the device. Inset: FFT magnitude traces for each of the textures, scaled for visibility. Red: commanded friction reduction spectral magnitude; purple: lateral fingertip velocity spectral magnitude acquired via Laser Doppler Vibrometer (LDV). (Note that dB units are calculated with a reference value of 1 cycle/mm.)	69
5.2	Graphical User Interface (GUI) for the Free Exploration test.	73
5.3	Graphical User Interface (GUI) for the Quantization test.	75
5.4	Guided Exploration test results for one subject. Lines connect subspace centroids throughout the optimization algorithm's progress (lines match the Pursuit Vector for each trial) and solid dots indicate termination points.	78
5.5	Graphical User Interface (GUI) for the Guided Exploration test matching trials.	79
5.6	Graphical User Interface (GUI) for the Guided Exploration test rating trials.	80
5.7	a) Central Pitch and b) Irregularity choices made during the Free Exploration Test, color-coded by texture family. Superimposed points: parameter choices of the trained rater.	81
5.8	Collected a) Central Pitch and b) Irregularity choices made during the Guided Exploration Test. For comparison, Free Exploration results are superimposed.	82
5.9	Central Pitch and Irregularity choices made during the a) Free Exploration Test and b) Guided Exploration Test for the Family Exemplars. Superimposed ellipses represent area of 95% confidence around the expected mean of a given texture.	83
5.10	Confidence ratings for all textures made during the Free Exploration Test versus distance between a test subject's chosen Texel parameters	

	and those made by the trained rater for the same target reference texture.	84
5.11	Similarity ratings between reference textures and Texel-rendering matches made during the a) Free Exploration Test and b) Guided Exploration Test. In c), a matrix portrays the Similarity ratings made during the Guided Exploration Test between each reference texture (columns) and a Texel-Rendering texture designed to match the named reference texture (rows).	84
5.12	Change in similarity ratings (between reference textures and Texel-rendering matches made during the Free Exploration Test) for each amount of quantization as compared to the similarity rating made in the unquantized state. Superimposed dotted line highlights zero change.	85
6.1	Texel algorithm used to render multiple textures using identical fine texture parameters (Central Pitch: 10 dB; Irregularity: -5 dB; Texel length: 0.25 mm). a) Rendering with amplitude held constant across texels (purely fine texture). b) Another rendering with amplitude held constant across texels. The stochastic nature of texel pitch selection results in another texture with similar statistical characteristics but different spatial characteristics. c) Rendering with amplitude modulated to include coarse sinusoidal features with period 5 mm (spatial frequency: -7 dB). d) Another rendering with identical amplitude modulation. Stochastic texel pitch selection does not affect coarse textural feature location.	95
6.2	Reference textures provided for the coarse texture discrimination test. Values given for "Friction" correspond to the friction range that the TPaD is capable of producing.	96
6.3	Examples of virtual textures demonstrating independent scaling of fine strength (horizontal axis) and coarse strength (vertical axis) for a given set of Texel rendering parameters (Central Pitch: 10 dB; Irregularity: -5 dB; Texel length: 0.25 mm; sinusoidal coarse features with spatial frequency: -7 dB). Values given for "Friction" are normalized.	97
6.4	Graphical User Interface (GUI) for the coarse feature discrimination test.	98
6.5	Subject performance in the coarse texture discrimination test for test segments with fine strength held constant at a.) 0%, b.) 50%, and c.) 100%,. Note that all lines have been offset by less than 1% for visibility.	100

6.6 Convergent coarse strength results for each test segment. a.) Individual subject performance given as points with lines connecting performance across segments for each subject. All points and lines have been offset by less than 1% for visibility. b.) 95% confidence intervals around mean convergent coarse strength values for each test segment.

CHAPTER 1

Introduction

1.1. The Sense of Touch

It is easy to block the visual channel by closing [your] eyes, yet slightly difficult to block the auditory channel by plugging your ears. However, it is almost impossible for a normal person to close his sense of touch. This shows that tactile sensing is a very fundamental perception modality for humans. . . [42]

Touch, as a means of exploring the world around us, is at once both a dividing and unifying sensory modality. Dividing, because, as it requires physical contact with the object being sensed, touch naturally gives an incomplete picture of the body's surroundings, leaving undetected everything outside of direct contact. On the other hand, the sense of touch unifies as it brings the *external* world as close to *internal* as possible, rendering sensations of immediate importance to the body's continued safety (most notably, via pain). In fact, touch has the demonstrated capacity to pull cognitive focus towards a previously ignored stimulus, effectively hotwiring the nervous system to momentarily place touch at the top of the sensory hierarchy [27]. By providing a person with sensory information immediately and without willful attention, the sense of touch stands as a powerful tool for richly experiencing the immediate environment.

It should come as no surprise, then, that the sense of touch exists as a viable feedback mechanism for effective activities in the modern world. From the kinesthetic haptic feedback used in teleoperation of robotic machinery to customizable virtual textures added to interactive computing devices, the quantity, scale, and complexity of haptic feedback has grown exponentially in recent years, and stands to only grow further [17].

Despite this, our understanding of tactile perception, and the means by which to generate rich synthetic tactile experiences, currently stands at a relatively immature state as compared to the wealthy knowledge bases of both visual and auditory perception. Closing this gap is absolutely necessary to enable optimal sensory feedback for tomorrow's technology.

1.2. Touch in Technology

The benefit of incorporating touch feedback into interactive devices is well-established. The addition of touch feedback mechanisms has been demonstrated to significantly reduce simple task completion time [65], expedite learning of physical skills [41], and even promote longevity of learning [76]. As the advance of technology allows for touch feedback in a growing number of applications, a greater share of the existing human-machine interface devices can benefit from this improvement.

Central to this movement are devices associated with touchscreens (e.g., smart phones and tablet computers). The ubiquity of devices in this category can hardly be understated. With the exception of some rudimentary vibrational pulses, however, sensory feedback from touchscreen-enabled devices has long been dominated by the visual and auditory modalities [7]. If a standardized means of adding touch feedback to touchscreen-enabled

devices can be developed, the sheer scale of the affected user base would be so massive as to be nothing short of revolutionary. To reach this point, an efficient, perceptually-rich, and compact means of defining tactile stimuli in a digital format is necessary.

1.3. Thesis Overview

This thesis begins with an overview of the mechanisms enabling the sense of touch and historical research elucidating the process of tactile texture perception. Multiple technologies capable of rendering virtual textures are reviewed, and the technology chosen for the present study is described. Next, motivation for the compression of digital tactile signals is given, followed by a review of state-of-the-art compression techniques. The concept of statistical texture representations is introduced with a review of research applying this concept to textures in both the auditory and tactile realms. Chapter 3 explores the representation of coarse and fine textures in working memory, including a psychophysical experiment demonstrating that the nature of the working memory segment used depends on the length-scale of the texture being explored. Chapter 4 introduces a novel rendering algorithm for virtual fine textures, utilizing a stochastic process to generate textures that, despite having differing spatial structures, remain perceptually identical, as confirmed via psychophysical study. Chapter 5 builds on the rendering algorithm, demonstrating that it can be used to produce varied families of fine textures. Using a novel guidance algorithm, it was shown that even users unfamiliar with virtual texture rendering technologies can repeatably use the algorithm to tune the rendering parameters and achieve varied output textures. Additionally, a psychophysical study was performed to demonstrate that perception of the rendered textures is robust to spectral quantization. Chapter 6 explores the

addition of coarse feature rendering capability to the algorithm, including a psychophysical study confirming that coarse feature discrimination is not affected by stochastic fine texture elements. Finally, the thesis is concluded with an overview of the studies' impact as well as a discussion of future work.

Several passages of this thesis are reproduced, with permission of the co-authors, from previously-published material in [11] (Copyright © 2021, IEEE) and [12] (Copyright © 2022, IEEE).

CHAPTER 2

Background

2.1. Touch Sensing

2.1.1. Somatosensory System

The human somatosensory system is a network of structures in the brain and throughout the body that allows detection of various physical stimuli. While also responsible for sensing pain, differences in temperature, and body position, the somatosensory system is the primary sensory nervous system responsible for the sense known as *touch*: intuitively, the means by which we explore the surrounding environment via physical contact.

The sense of touch can be further subdivided into *tactile* and *haptic* sensing [42]. While tactile sensing utilizes cutaneous (skin-bound) nerve endings for the detection and processing of contact parameters between the body and some target surface, haptic sensing involves information from both cutaneous and kinesthetic (body position) systems to model the two-way interaction between the body and some target object. As sensing in either modality is sometimes difficult to differentiate, for the present study we will consider touch sensing as a whole. When exploring a virtual texture by swiping the finger, both the cutaneous excitation of the fingertip by the texture as well as the awareness of the position and speed of the finger's movement are integrated automatically into a rich touch sensation.

The physical structures in the body responsible for producing neural signals in response to touch are known as mechanoreceptors. These structures are densely populated throughout human skin, with upwards of 17,000 such afferents in the hand alone [34]. In the context of exploring texture via touching with the fingertips, the relevant mechanoreceptors are those that populate glabrous (non-hairy) skin, which are summarized in Table 2.1.

Table 2.1. Cutaneous Mechanoreceptors of Glabrous Skin, adapted from [2] and [36]

Type	Adaptation Speed	Receptive Area	Optimal Stimulus	Physical Receptor	Peak Response Frequency Range
SAI	Slow	Small	Skin Indentation	Merkel cell	2 - 32 Hz
SAII	Slow	Large	Skin Stretch	Ruffini ending ¹	<8 Hz
RAI	Rapid	Small	Skin Movement	Meissner corpuscle	8 - 32 Hz
RAII	Rapid	Large	Skin Vibration	Pacinian corpuscle	128 - 400 Hz

These mechanoreceptors can be classified by their adaptation speed and receptive area. Slow-adapting (SA) receptors generate response signals while a relevant stimulus is continuously applied, whereas rapidly-adapting (RA) receptors only generate response signals at the initiation of a new (or changed) stimulus. Furthermore, Type I receptors respond to stimuli in a relatively smaller area of the nearby skin, whereas the responsive area of type II receptors is relatively larger. Due to these differences, as well as the differences in physical construction of the mechanoreceptors themselves, each of the four receptors in Table 2.1 is most effective at responding to a specific type of stimulus.

¹Although it has been theorized that the Ruffini ending is the physical receptor responsible for SAII responses in human glabrous skin, this has not been conclusively determined with direct evidence [2].

Although each of the mechanoreceptors mentioned demonstrate varied utility during diverse tactile exploration activities, of special note for this work are the means by which these afferents communicate different aspects of a virtual texture displayed on a physically flat surface when explored via lateral swiping of a single fingertip. The RAI units are active in communicating bulk finger vibration information at higher frequencies (>100 Hz), a salient stimulus associated with fine-grained textures. The SAI units, tuned to fire with skin indentation, may primarily provide information about whether the finger is in contact with the surface. The SAII units are active in detecting lateral tension in the skin, especially around the nail bed, which may provide information about the direction and magnitude of finger swipe direction [35]. Finally, the RAI units, with peak frequency response in the 8 - 32 Hz range, likely communicate information about lower-frequency vibrational stimulus during exploration.

With this limited use case in mind, we can exploit the strengths and limitations of the cutaneous mechanoreceptors to endow virtual textures with a perceptually rich touch sensation.

2.1.2. Duplex Theory of Texture Perception

The nature of tactile perception, in particular the distinction of sensing modes between fine and coarse textures, has a long history of research. In 1925, David Katz proposed what was later termed the *duplex theory of texture perception*: both spatial and non-spatial (vibratory) cues contribute independently to texture perception [38]. In later research, it was further confirmed that when textures are fabricated with features of sufficiently

small length scale, spatial feature arrangement can be safely discarded without effect to the user [31].

Specifically, the research of Hollins et al. [33] noted that spatial cues gave very little discriminatory ability to so-called *fine* textures (here, indicating particle sizes below 20 micrometers), evidenced by poor discrimination performance when testing textures with a static touch. This ability was greatly boosted, however, by the addition of finger movements, which produce non-spatial cues via finger vibration. This is in contrast to the discrimination performance of so-called *coarse* textures (here, indicating particle sizes above 100 micrometers), which was comparable during both static and dynamic touch. In another study, sensitivity to vibratory cues was decreased via vibrotactile adaptation, which resulted in a drop in discriminatory ability for fine textures but had no effect on this ability for coarse textures [32]. Below a certain threshold of feature spacing, it would seem, temporal cues dominate texture perception, whereas spatial cues are more important above this threshold.

Since then, several length-scale benchmarks have been set for the boundary between fine and coarse textures with respect to the nature of their perception. Estimates on the lower bound of spatial tactile acuity include 1.6 mm [56], 0.25 mm [53], 0.8 mm [15], and 1.0 mm [80].

This boundary marks the length-scale below which spatial information is no longer processed during perception. This “lossy” mode of perception is of particular use in the efficient design of virtual textures: fine textures can be represented in a non-spatial manner (e.g., as a sum of sinusoids with randomized phase) without producing any detectable

difference to the user. To put this knowledge to use, we require technology capable of generating controllable touch stimuli: a virtual texture rendering device.

2.2. Surface Haptics Technology

The long-held reliance on visual and auditory stimuli to mediate the human-machine interaction of modern technology (i.e., via screens and speakers) is rapidly being altered by the addition of another information transfer modality: the sense of touch. In recent years haptic feedback has been enabled for a wide variety of applications (such as machine teleoperation, advanced interactive computing, and virtual reality rendering), relying on myriad technologies to generate touch stimuli [17]. The sub-field of kinesthetic haptic feedback, focusing on technologies capable of applying forces about a user’s joints to simulate the force-displacement relationships of rendered objects and scenes, has historically been a major focus of haptics research and thus currently has a large body of established technologies and applications [23]. An alternate area of rich opportunity for the deployment of haptic feedback is that of touch-controlled devices, where fingertip interaction (rather than bulk body movement) is the primary source of input from the user. Owing both to the ubiquity of touchscreen-bound devices in today’s world as well as the advent of numerous touch-feedback and texture-rendering technologies in recent years [68], it is reasonable to expect the proliferation of these applications to only increase with time.

Of particular interest to designers striving to add touch feedback capability to fingertip-controlled devices is the field of surface haptics: in short, any technology with the core intent of providing touch-sense stimuli to the user via control of the forces between the user’s fingertip and the device’s active surface. Here, it is important to note that the

key differentiation between general haptic feedback and surface haptics is, logically, the constraint of *surface* interactions: in surface haptics applications, the human-machine interaction occurs primarily through exploration of a surface with a fingertip. The authors of [7] logically sort the varieties of surface haptics interactions into groups that share the same direction of force in which the fingertip is stimulated (either normal or tangential to the touch surface) and, further, into groups sharing similar means of generating this force.

Modern applications of normal-direction surface haptics are dominated by vibrotactile stimulation, where the out-of-plane vibration of the active surface couples directly to the vibration of the user’s fingertip and generates a sensation. While this type of stimulation is relatively simple to implement (as evidenced by the ubiquity of vibrotactile feedback protocols in today’s smartphones and handheld electronics), it remains a difficult task to actively control this type of sensation across an entire active surface [47]. While significant progress has been made towards solving this localization issue through exploitation of a touch surface’s harmonic behavior with an array of vibrotactile actuators as in [85], this technology has not yet been demonstrated as optimal for virtual texture rendering.

Tangential-force surface haptics, on the other hand, has dominated in the technologies used to render virtual textures. Studies such as that in [55] demonstrate that virtual textures rendered purely through modulation of a tangential force on the fingertip enable texture discrimination performance matching that when of physical textures of the same type. While other modalities exist, a significant body of literature in this area focuses on the concept of friction modulation: the direct control of sliding friction, in this case between a user’s fingertip and the touch surface. Likely due to the relative ease with

which friction modulation capability can be deployed within existing touch-controlled device architectures to provide tangential force feedback [7], this variety of surface haptics remains highly utilized in both academia and industry.

2.2.0.1. Electrostatic Friction Modulation. To achieve control of the frictional force between a sliding fingertip and a static touch surface, several methods have been developed. One of the most popular methods, electrostatic friction modulation, leverages an effect known as electroadhesion. First reported by Johnsen and Rahbek, this effect is the attractive force between two contacting surfaces with an electrical potential applied between them [37]. Attributable to Coulombic attraction between the surfaces across the small air gap that necessarily exists between them [69], this effect can be used to attract an electrically-grounded fingertip to a surface onto which a voltage is applied, effectively increasing the fingertip’s normal force (and thus magnitude of frictional force) [4]. Such a design was first used to generate virtual textures in [74], but the technology has since become widely popular in surface haptics applications [28].

2.2.0.2. Ultrasonic Friction Modulation. An alternate path towards friction modulation, and the texture rendering technology used for the present study, leverages an effect known as ultrasonic friction modulation. By vibrating a surface out of plane at ultrasonic frequencies with displacement on the order of a few micrometers, (a feat usually achieved by exploiting the resonant frequency of the vibrating surface,) the friction coefficient between a sliding fingertip and this surface is decreased significantly [79]. This effect has been established as dynamic levitation caused by both a squeeze film of air trapped beneath the fingertip surface [82] [9] and the intermittent contact between fingerpad asperities and the vibrating surface [78] [67]. While the magnitude of the squeeze

film thickness, and consequently that of the reduction in friction coefficient, has been demonstrated to correlate with combined mechanical properties specific to the fingertip-surface interaction (such as damping ratio) [21], calibration has been shown to be effective to correlate friction coefficient to fingerpad-specific dynamics based on limited required measurements [75].

2.2.1. TPaD: Tactile Pattern Display

To utilize ultrasonic friction reduction for the purpose of rendering virtual textures, the Tactile Pattern Display (TPaD) was developed. A glass surface vibrated by piezoelectric actuators at ultrasonic frequencies was found to reduce sliding friction with a human fingertip proportional to the amplitude of the vibration [84] and the mechanical properties and mode shapes of the vibrating glass were carefully tuned to maximize the friction-reduction magnitude and active area [86] [45]. To render virtual textures, this effect is coupled with finger position sensing, effectively allowing texture designers to specify a unique friction coefficient for every point on the active surface [57], with resolution bounded by that of the position-sensing system.

The TPaD equipment used for the present study is pictured in 2.1, developed in the author’s lab to maximize both response time and friction modulation magnitude [83]. The active borosilicate glass surface (104 mm x 25 mm x 3.2 mm) is glued to two piezoelectric actuators that vibrate with the surface’s resonant frequency, 34 kHz. At peak actuator voltage (50 V), the surface exhibits peak out-of-plane displacement of 4.8 μm . A CCD infrared light sensor array (102 mm length) is affixed above the active surface which detects light from an infrared LED affixed below the active surface. When in use, the shadow cast

by the user's finger, detected by the light sensor array, is used as one-dimensional position sensing. (The finger position in the direction corresponding to the 23 mm dimension is not sensed.) The resolution of position sensing is 0.0053 mm, allowing for 19,200 discrete points across the active surface. At each of these points, the friction coefficient between the sliding fingertip and the active surface is modulated via amplitude of surface vibration.



Figure 2.1. High-Performance TPaD surface haptics device

An on-board microcontroller (PIC32) can store a (1-D) virtual texture for playback in the form of a 19200-element array. Each element gives the relative numerical amplitude of the carrier signal with 65535 points of resolution. During playback, the amplitude of

vibration is set by the value assigned to the element corresponding to the current finger position (as determined by the CCD sensor) at a refresh rate of 8333 Hz.

2.2.1.1. Advantages. Ultrasonic friction reduction is, in many ways, optimal for virtual texture rendering in touch-enabled devices. Since the active surface can be entirely composed of a transparent material like glass, it can be positioned above an existing touchscreen without any loss of visual feedback [22]. Furthermore, recent advances in the production of fully transparent piezoelectric actuators will only further minimize the additional weight and thickness of additional hardware to add surface haptics capability to an existing touchscreen-bound device [24] [25].

Electrostatic friction modulation provides many of the same benefits as ultrasonic friction modulation, with the two demonstrating effectively identical perceptual effects to a user [77]. It is notable, however, that while electrostatic friction modulation increases the sliding friction force from the surface material’s unmodulated value, ultrasonic friction modulation decreases the friction force from the surface material’s unmodulated value. By choosing a surface material with relatively high friction coefficient when touched with a fingertip, a wide range of friction values is available to the virtual texture designer when using ultrasonic friction modulation [55]. In fact, a review of surface haptics technologies indicates that ultrasonic friction modulation applications provide a greater range of friction force values as compared to that of electrostatic friction modulation [7].

Additionally, the power consumption of ultrasonic friction reduction devices was shown to be amenable to optimization through use of a theoretical framework [81], suggesting that deployment of the technology in industry at mass scale is feasible.

2.2.1.2. Limitations. One measure by which electrostatic friction modulation is superior to ultrasonic friction modulation is bandwidth. Due to the dynamic nature of ultrasonic vibration and the requirement to actuate at the surface’s resonant frequency, ultrasonic friction modulation tends to have limited bandwidth, a shortfall not shared by electrostatic friction modulation [55]. For the present study, this bandwidth limitation was not detrimental to the virtual textures explored, but utilization of the technology for very high frequency friction modulation will require further development.

Additionally, the nature of friction modulation necessarily generates stimulation across the entire fingerpad simultaneously. While this has not impeded the development of a wide and varied group of virtual textures renderable using friction modulation, recent studies have demonstrated that spatially-distributed stimuli across the fingerpad are perceivable, even at frequencies traditionally associated with fine textures [26]. This suggests that, while friction modulation is an efficient means of providing tactile feedback to a user, there remain a subset of real textures that cannot be rendered using this methodology.

2.3. Compression of Tactile Signals

Virtual textures are, in a very literal sense, composed of data. Unlike their physical counterparts, virtual textures have no mechanical structure nor material properties. Rather, virtual textures are merely a set of instructions for a rendering device to follow during interaction with a user. Paramount to the creation of a rich virtual texture, then, is a deep understanding of what data is important to include in this interaction. Put another way: an ideal virtual texture enables a rendering device to transfer all relevant touch

information to the user without wasting processing and rendering time on any unneeded data (such as the textural information that is not perceived by the user).

2.3.1. The Tactile Internet

The discarding of unneeded textural information is of great interest to researchers developing Haptic Codecs for the Tactile Internet. The Tactile Internet, first proposed by Fettweis in 2014 [18], is an envisioned communications network with latency low enough to enable perceptually “real-time” interactions (in practice, less than 1 ms latency). The Tactile Internet is so named in reference to the power of latency this low: while consumers of visual and auditory media can accept a delay on the order of 10 and 100 ms, respectively, without perceptual effect, closed-loop force-feedback cues require a delay of 1 ms or less before latency can be ignored. A Tactile Internet, then, would provide the infrastructure allowing touch-based feedback over great distances. In their review of the concept, Simsek et al. term this paradigm shift “the democratization of skill”: the global sharing of expert ability in industries where, previously, live presence was required [70]. The advent of 5G communication networks, and development of 6G, both cast the Tactile Internet as a quickly-approaching reality [39].

To achieve this level of communication speed, it will be imperative to transfer information in the most data-efficient format possible, a specific goal of the Haptics Codec Task Group [1]. For tactile information, this necessitates two related goals. First, the neural substrate and coding that underlies human tactile sensation must be understood. Second, an algorithmic framework must be developed that selectively captures those elements of tactile stimulation that are perceived, while eliminating any elements that do not alter

the perceptual outcome. These goals, although not yet fully realized for touch-based sensation, are analogous in nature to the efforts that produced audio compression algorithms (such as MP3).

2.3.2. Auditory Compression Techniques

Several concepts utilized in audio compression algorithms are relevant to the challenge of tactile data compression [87]. In particular, the techniques involving absolute perceptual thresholds, spectral quantization, and simultaneous masking seem to have immediate utility towards tactile data compression. In each of these cases, the goal is to remove audio content (and thus reduce the amount of information in the final data structure representing the audio signal) without perceptual effect.

The absolute threshold of hearing represents the minimum energy required by a pure tone to be detected by a listener without background noise, collected for all audible frequencies. For the average healthy individual, this threshold is lowest for frequencies between 3000 and 4000 Hz, higher at frequencies above and below this range [71]. To utilize this biological limitation towards audio compression, an algorithm can remove any components of an audio signal that fall below this threshold. As this audio could not be perceived via human hearing, it will not have a perceptual effect when removed.

Spectral quantization relates to the discretization of frequency-domain information. In practice, this generally involves the modification of the Fourier Transforms of the segmented time-domain audio signal. When used for the purpose of audio compression, this can take two forms: decreasing the total number of frequency bands into which to store spectral magnitude information, or decreasing the total number of magnitude values

allowed for each frequency band [72]. Note that for digital audio signals, spectral data is natively discretized, so compression via spectral quantization in this case refers to a decrease in quantization level.

Simultaneous masking pertains to the phenomenon of one auditory signal causing another auditory signal, occurring at the same time, to be inaudible by a human listener. In general, if a sound has a sufficiently larger amount of energy than another sound occurring simultaneously, the latter will be inaudible, although the true masking behavior depends not only on the difference in energy but also the frequency bandwidths of both sounds [71]. In an audio compression algorithm, any sound that is being masked by another sound can be safely removed without perceptual effect.

2.3.3. Tactile Compression Techniques

In fact, numerous studies have been conducted to test the efficacy, for tactile compression, of techniques and algorithms used historically for auditory data. A study by Okamoto et al. demonstrated a tactile compression technique taking advantage of human sensory detection thresholds [61]. The same researchers went on to leverage quantization of the spectral components of a texture signal to compress the amount of texture data required, without significant perceptual effect [62]. Similarly, Chaudhari et al. used a methodology inspired by MP3 audio compression to demonstrate perceptual masking in tactile perception of multi-frequency tactile signals [14].

While existing audio-compression techniques are logical starting points for a tactile compression algorithm, they are certainly not the only avenues being investigated [73]. Hassen et al. have developed an algorithm that, in conjunction with a model of human

detection sensitivity, is capable of compressing a tactile signal using a linear prediction scheme [30]. Noll et al. have developed a compression scheme centered around discrete wavelet transforms [59] and significantly increased this algorithm’s performance through improvement of the underlying human tactile sensitivity model and the quantizing methodology [60]. Knowing that haptic feedback will be paired with audio/visual feedback in a large number of applications, Yuan et al. demonstrated that haptic signals can be simplified with little perceptual effect when paired with audio/visual signals that have strong semantic correlation [88].

An alternate path towards achieving low-data-load tactile signals is building a virtual texture from elementary tactile features, adding complexity until perceptual similarity to the target reference is achieved. This approach is apparent in [29], where the researchers deliberately generated new virtual textures by connecting psychophysical perceptual space and the space defined by measured tactile quantities (like normal force and swipe velocity). While this technique is effective at generating new realistic-feeling virtual textures, the need to use previously acquired data may limit the algorithm’s ability to render any arbitrary virtual texture conceived by a texture designer, especially those that are not represented by a physical texture.

2.4. Statistical Signal Representations

With the goal of compressing virtual texture signals, it is imperative to understand the data necessary to achieve a rich touch sensation, allowing a compression algorithm to remove all data extraneous to this purpose. The duplex theory of texture perception suggests that, while spatial data must be retained for accurate reproduction of coarse

textural features, fine textures may suffer no perceptual effect from representation in a purely statistical (space-averaged) format. Textures in multiple sensory modalities have been successfully represented in statistical formats, and these results are informative towards the development of a statistical representation of fine textures. Notably, research in the audio domain has particular utility to this task.

2.4.1. Statistical Audio Signals

A vital source of inspiration for the present study is the work of McDermott, Simoncelli, and colleagues on the topic of *sound textures*: audio signals with perceptual properties that remain constant over time. These researchers noticed that sound textures are common in nature, arising from the combination of numerous overlapping acoustic events. The sound of water running through river, a burning fire, and a large population of chirping insects are evocative examples of varied sound textures, and although these sounds are dissimilar to one another, they all share the property of forming auditory *scenes* that generate a constancy of sensation through time. Such constancy suggests that time-averaged statistical, rather than discrete event-based, representations dominate the perceptual form of these stimuli. To test this hypothesis, the researchers demonstrated that Gaussian noise modified by a relatively few statistical measures (such as frequency sub-band variance and kurtosis as well as sub-band envelope cross- and auto-correlation) were sufficient to reproduce realistic synthetic versions of several varied natural sound texture recordings [48]. By representing a rich sound texture through a few statistical parameters, significant data compression is achieved as compared to a time-domain sound texture recording. In addition to achieving perceptual similarity to natural sound texture

recordings as demonstrated via psycho-physical studies (the *realism* metric), a key test of the validity of these sound texture renderings is that any two synthetic textures rendered from the same statistical parameters, while representing different time-domain signals, must necessarily remain in the same perceptual class of sound texture (the *homogeneity* metric). In further development of this rendering algorithm, the researchers found that both the realism and homogeneity metrics were high only when the statistical parameters used for rendering were based on biologically-plausible mechanisms [50]. Furthermore, the researchers demonstrated that synthetic sound textures for certain classes of sounds; such as speech, music, and highly rhythmic sound patterns; did not achieve high realism. Such a result suggests that non-statistical mechanisms may handle perception of sounds that rely heavily on specific pitch and timing [50].

Further investigation of these synthetic sound textures revealed additional perceptual properties and suggested pertinent sensory mechanisms. Counter-intuitively, it was found that sound textures, while being easily discriminable from other sound textures even at short durations of playback, were more difficult to discriminate from different sections of the same sound texture as the duration of playback increased [49]. Such a result supports the hypothesis that the perception of sound textures of sufficient length is mediated by an automatic (and unavoidable) integration of summary statistics, replacing the temporal structure of perception for short or non-texture sound signals (like speech or music). Through rigorous psycho-physical testing utilizing their synthetic sound texture rendering algorithm, the researchers demonstrated that a finite integration window, several seconds in duration, accounted for the summary statistics of the human perception of a sound texture [51]. This window was found to increase in duration for sound textures with higher

variability, suggesting that the automatic process of perceptual integration is biased by the statistics of the sound texture being sensed. Another unique property of sound textures is the demonstrated capacity to produce an “illusory” perception of continuity over several seconds when interrupted by a masking sound (despite being physically absent). Such a long-duration illusion could not be achieved by non-texture sounds (such as speech), which produce a sensation of continuity only when interrupted up to a duration of a few hundred milliseconds [52]. This phenomenon may indicate that the perceived homogeneity of continuous sound textures, in addition to being statistical in nature, is a direct result of the automatic disregarding of some sensory inputs.

The elucidation of statistical audio representations made by McDermott, Simoncelli, et al. form a sturdy foundation on which to investigate analogous representations of touch-based textures, especially in the context of data compression. Specifically, the concepts of texture homogeneity and a perceptual nature tied to summary statistics show great promise in extension into the realm of touch sensation. Similarly to the process of building synthetic sound textures by adding statistical complexity to a simple base (Gaussian noise) signal, a wide class of virtual tactile textures could be realized through the process of starting with elementary tactile features and adding complexity to achieve a perceptually rich texture: the “bottom-up” strategy.

2.4.2. Statistical Touch Signals

A major motivating factor in deciding to face the problem of tactile data compression from such a “bottom-up” perspective lies in the relatively insensitive mode of human touch perception. This insensitivity has been previously exploited towards the goal of simplifying

haptic control. Martinez et al. demonstrated that single-frequency vibrotactile signals, applied on physically separate points of the arm, cannot be reliably discriminated from a multi-frequency signal applied at a single point [46]. This result suggests that individual space-isolated vibrotactile signals are not being perceived: rather, a combination of all signals within the receptive field are combined during perception. Friesen et al. further demonstrated that dual-frequency friction modulation signals on the fingertip are reliably matched with single-frequency tactile signals [20]; more specifically, the single-frequency matches correlated well with the amplitude ratio of the dual-frequency signals, indicating that a predictable single-frequency perceptual representation exists for multi-frequency signals. Such a result points towards a simplified cognitive representation of tactile signals, where precise vibration frequencies are discarded in favor of a weighted combination of these individual components.

The authors of [20] further demonstrated that the parameters of frequency modulation signal amplitude, central frequency, and spectral width represent potent parameters for rendering a wide array of *fine* textures (this term being used here to classify textures where feature spacing does not exceed 1 mm) [19]. This study found that a frequency/amplitude combination and an irregularity-dependent term represented the primary and secondary perceptual dimensions, respectively, demonstrating that an amplitude-weighted estimate of vibrational *pitch*, alongside an awareness of the *noisiness* of the signal, exist as salient points of signal discrimination during touch perception, especially for fine textures.

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 our understanding 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 is 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* [54], the term used in this case to represent a single finite-length textural unit composed of 101 spectral components. Briefly, the novel texture rendering algorithm generates the superposition of several overlapping 1 mm-long friction modulation signals, each with a different specification of amplitudes for the 101 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 [53], so the phase information of all spectral components could be safely discarded. Further, the lateral (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, thereby significantly decreasing the amount of data required to store a representation of this texture. While this greatly reduces the data necessary to render a virtual fine texture, a “bottom up” strategy might reveal further gains by leveraging lossy human tactile perception.

CHAPTER 3

Spatial perception of textures depends on length-scale

The content of this chapter was previously published in [11] (Copyright © 2021, IEEE). It is reproduced here with the permission of the co-authors.

3.1. Tactile Working Memory

The ability to compare two textures explored by touch necessitates the existence of some texture representation in human memory. The Working Memory model, first proposed by Baddeley and Hitch in 1974 [5], remains a leading explanatory schematic of the link between perception and memory. Succinctly, the model describes an internal *central executive* that redirects sensory input to one or more independent systems for temporary storage [66].

At the model's conception, these systems included the *phonological loop*, governing information stored via subvocal rehearsal (such as speech or rhythmic patterns); and the *visuospatial sketchpad*, capable of storing information in a spatial format.¹ Importantly, these systems are susceptible to independent suppression. Further research has demonstrated methods to interfere with these systems, such as articulatory suppression preferentially affecting the phonological loop [6] and spatial imagery affecting the visuospatial sketchpad. Articulatory suppression typically requires the test subject to verbally repeat

¹We note that while the term *visuospatial sketchpad* suggests exclusively visual inputs, this component could constitute a more general form of spatial working memory, compatible with non-visual sensory modalities.

a sequence of meaningless syllables, competing with the inner phonological loop used to retain verbal or rhythmic information in working memory. Spatial imagery distracting is exemplified by the Brooks Matrix Task, which requires memorizing a spatial array of numbers and later recalling the numbers stored at specific locations [10].

An extensive body of research has investigated visual and auditory stimuli within the framework of working memory, but the analogous pathway of touch stimuli remains rather unknown. It is logical to surmise that, given the demonstrated duplex theory of texture perception, there exist similarly dual modes of representing textures in memory. Due to the demonstrated importance of spatial information in textures with wide feature spacing, in contrast to the diminished importance of spatial information in textures with narrow feature spacing, we argue that the visuospatial sketchpad is a likely candidate for a working memory system that is effective for some, but not all, textural stimuli.

While earlier studies have demonstrated aspects of tactile perception that indicate the dominance of spatial and non-spatial cues for coarse and fine textures, respectively, this is the first study, to our knowledge, that directly relates length-scale-specific perception to spatial working memory. The demonstration of selective discrimination suppression of textures with a longer length scale using a spatial distractor represents both a novel mode for controlling texture perception and suggests a previously-untested pathway for texture representation within human memory.

3.2. Experimental Design

3.2.1. Overview

To study tactile perception at varying length scales with and without the effect of spatial working memory distraction, a test was designed with the following properties. 40 total test trials were presented to the subject, each requiring the subject to confirm whether two touch-explored virtual textures were identical or different in some way. The first 8 trials were considered training trials, and the subsequent 32 regular trials were recorded for results, divided evenly into four groups: 1.) large feature spacing without spatial working memory distraction, 2.) large feature spacing with distraction, 3.) small feature spacing without distraction, and 4.) small feature spacing with distraction. Exactly half of the trials presented different textures to the user, while the other half presented identical textures (as control).

3.2.2. Subjects

There were 24 test subjects, 16 males and 8 females, median age: 28 years. Subjects used a single finger of their choice to interact with the stimulus surface throughout the duration of testing, after pre-test training allowed them to try several fingers to find the strongest sensation. The protocol was approved by the Northwestern University Institutional Review Board, all subjects gave informed consent, and all subjects were paid for participation.

3.2.3. Apparatus

Virtual textures were displayed using a TPaD haptic display, as described previously. Friction maps were programmed and commanded over USB using a Microsoft Surface Pro 4 tablet running MATLAB R2019b. This tablet was also used to record test subject responses and advance test trials.

3.2.4. Stimuli

Pilot testing indicated that, for average inter-feature distances below 5 mm, subjects could not reliably distinguish textures of equal average inter-feature distance with different individual feature placements. As a spatial rearrangement of textural features is imperceptible at this length scale, it can be concluded that this specific family of textures at this characteristic length scale does not provide spatial information during perception. Above average inter-feature distances of 10 mm, however, textures with rearranged feature placements were reliably identifiable by subjects. By the same reasoning, this characteristic length scale must provide some spatial cues during perception. To target an average (non-distracted) texture discrimination performance of approximately 75% correct, subjects were asked to make same/different judgments between either: 1.) two textures with average length scales of 13.5 mm (wide spacing), or 2.) one texture with average length scale of 2.5 mm and one with an average length scale of 4.6 mm (narrow spacing). In both cases, the "same" condition represented the first texture being displayed twice.

All textures were composed of several 2.6 mm-long patches of minimum friction separated by varying-length patches of maximum friction. (See Figure 3.1.)

One pseudorandom texture of each inter-feature distance was generated by pulling several integers from a uniform distribution to assign inter-feature distances. The center of each distribution was equal to either 13.5 mm, 2.5 mm, or 4.6 mm. The width of each distribution was equal to twice its center, so that each distribution covered integers from 0 to twice its center. In this way, the three textures approximated scaled versions of each other, having equal width-to-center ratios for the uniform distributions from which the inter-feature distances were drawn. For each texture length scale, one virtual texture was produced by placing fixed-length minimum-friction features between patches of varying-length maximum-friction patches, the lengths of which were pulled from the distributions described above. (See Figure 3.1).

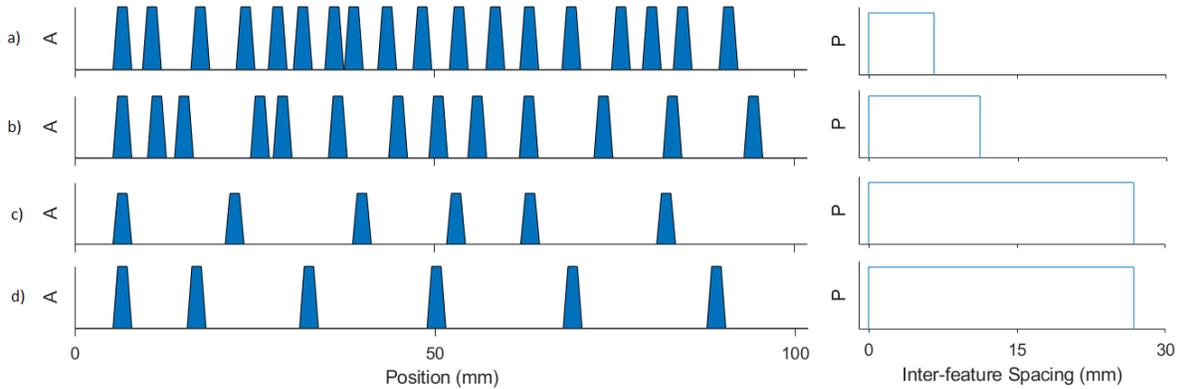


Figure 3.1. Sample texture maps and corresponding probability distributions for three values of average feature spacing: 2.5 mm (a), 4.6 mm (b), and 13.5 mm (c) and (d). 'A' represents normalized TPad friction reduction amplitude, and 'P' represents normalized probability amplitude. (All probability distribution widths are twice the given average feature spacing.) Note that, during a trial with the *different* condition, texture (a) and texture (b) would be compared for the narrow spacing group, and texture (c) and texture (d) would be compared for the wide spacing group.

The three textures as produced by the process above were given random rearrangements to produce additional pseudorandom textures. The order of inter-feature distances was shuffled, producing a different texture with inter-feature distance mean and standard deviation identical to the original texture. In this way, 64 unique textures were created.

3.2.5. Main Task and Distraction

The main task of the experiment was to compare two virtual textures presented at different times and decide whether they were identical matches with one another or different in the arrangement of fixed-length features.

Spatial working memory suppression was achieved using a technique derived from the Brooks Matrix Task. A 4-by-4 matrix of cells was presented to the user with five arbitrary cells containing five pseudorandom integers between and including 1 and 9. Two integers of the same value were never displayed in the same matrix. To apply a similar visual load during the non-distracted (control) condition, this matrix was completely filled with the digit 1, and test subjects were informed that a fully-filled matrix would always only contain the digit 1. In this way, a similar amount of visual information was displayed as compared to the distracted (test) condition without requiring any significant load on spatial working memory.

Due to the competing nature of the texture discrimination and spatial working memory distraction tasks, it was possible for subjects to ignore the main task in favor of performing well on the distraction task. As the purpose of the study was to investigate the effects of distraction on main task performance, subject data were considered outliers if the number

of main task errors exceeded two standard deviations from the population mean without any increase in distraction task errors.

Test subject input was recorded using the tablet PC containing the TPaD texture data. A simple Graphical User Interface (GUI), pictured in Figure 3.2, was designed that allowed test subjects to view the current phase of the test trial and the trial number, to view the distractor matrix as well as an indication of which number to recall, and to input decisions in both the texture discrimination task and matrix distraction task. Test subjects used the touchscreen interface to input decisions. Test subjects were required to use one finger for all TPaD touch interactions but had the freedom to use any finger on any hand for touch input into the tablet PC.

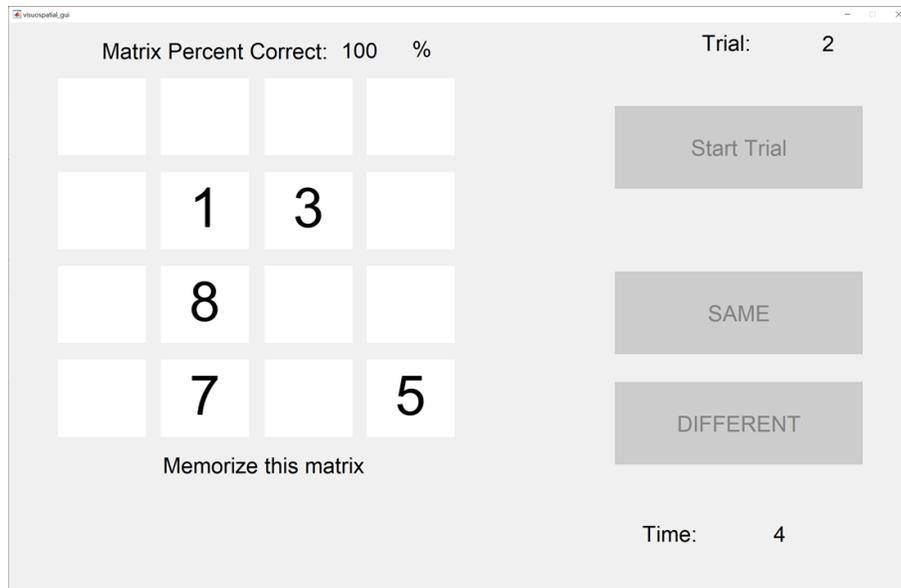


Figure 3.2. Graphical User Interface (GUI) presented to subjects in the texture discrimination with distraction test.

3.2.6. Experiment

Test subjects were introduced to the TPaD device and allowed to freely explore four sample textures for approximately five minutes. Sample textures consisted of friction maps not used in the experiment and were designed for the purpose of confirming that test subjects were able to detect friction changes when swiping a finger across the TPaD active surface. Following this, subjects were introduced to the GUI and given instructions for the test, at which point they were allowed to begin. Including sample texture exploration, total testing time per subject was approximately 40 minutes.

For each of 40 trials the following procedure was followed. The test subject started the trial and the first texture was produced on the TPaD surface, allowing the subject to inspect it freely for 10 seconds. Following this, the TPaD surface was turned off and the distractor matrix was displayed for 10 seconds. Next, the distractor matrix was hidden, and the second texture was produced on the TPaD surface, allowing the subject to inspect it freely for 10 seconds. After this, the TPaD surface was turned off and the subject was prompted to answer the first question: were the two textures displayed exact copies of one another or different in the arrangement of fixed-length features? Upon answering, the test subject was prompted to answer the second question: what was the number contained in a single indicated cell of the matrix? (For non-distracted control trials, where all cells in the matrix contained the number 1, all cells were indicated during this decision prompt). This ended the trial.

3.3. Results

The average number of errors made by subjects in the main task (texture discrimination) was 10.63, standard deviation: 3.36 (as compared to 16 errors or 50% expected by chance). The average number of errors made by subjects in the distractor task (matrix recall) was 2.04, standard deviation: 1.83 (as compared to 14 errors or 89% expected by chance, given that there were 9 choices). A weak negative trend between main task errors and distractor task errors was observed ($R^2 = 0.0747$). This effect is consistent with dual-task interference; that is, an increase in attention to one task would reduce attention to the other.

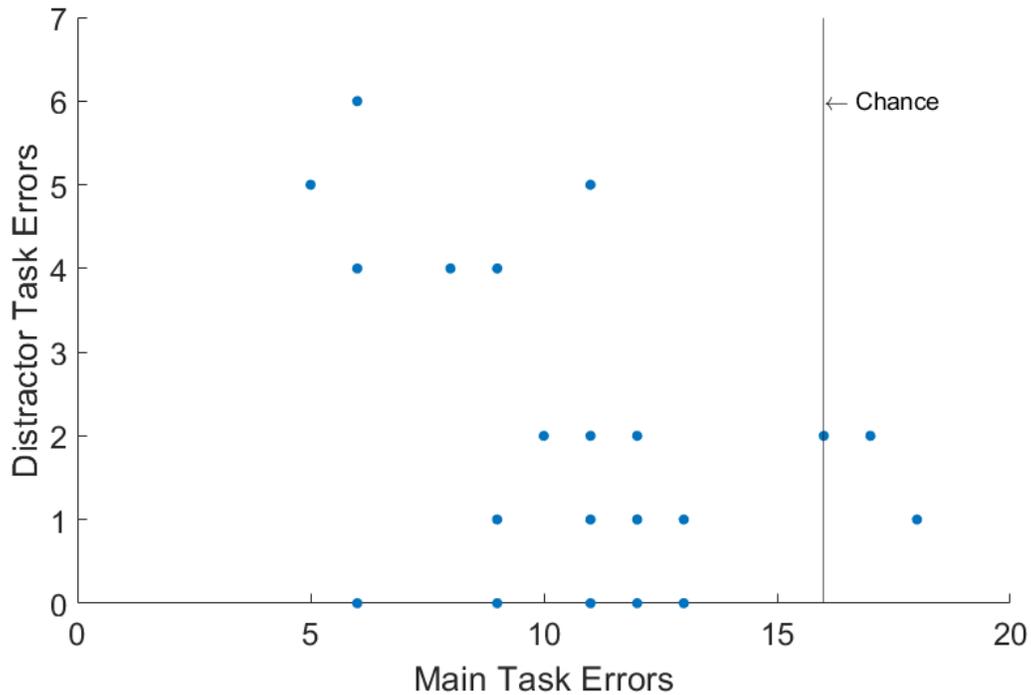


Figure 3.3. Subject performance in Main Task (texture discrimination) and Distractor Task (matrix recall). Vertical line indicates performance at chance for Main Task.

Table 3.1. Repeated-Measures ANOVA Results

	Sum Sq.	Error SS	Df	F	p-value
Length Scale	26.18	48.82	21	11.26	0.003
Distraction	11.64	30.36	21	8.05	0.010
Interaction	4.55	17.46	21	5.47	0.029

Two subjects exhibited especially low performance in the main task, with main task errors of 17 and 18. These subjects performed well in the distractor task, with distractor task errors of 2 and 1, respectively. This main task performance, below the level of chance, was near two standard deviations from the mean performance computed using all subjects. Given the pattern of the data, we conclude that these two subjects devoted nearly all concentration to the distractor task at the cost of the main task, resulting in chance main task performance and high distractor performance. As the purpose of this study is to investigate the effects of distraction on main task performance, we consider these two subjects outliers and remove them from the analysis below.² The texture discrimination performance among the remaining 22 subjects was sorted into four categories along two variables: *Wide* or *Narrow* feature spacings (corresponding to 13.5 mm and 2.5/4.6 mm textures, respectively) and *No Distractor* or *With Distractor* trials. The average performance in each of these categories is shown in Figure 3.4.

Univariate Type III Repeated-Measures ANOVA (assuming sphericity) was performed using the R programming language to identify the significance of the interaction between texture length scale (*Narrow* or *Wide*) and Distraction (*No Distractor* or *With Distractor*) on texture discrimination performance. The results are shown in Table 3.1.

²If the two outlier subjects are included, the ANOVA indicates a lack of power to detect the interaction ($p = 0.139$), but the t-test comparing distractor to control yields the same effects: nonsignificant for narrow textures ($p = 0.198$) and significant for wide textures ($p = 0.007$).

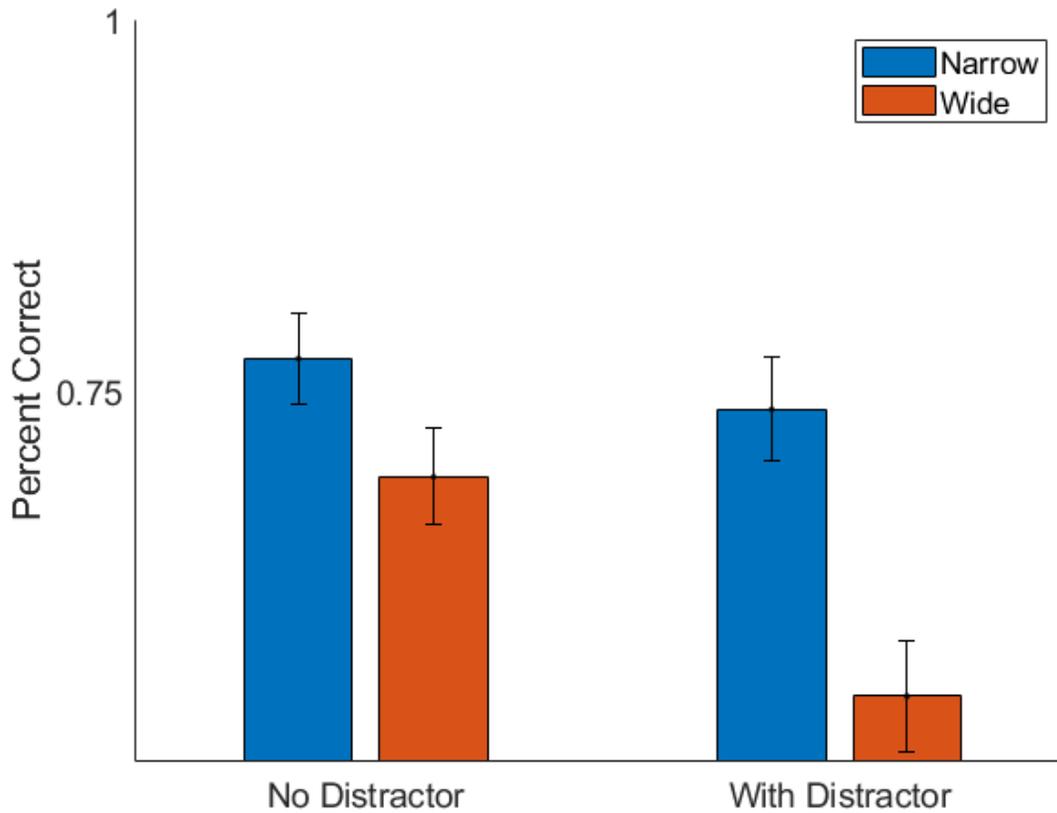


Figure 3.4. Average texture discrimination performance for each of two texture scales, with and without distractor. Bar: average performance; whisker: standard error of the mean

3.4. Discussion

The results demonstrate that suppression of spatial working memory decreases performance in texture discrimination significantly for wider inter-feature spacing but not for narrower inter-feature spacing. The distractor produced a 14.8% drop in average performance for wide textures, as compared to a nonsignificant 3.3% drop for narrow textures.

These results are in line with the duplex theory of texture perception, confirming that widely-spaced features in a texture rely heavily on spatial cues during perception, while narrowly-spaced features do not. To our knowledge, however, the connection between this theory and the concept of dedicated working memory components has not yet been established. By demonstrating that spatial working memory is a key to storing only widely-spaced textures, we reach two novel conclusions. First, we suggest that the perception of widely-spaced textures, while touch-sensed in origin, is converted to a spatial format in its cognitive representation. Second, we note that the perception of narrowly-spaced textures, while affording discrimination performance well above chance, must rely on some other non-spatial memory representation, potentially used for retention of statistical summary aspects of touch-based perception. This would be consistent with the assumption that fine textures are coded by the neural system as a vibrotactile signal.

The existence of a length-scale-governed threshold between textures that make use of the visuospatial sketchpad and those that do not will give guidance to future designers of virtual textures. By enforcing feature density to remain above this threshold, a designer can be confident that spatial information is not being perceived by the user and can be safely discarded without risking any perceptual difference. This high-information-load spatial data (e.g., exact location of each textural feature) can be replaced with a statistical representation of the texture (e.g., given a mean and standard deviation of feature spacing within a given area, features placed stochastically) that can be stored using significantly less information.

It is not necessary to our conclusions that the specific model of working memory described by Baddely and Hitch be accepted. For example, a competing theory described by

Cowan [16] rejects the existence of several specific-purpose short-term memory subsystems in favor of a limited-capacity focus of attention. In this model, the Focus of Attention, capable of storing a small number (typically 3-5) of cognitive items, is directed voluntarily and involuntarily by a Central Executive. The phenomenon of short-term memory is apparent in the finite lifetime of Activated Memory, where, without the rehearsal provided by the Focus of Attention, storage lasts on the order of seconds. To accommodate results like these that show modality-specific interference, this model would have to be augmented with another mechanism, for example, modality-specific application of the Focus of Attention.

CHAPTER 4

A low-parameter rendering algorithm for fine textures

The content of this chapter was previously published in [12] (Copyright © 2022, IEEE). It is reproduced here with the permission of the co-authors.

4.1. Single-Frequency Texel Rendering

The demonstrated non-spatial nature of fine texture perception, coupled with previous success in rendering virtual fine textures in statistical [54] and low-parameter [19] formats necessitates an investigation of combining these strategies towards the end goal of minimizing the data required to represent a virtual fine texture.

This chapter introduces a significant simplification of the *texel* described in [54]. Rather than allowing each spatial texel to embody several spectral components, each texel is constrained to a single sinusoidal frequency. Texture *irregularity* is achieved through variation of this frequency texel-to-texel, governed by an underlying distribution where individual texel frequencies are pulled via a stochastic process. For the remainder of this thesis, the term “texel” will be used 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 Figure 4.1. We note that, while the abrupt transition in frequency at the border between texels was not found to be perceptible in this study, improvement to the algorithm in the next chapter includes a provision towards eliminating potentially perceivable texel transitions by ensuring continuity of oscillation phase at texel boundaries.

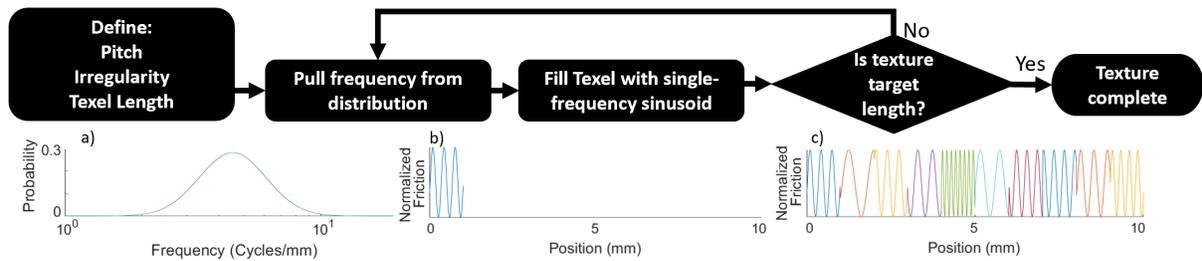


Figure 4.1. Flowchart depicting rendering algorithm. a) Texel frequency distribution. b) First texel with stochastically-drawn frequency using a 1-mm texel length. c) Nine additional texels with stochastically-drawn frequencies.

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 [53] and 1.0 mm [80]. 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.

This rendering scheme relies on the idea that individual tactile frequencies are temporally merged during perception, forming cognitive formats that combine the spectral characteristics of many signals over some time frame, analogous to the audio texture integration window observed in [51]. If such an assumption is taken as true, it is straightforward to assert that a fine texture composed of texels contains sufficient spectral similarity to an amplitude-weighted multi-frequency texture with the same spectral components.

To do so, we take a finite-length multi-frequency signal y_1 composed of N individual spectral components, each with frequency f_i and phase ϕ_i , to be defined by the equation given in equation 4.1, where L and A are the total signal length and amplitude, respectively.

$$(4.1) \quad y_1 = \begin{cases} 0 & \text{if } x < -\frac{L}{2} \\ \sum_{i=1}^N \frac{A}{N} \sin(f_i x + \phi_i) & \text{if } -\frac{L}{2} \leq x \leq \frac{L}{2} \\ 0 & \text{if } x > \frac{L}{2} \end{cases}$$

Similarly, we take a texel-based simplification of this texture y_2 to be defined by the equation given in equation 4.2, where $f(x)$ and $\phi(x)$, the texel frequency and phase, respectively, modulate instantaneously between the individual frequencies of f_i and phases of ϕ_i in equation 4.1 at the border of each texel of length $\frac{L}{N}$.

$$(4.2) \quad y_2 = \begin{cases} 0 & \text{if } x < -\frac{L}{2} \\ A \sin(f(x)x + \phi(x)) & \text{if } -\frac{L}{2} \leq x \leq \frac{L}{2} \\ 0 & \text{if } x > \frac{L}{2} \end{cases}$$

We take the Fourier Transform of both equations. The multi-frequency signal yields equation 4.3.

$$(4.3) \quad \mathcal{F}(y_1) = \frac{L}{2N} \sum_{i=1}^N \left(\text{sinc}\left(\frac{L}{2}(f_i - \omega)\right) + \text{sinc}\left(\frac{L}{2}(f_i + \omega)\right) \right)$$

We note that the texel-based signal is a train of finite sinusoids, each of length $\frac{L}{N}$. We denote δ_i the spatial offset from the origin to the beginning of the texel with frequency f_i , allowing us to write the Fourier Transform of equation 4.2 as in equation 4.4.

$$(4.4) \quad \mathcal{F}(y_2) = \frac{L}{2N} \sum_{i=1}^N \left(\text{sinc}\left(\frac{L}{2N}(f_i - \omega)\right) e^{i\omega\delta_i} + \text{sinc}\left(\frac{L}{2N}(f_i + \omega)\right) e^{i\omega\delta_i} \right)$$

By constraining our frequencies to remain above the limit of *fine textures* where spatial information is not detected by human tactile perception, we focus only on the magnitudes of the spectra, allowing us to ignore the phase-only contribution of the $e^{i\omega\delta_i}$ terms. With this simplification, we note that the spectrum of the texel signal differs from that of the multi-frequency signal in the frequency of oscillation of the *sinc* terms: notably, the spectral peaks lie at the same frequencies with the same amplitudes. The difference,

which constitutes different amplitudes of the frequency components surrounding the spectral peaks, is likely lost in perception, considering the lossy tactile frequency perception demonstrated in [46] and [20].

4.2. Experimental Design

4.2.1. 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 subjects to differentiate multiple texel-based textures rendered from the same underlying frequency distribution while scaling the texel length. Virtual texture rendering was achieved using the TPaD ultrasonic friction modulation device.

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 approximately 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 the TPaD was turned off and 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.

4.2.2. 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 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

Table 4.1. Reference Texture Parameters (Second Experiment)

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

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 [19]: 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.

4.3. Results

4.3.1. 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 Figure 4.2.

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

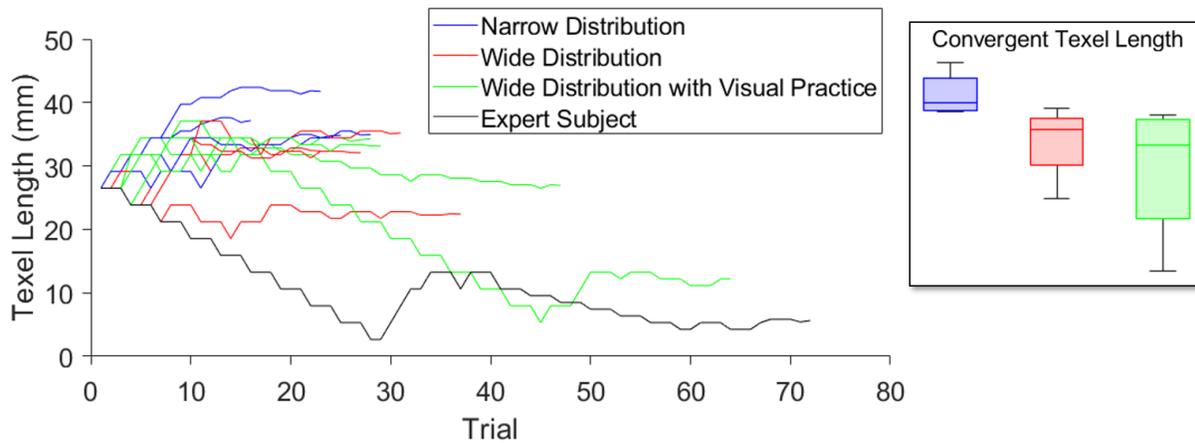


Figure 4.2. Individual subject performance in experiment 1. Expert subject performance (using "Narrow Distribution") shown for reference. Box plots collect convergent texel lengths for each group.

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 4.3, with two sample subjects' directional proportion values displayed for each trial.

4.3.2. Texel Parameter Mapping

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

In Figure 4.4 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

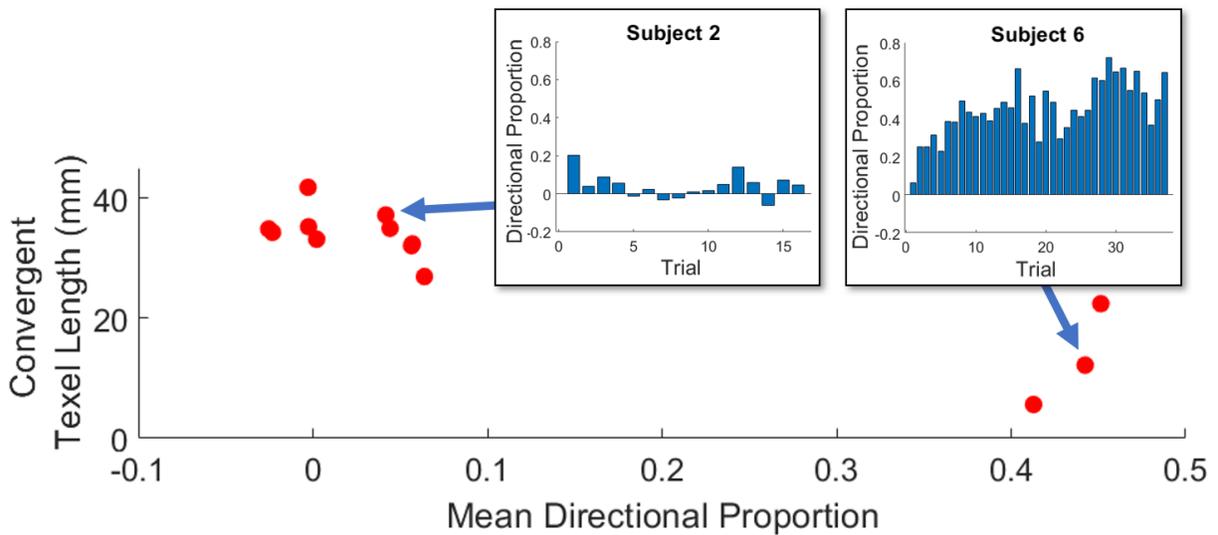


Figure 4.3. Average values of directional proportion for each subject across all trials. Inset: sample plots of directional proportion for each trial for two sample subjects. (No data points are hidden by the inset.)

in both axes. In Figure 4.4 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 Figure 4.5 a), box plots depict the set of similarity ratings between textures with the highest and lowest reference central frequency, with ζ equal for both, for each value of ζ . Similarly, in Figure 4.5 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).

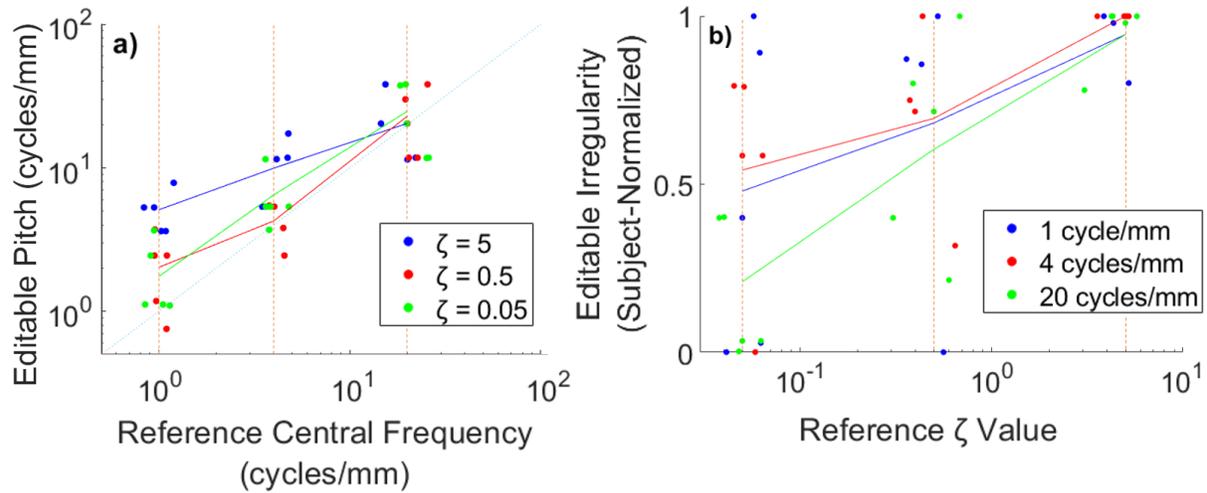


Figure 4.4. a) Subject-matched values for Pitch. b) Subject-matched values for Irregularity. Solid lines indicate average values across all subjects. Points offset horizontally for clarity. Vertical dashed lines indicate actual reference values.

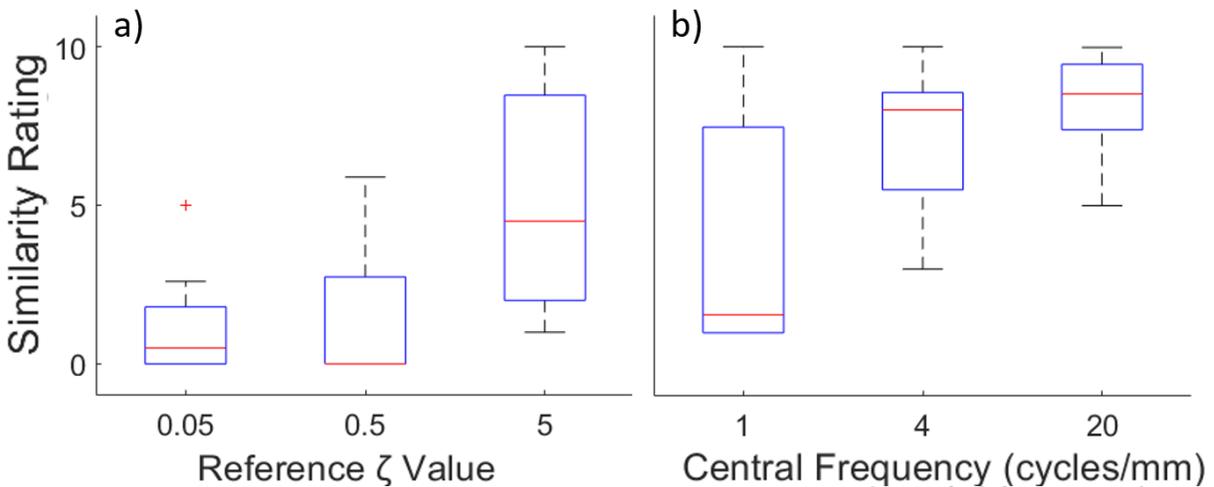


Figure 4.5. a) Box plot depicting similarity ratings when comparing the lowest- and highest-central-frequency textures for a given value of ζ . b) Box plot depicting similarity ratings when comparing the lowest- and highest- ζ textures for a given value of central frequency.

4.4. 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 Figure 4.2), but no subject was capable of performing this discrimination near the 1-mm mark. The potential to produce textural surrogates is analogous to the *homogeneity metric* of sound textures [50].

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 [3].

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 Figure 4.3). 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 Figure 4.4 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 Figure 4.4 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 Figure 4.5 b) by the similarity ratings between minimum- and maximum-Irregularity versions of a texture with 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 [44]. 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 Figure 4.5 a) by the high mean similarity rating between minimum- and maximum-Pitch versions of a texture with 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.

CHAPTER 5

The Single-Pitch Texel: A flexible and practical texture rendering algorithm

5.1. Texel Rendering: Rigorous Testing

The Single-Pitch Texel rendering algorithm takes as input the physical length of a *texel*, along with a spectral distribution location and width, and generates as output a virtual friction-modulation texture composed of a series of texels, each with a single frequency of friction oscillation pulled from the (lognormal) spectral distribution. Here, we refer to this oscillation as the texel's *pitch*. In the last chapter, this rendering algorithm was shown to efficiently recreate the sensation of a virtual fine texture with multiple frequency components present at all positions, provided that a Texel length below 1 mm was used. Furthermore, the spectral parameters of distribution location and width (referred to here as *Central Pitch* and *Irregularity*, respectively) were reliably tuned by participants to recreate reference textures, thereby representing effective dimensions for texture design.

In this chapter, we further test the utility of this rendering algorithm through matching tasks with a variety of reference textures encompassing a diverse set of wide-band spectral characteristics. Families of textures investigated included superpositions of multiple sinusoidal signals, spectrally-shaped white noise signals (so-called *colors of noise*), and periodic non-sinusoidal waveforms. Each of these bear spectral characteristics not readily reproduced by a single lognormal distribution in spectral space, as is required by

the Texel rendering. Thus, achieving perceptually-similar versions of these textures using the Texel rendering scheme represents a significant test of the algorithm’s extensibility.

Furthermore, we investigate the effect of quantizing the spectral components of a Texel-rendering texture on the perceptual outcome, a process analogous to that used widely in audio signal compression [72]. By leveraging the demonstrated inaccuracy in frequency judgement during human vibrotactile perception [46] [20], the discretization of oscillation frequencies available to the rendering algorithm could greatly reduce the amount of data required to store the texture without perceptual effect.

The possibility of an artifact left by the sudden change in friction level at the boundary of two texels was removed by enforcing a continuity of phase at the boundary of every texel. (The phase of the first texel in a rendering is randomized.) Otherwise, the algorithm used here was identical to that in the previous chapter.

5.2. Optimization techniques

While the *Central Pitch* and *Irregularity* rendering parameters were previously shown to give texture designers enough flexibility to predictably match various reference fine textures composed of white noise with Gaussian-shaped spectral amplitude, navigating the two-dimensional parameter space to recreate any arbitrary virtual texture is a non-trivial task. Modifying the *Central Pitch* and *Irregularity* parameters individually is referred to in this work as *free exploration*, which was found to be a difficult-to-use methodology for users in a texture matching task, especially for those with little experience tuning virtual

textures. Accordingly, we devised an assistive human interface based on established optimization algorithms that simplified the required user input while converging on a Texel rendering that the user perceives as a good match for a reference texture.

The use of optimization algorithms in turning virtual texture parameters has previously found success. In [64], the researchers use an evolutionary optimization algorithm to personalize the vibrotactile signals generated by a wearable device to maximize discrimination and positive sensation for the subject. Notably, the algorithm required only three pieces of preference information from the subject per iteration.

In [63], the researchers approach the problem of improving vibrotactile signal discrimination by using a genetic optimization algorithm to tune the stimulus to an individual subject's sensory preferences. While discrimination performance universally improved, the tuning parameters generated by the group of users differed greatly. This approach demonstrates the success of a genetic optimization algorithm responsive to preference input from a test subject.

The researchers of [43] combine an evolutionary algorithm with a Generative Adversarial Network (GAN) to produce perceptually-matching virtual textures for real reference textures. In this approach, when a user selects the closest match to the reference among a group of virtual textures generated by the GAN, the preference drives the evolutionary algorithm to select optimal texture parameters (input to the GAN) towards matching the reference texture. This process repeats until convergence is detected. This technique effectively approaches the best-match texture parameters using only preference data (in

this case, best match to a reference among several candidates), while presenting a substantially simpler task to a user as compared to free exploration of multiple parameter dimensions simultaneously.

5.3. Methods

5.3.1. Single-Pitch Texel Rendering

For clarity, an outline of the Texel rendering algorithm is repeated here.

Three parameters are first chosen by the texture designer: the physical length of a single Texel, the spectral mean of the pitch distribution (“Central Pitch”), and the spectral width (standard deviation) of the pitch distribution (“Irregularity”). These parameters are used to build a pitch distribution that is Gaussian in log-frequency space. Starting at the left end of the texture, a value is drawn from the pitch distribution and used as the frequency for the Texel’s “pitch”, that is, a sinusoidal oscillation filling a single Texel length (initially with random starting phase). This process is repeated for additional Texels, continuing to the right on the display (with phase continuous at the borders of the Texels) until the target texture length is achieved. For this study on fine-scaled textures, the amplitude of oscillation is held constant throughout.

Due to the log-scale nature of the input parameters of Central Pitch and Irregularity, in this work we will report both in units of dB with a reference value of 1 cycle/mm, abbreviated herein as simply “dB”. Values are computed using $X_{dB} = 10 \log_{10}(X_{cpm}/(1 \text{ cycle/mm}))$, where X_{dB} and X_{cpm} are a parameter value in dB and cycles/mm, respectively. (Here, the power-quantity formula for decibels is used for the convenience of defining 10 cpm = 10 dB.)

5.3.2. Reference Textures

To test the viability of the Texel rendering algorithm, a diverse set of multi-frequency reference textures was sought, summarized in Figure 5.1. Rigorous testing of the algorithm would ideally include textures that don't intuitively lend themselves to successful imitation by a string of single-frequency texels stochastically-drawn from a Gaussian distribution (in log-frequency space). To this end, three families of fine textures were identified. Note that all textures were scaled to have the same maximum amplitude.

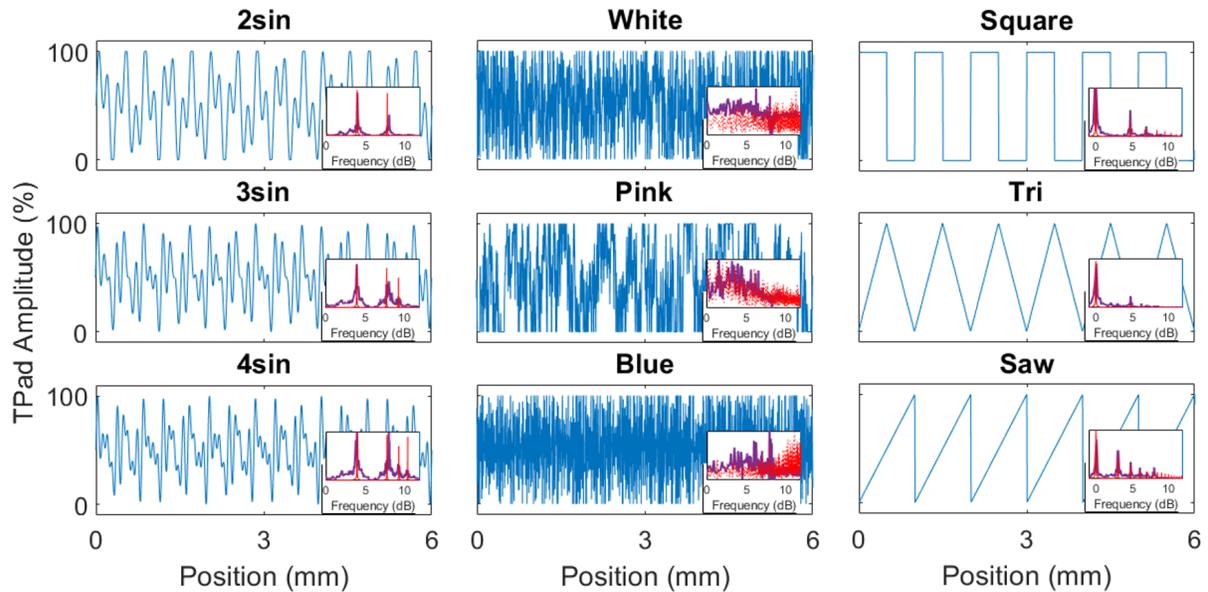


Figure 5.1. The nine reference textures used in this study represented by 6-mm-long samples of the space-domain TPad texture maps. TPad Amplitude refers to the friction reduction amplitude achieved at any point through ultrasonic vibration between the TPad surface and the test subject's finger, scaled by the maximum amplitude of the device. Inset: FFT magnitude traces for each of the textures, scaled for visibility. Red: commanded friction reduction spectral magnitude; purple: lateral fingertip velocity spectral magnitude acquired via Laser Doppler Vibrometer (LDV). (Note that dB units are calculated with a reference value of 1 cycle/mm.)

The first family, *Sums of Sinusoids*, consisted of the equally-weighted superposition of two or more sinusoidal components of differing frequencies. This resulted in oscillating, but dramatically non-sinusoidal, textural signals. Care was taken in the choice of the discrete frequencies to avoid beat frequencies below 1 cycle/mm, which would be perceivable as coarse textural features. The first texture in this family was an equally-weighted sum of a 2.5 cycles/mm sinusoid with a 6 cycles/mm sinusoid. The second summed 2.5, 6, and 8.5 cycles/mm. The third summed 2.5, 6, 8.5, and 11 cycles/mm. (These frequency values correspond to 4.0, 7.8, 9.3, and 10.4 dB, respectively.) For brevity, these textures are referred to as *2sin*, *3sin*, and *4sin*, respectively, for the remainder of this chapter.

The second family, *Colors of Noise*, collected three wide-band spectral signals inspired by analogous audio signals of the same name. Specifically, these textures included White Noise (with constant power spectral density across all renderable frequencies), Pink Noise (with power spectral density inversely proportional to frequency), and Blue Noise (with power spectral density proportional to frequency), referred to as *White*, *Pink*, and *Blue*, respectively, for brevity.

The third family, *Shaped Waveforms*, consisted of three textures where a particular arrangement of spectral amplitude and phase produce oscillating signals with recognizable shapes in the spatial domain. All of these textures were constrained to have a fundamental frequency of 1 cycle/mm (0 dB). The three waveforms selected were a square wave (containing odd-integer harmonics decreasing at a rate of 6 dB/octave), a triangle wave (containing odd-integer harmonics decreasing at a rate of 12 dB/octave), and a sawtooth wave (containing all integer harmonics decreasing at a rate of 6 dB/octave), referred to as *Square*, *Tri*, and *Saw*, respectively, for brevity.

In cases where a single representative texture from each family was required to ensure testing wasn't prohibitively long, the following three were chosen on the basis of pilot studies indicating they were sufficiently dissimilar from one another: The sum of four sinusoids from the Sums of Sinusoids family (*4sin*); Pink Noise from the Colors of Noise family (*Pink*); and the sawtooth wave from the Shaped Waveforms family (*Saw*). These will be referred to as the Family Exemplars.

5.3.3. Equipment

All testing was performed using a TPAD device. Texture maps and test subject GUI controls were controlled via MATLAB R2021b on an attached PC.

Due to the limitations set by the hardware used, Pitch and Irregularity were limited for all testing to the following ranges. Pitch: 0.01 cycles/mm to 94 cycles/mm (-20 to 19.7 dB); Irregularity: 0.001 cycles/mm to 3.16 cycles/mm (-30 to 5 dB).

To eliminate any contribution towards textural perception made by sound cues, all subjects wore headphones playing Pink Noise audio during all tests and confirmed before testing that the volume level made inaudible any sounds produced by the TPAD equipment.

5.3.4. Experimental Design

5.3.4.1. Free Exploration. To study the ability of a user to tune the two-parameter Texel Rendering space to generate a desired tactile effect, the *Free Exploration* test was designed, so-called as it tasked subjects with matching Texel renderings to reference textures by freely modifying the Central Pitch and Irregularity variables of the Texel rendering

(with texel length held constant at 0.25 mm). Note that, due to the constraint of continuous phase between texel boundaries, frequencies with period above the texel length can be rendered. For example, if a continuous series of several texels are assigned similar frequencies that fall below the texel length, the oscillation is rendered continuously across these texels.

The test sequence was as follows. During each of nine trials, the left half of the TPaD active surface (50.88 mm long) presented one of nine reference textures, while the right half presented the Texel rendering generated by the subject's current Central Pitch and Irregularity choices. Using two slider controls on the GUI (Figure 5.2), the subject was free to alter both Central Pitch and Irregularity with the stated goal of matching the reference texture to the Texel rendering. When the subject was satisfied with the match, confidence in the match was recorded on another slider control from 1 (least confident) to 5 (most confident). Following this, the Central Pitch and Irregularity settings could be submitted and the trial ended. This was repeated for each of the nine reference textures.

5.3.4.2. Quantization Test. In the last chapter, the stochastic nature of the Texel rendering algorithm was demonstrated to have little perceptual effect when the texel length was sufficiently small, effectively transforming a series of stochastically-generated vibration frequencies into the sensation of a homogeneous fine texture. To further test the robustness of the rendering algorithm to changes in the choices of texel oscillation frequencies, the *Quantization Test* was designed. Here, *quantization* refers to the discretization of available texel pitch choices.

Quantization was performed by limiting texel oscillation frequencies to a discrete set of values, effectively reducing the frequency resolution of the texture. In practice, this was

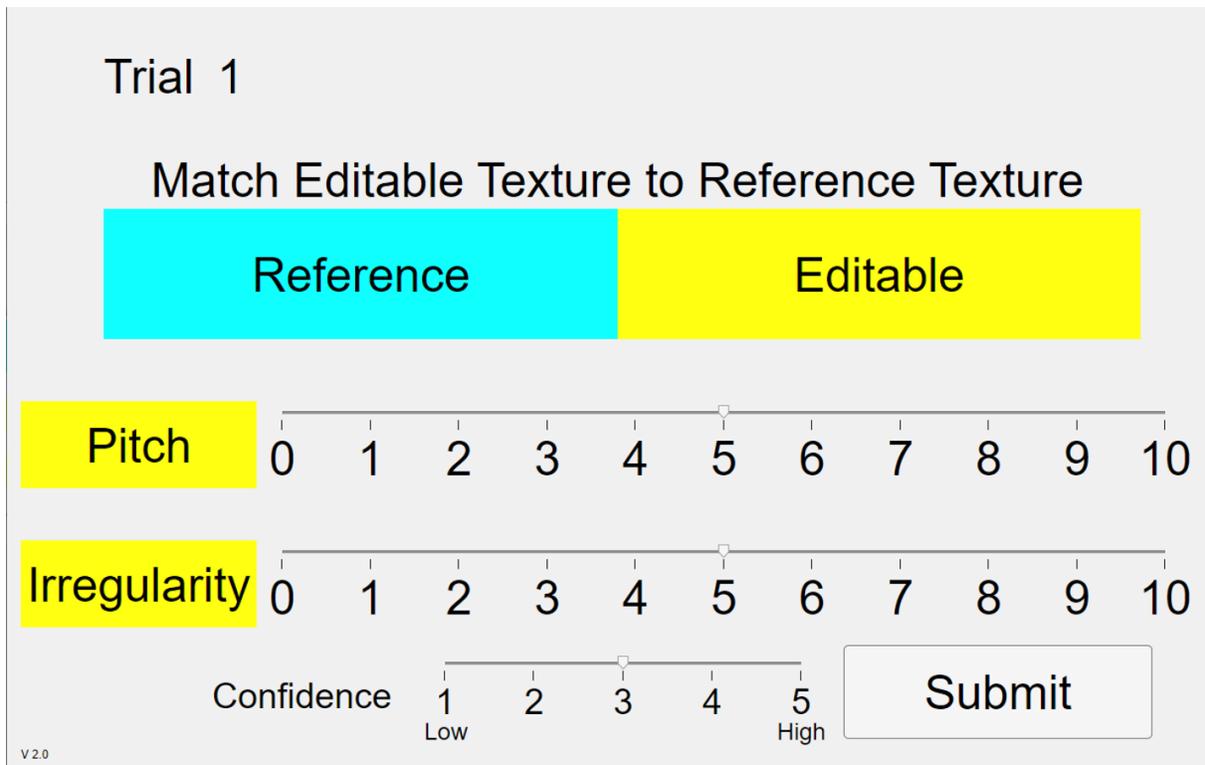


Figure 5.2. Graphical User Interface (GUI) for the Free Exploration test.

achieved by rounding the frequency to the nearest allowed value following its stochastic generation. The amount of quantization was controlled via the distance (in log space) between allowed values. Log-spaced values (rather than linear-spaced values) were chosen to explicitly match the log-normal nature of the underlying distribution itself, which was determined in the previous chapter to perform well for both low-frequency and high-frequency means. For example, if a quantization distance of 1 dB was chosen, frequencies are rounded to the nearest multiple of $1 \text{ dB} \approx 1.26 \text{ cycles/mm}$ following stochastic generation. Note that, in the non-quantized state, the frequency resolution is bounded only by the limits of the textural rendering hardware (in this study, $\sim -20 \text{ dB} \approx 0.01 \text{ cycles/mm}$).

The test sequence was as follows. During each of twenty-seven trials, the left and right halves of the TPaD active surface presented two different textures. The left half of the display presented one of nine reference textures while the right half displayed a Texel rendering based either on the test subject’s Central Pitch and Irregularity choices for the matching reference texture (made in the Free Exploration Test) or a Texel rendering based on Central Pitch and Irregularity choices made previously by the *trained rater* (the author). During the trials based on the three Family Exemplars, the Texel rendering was either non-quantized or quantized to one of three levels: 1, 3, or 5 dB. During each trial, the subject was tasked with rating the similarity between the left (reference) and right (Texel) textures on a scale between 1 (completely different) to 5 (identical), input via a slider control on the GUI (Figure 5.3). Following this, the similarity rating could be submitted and the trial ended.

5.3.4.3. Guided Exploration Test. Based on the results of the Free Exploration test, it was evident that the navigation of the 2D parameter space defined by Central Pitch and Irregularity to accomplish a matching task was difficult for a subject naive to the Texel rendering algorithm. To alleviate this difficulty, the *Guided Exploration Test* was designed to allow a subject to navigate the parameter space with simple comparison decisions, aided by an assistive algorithm.

The assistive algorithm devised for this study utilized features of the gradient descent procedure [13] and the Nelder-Mead downhill simplex method [58]. With the goal of achieving the best perceptual match between some reference texture and a Texel rendering, the algorithm searches the 2-D texel parameter space (axes: Central Pitch and Irregularity) using as input comparison data from the human subject, iteratively shrinking

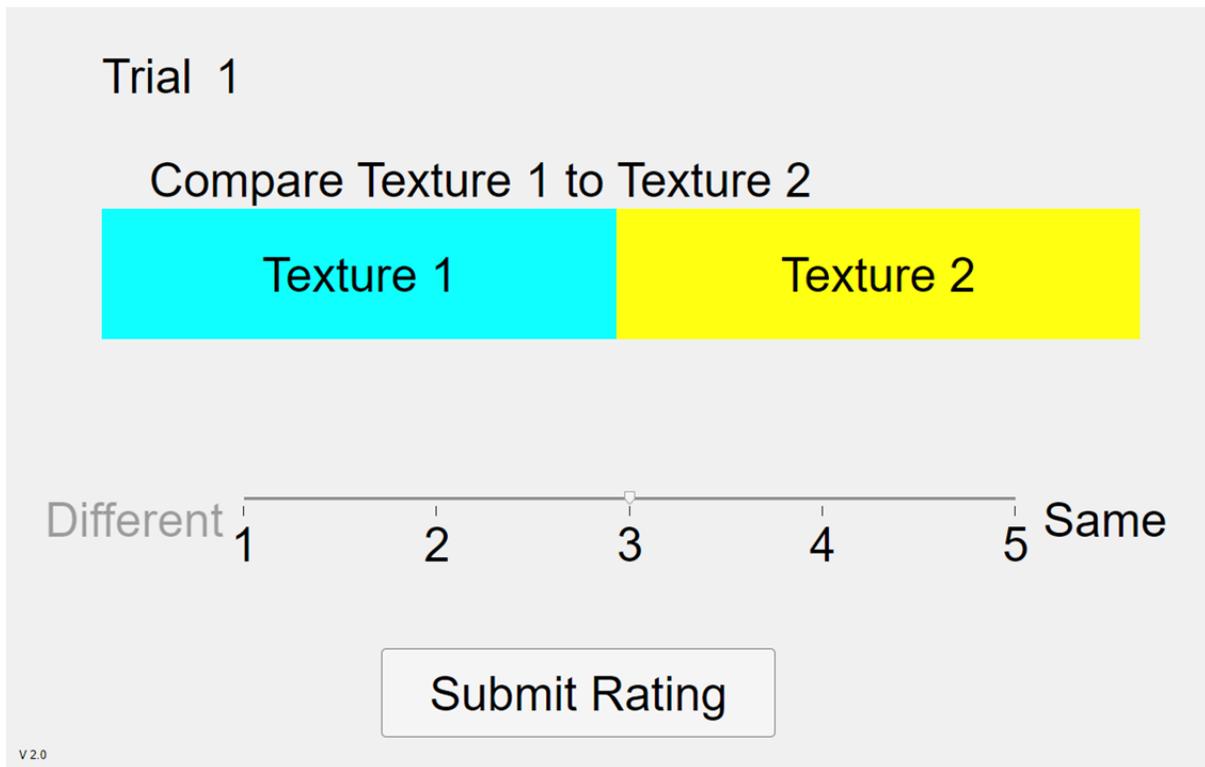


Figure 5.3. Graphical User Interface (GUI) for the Quantization test.

the target parameter space until a perceptual match is identified. During any one trial, there is an amount of this 2-D space remaining, inside which the ideal perceptual match is believed to exist. The algorithm draws vectors connecting the centroid of this remaining space to its edges along each direction parallel to an axis and along each diagonal. A comparison point is marked at each vector midpoint, totaling eight comparison points surrounding the remaining space's centroid. Each trial consists of four comparisons, as the subject is asked to choose whether a comparison point or its opposite (the comparison point lying on the parallel vector on the other side of the centroid) feels more similar to a reference texture. The subject can also choose "unsure", meaning both textures feel

equally similar to the reference. The polygon described by the eight comparison points is used analogously to the Nelder-Mead simplex [58], building a set of trial points to establish an ideal direction for optimization over the 2-D space. While the 2-D simplex generally uses the minimum of three points of function evaluation to establish an approximation of the function’s local behavior, the polygon here uses four “evaluations” (from binary comparisons of eight total points) for this purpose. This redundancy was found to perform well during pilot studies. Based on the trial’s four points of comparison data, the algorithm builds the so-called *Pursuit Vector*: the scaled sum of all of the vectors drawn from the centroid to each comparison point deemed “more like the reference texture” (without contribution from all “less like the reference texture” and “unsure” points). This vector, like the search direction in the gradient descent procedure [13], is an approximation of the direction pointing towards the optimal point (in this case, the texture parameters for an ideal match), and is the direction and magnitude of shift of the remaining parameter space’s centroid at the end of the trial. The Pursuit Vector’s scaling factor was chosen to constrain the vector’s maximum possible magnitude to be half the amount that the remaining parameter space is decreased per trial, which ensured that the remaining parameter space is always a subset of the space in the trial before it. At the end of each trial, the parameter space is decreased and its centroid moves according to the Pursuit Vector. The resulting space is used as the remaining parameter space for the next trial. The optimization algorithm terminates when the remaining space decreases to a single point, and this point is identified as the perceptual match to the reference texture. For algorithmic simplicity and speed of testing, the parameter space decreased by the same amount each trial (in this case, 20% of the initial maximum axes lengths)

along both axes, but future versions of this optimization algorithm could utilize adaptive shrink rates as well as convergence criteria not based solely on number of trials. Figure 5.4 visualizes the progress of the Guided Exploration optimization algorithm throughout the test for one subject.

The test sequence was as follows. For each of three Family Exemplars, twenty comparison trials were conducted. In each trial, the subject was able to freely select the current reference texture and two comparison textures using buttons on the GUI (Figure 5.5), causing the TPaD to present the selected texture. After exploring these three textures, the subject confirmed which (if either) of the two comparison textures felt more like the reference texture and submitted their response, ending the trial. This continued for twenty trials for each of three reference textures (totaling 60 comparison trials), after which the user completed nine rating trials. In each rating trial, the subject was asked to rate the similarity between one reference texture and one Texel rendering. These nine combinations included all permutations of the three reference textures and three Texel renderings based on the parameters identified as ideal matches by the optimization algorithm. The subject entered ratings on a scale between 1 (completely different) to 5 (identical), input via a slider control on the GUI (Figure 5.6). Following this, the similarity rating was submitted and the trial ends.

5.4. Results

There were 18 subjects for the Free Exploration Test, 8 female; 8 subjects for the Quantization Test, 4 female; and 17 subjects for the Guided Exploration Test, 6 female. 6 subjects participated in both the Free Exploration Test and Guided Exploration Test.

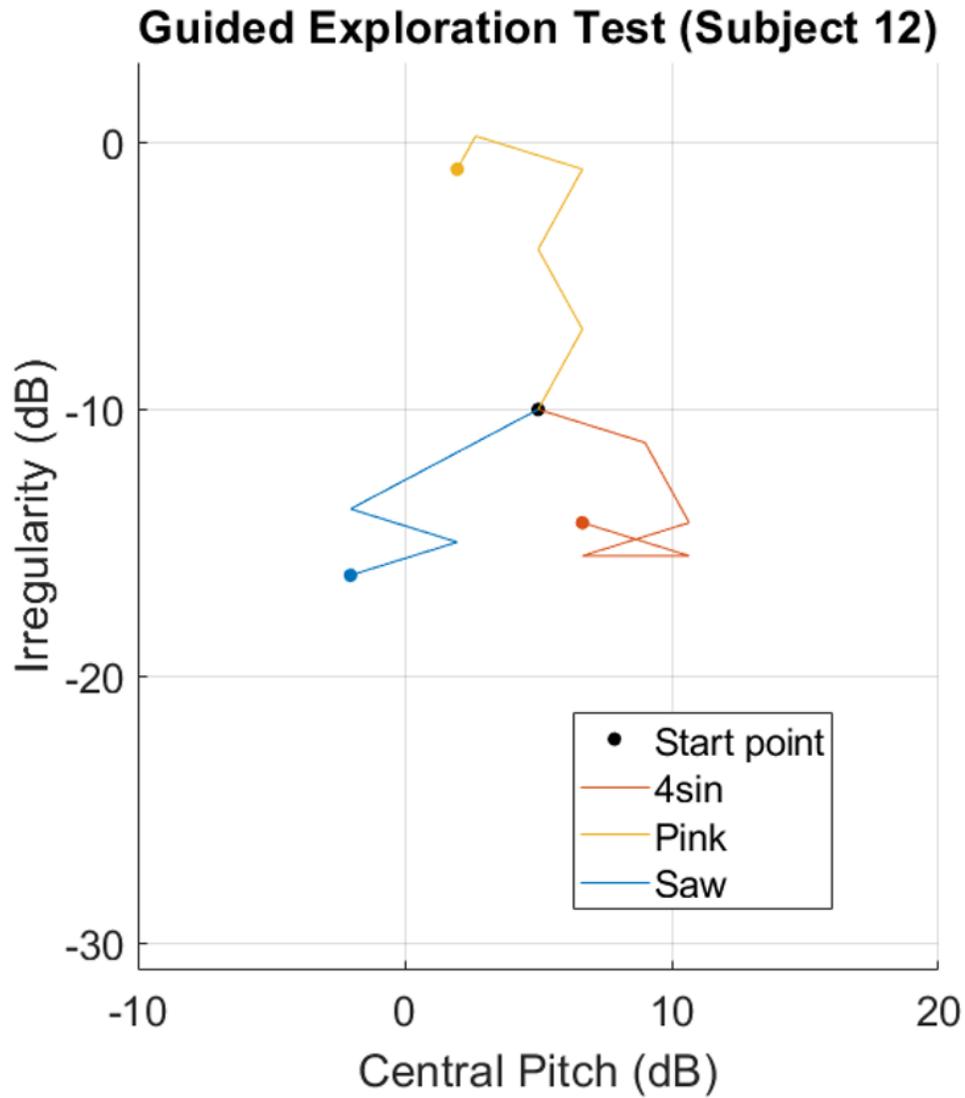


Figure 5.4. Guided Exploration test results for one subject. Lines connect subspace centroids throughout the optimization algorithm's progress (lines match the Pursuit Vector for each trial) and solid dots indicate termination points.

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

Trial 1 Tex 1

Reference

Texture A Texture B

Which texture feels more similar to the reference?

Texture A
 Texture B
 Unsure

Submit

V 1.0

Figure 5.5. Graphical User Interface (GUI) for the Guided Exploration test matching trials.

Northwestern University Institutional Review Board, all subjects gave informed consent, and all subjects were paid for participation.

In all box charts displayed here, filled boxes mark the data between the 0.25 quantile and 0.75 quantile (the interquartile range, or IQR), with the central horizontal line marking the median. Whiskers extend to the nonoutlier minimum and maximum, with outliers (marked as dots) identified as points greater than 1.5 times the IQR away from the top or bottom of the box.

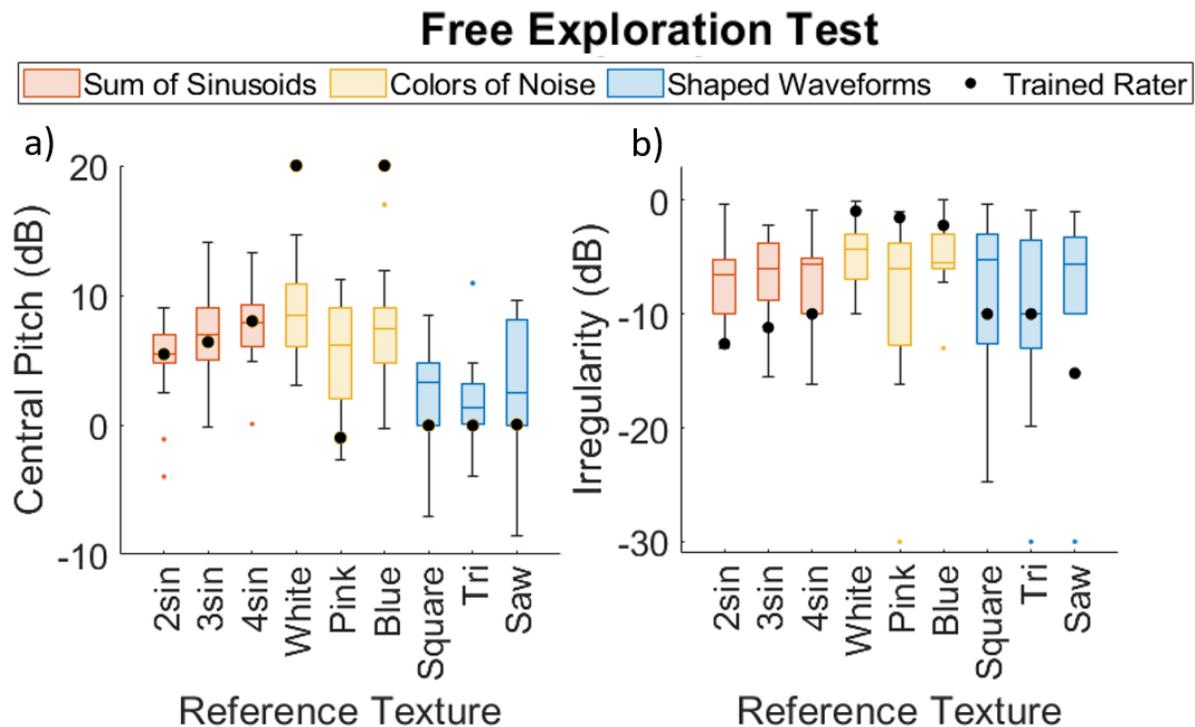


Figure 5.7. a) Central Pitch and b) Irregularity choices made during the Free Exploration Test, color-coded by texture family. Superimposed points: parameter choices of the trained rater.

Figures 5.9 a) and 5.9 b) present the parameter choices for the Free Exploration Test and Guided Exploration Test, respectively, on 2-D plots, superimposed with 95% confidence ellipses around the expected mean (adapted from [8]). Although Free Exploration was tested for all nine reference textures, for ease of comparison with Guided Exploration, we present only the Texel parameters matching the three Family Exemplars for each test.

Figure 5.10 summarizes the subjects' confidence in the Free Exploration Test aggregated across all textures. Confidence rating (divided into four bins) is plotted against distance between the test subject's Texel parameters and the trained rater's parameters.

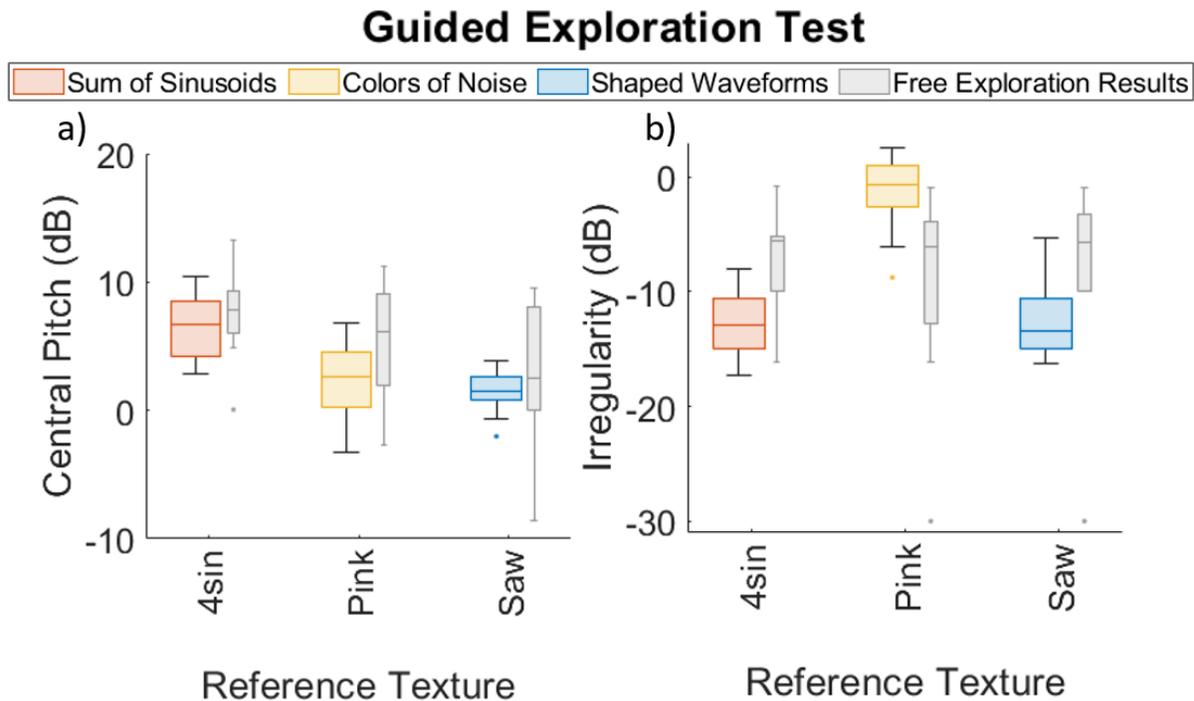


Figure 5.8. Collected a) Central Pitch and b) Irregularity choices made during the Guided Exploration Test. For comparison, Free Exploration results are superimposed.

5.4.1.2. Similarity Ratings. Figures 5.11 a) and 5.11 b) present the similarity ratings made during the Free Exploration Test and Guided Exploration Test, respectively, between the Texel rendering generated by the subject and its matching reference texture. (Note that while the Free Exploration Test used all nine reference textures, the Guided Exploration test used only the three Family Exemplars.) Figure 5.11 c) portrays in matrix format the average similarity rating between each reference texture (columns) and the Texel rendering generated by the subject in the Guided Exploration test for the named reference texture (rows).

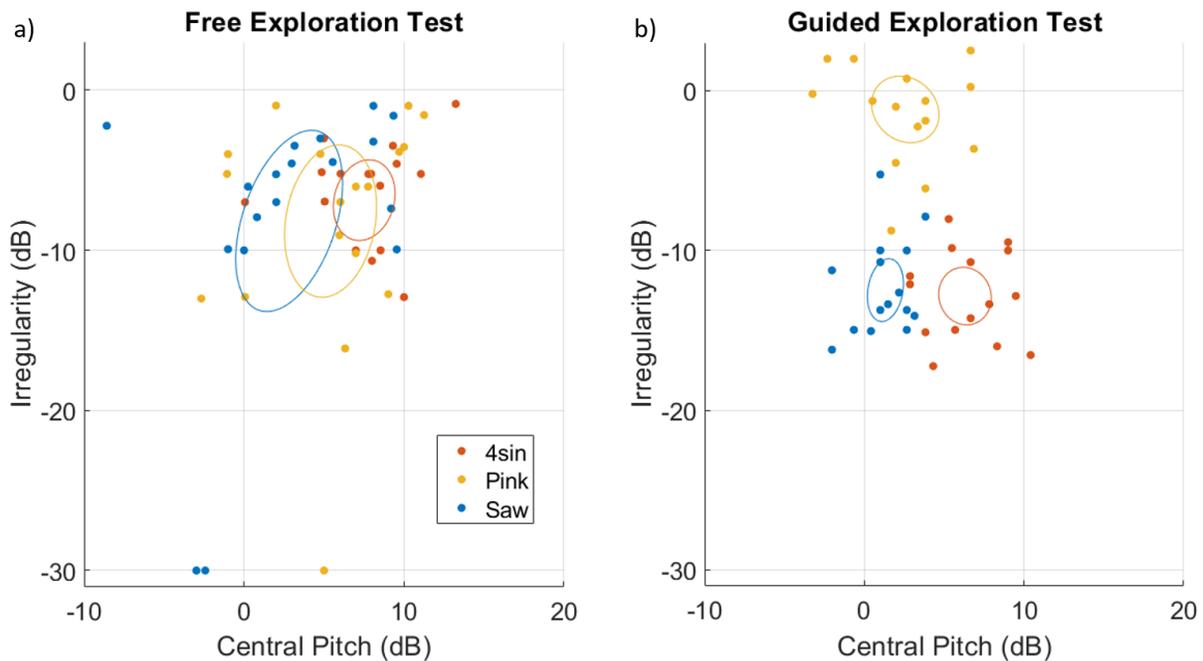


Figure 5.9. Central Pitch and Irregularity choices made during the a) Free Exploration Test and b) Guided Exploration Test for the Family Exemplars. Superimposed ellipses represent area of 95% confidence around the expected mean of a given texture.

5.4.2. Spectral Quantization

Figure 5.12 summarizes how pitch quantization affects the efficacy of the Texel rendering as observed in the Quantization Test. For each of the three Family Exemplars, the effect of quantization is calculated as the similarity rating for a reference texture and its unquantized Texel rendering minus the rating for the same reference texture and its quantized Texel rendering. In the figure, change in similarity rating is plotted for each texture at each of three levels of quantization.

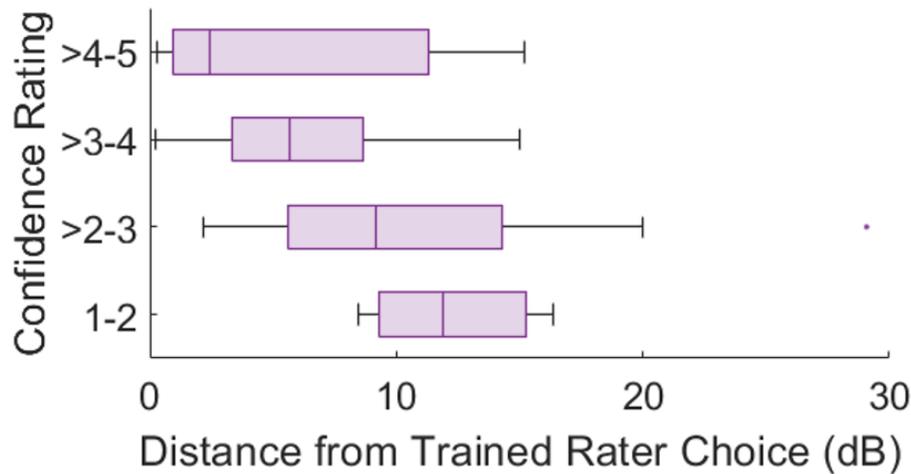


Figure 5.10. Confidence ratings for all textures made during the Free Exploration Test versus distance between a test subject's chosen Texel parameters and those made by the trained rater for the same target reference texture.

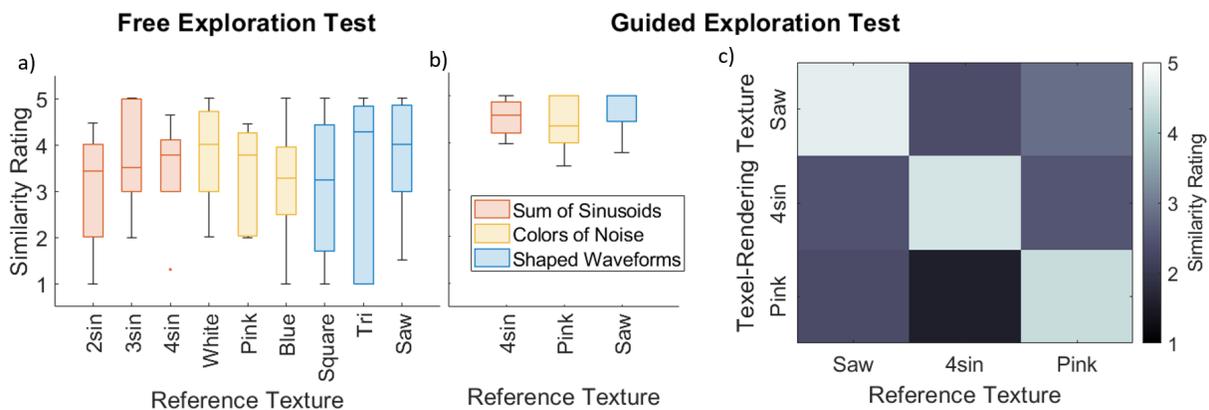


Figure 5.11. Similarity ratings between reference textures and Texel-rendering matches made during the a) Free Exploration Test and b) Guided Exploration Test. In c), a matrix portrays the Similarity ratings made during the Guided Exploration Test between each reference texture (columns) and a Texel-Rending texture designed to match the named reference texture (rows).

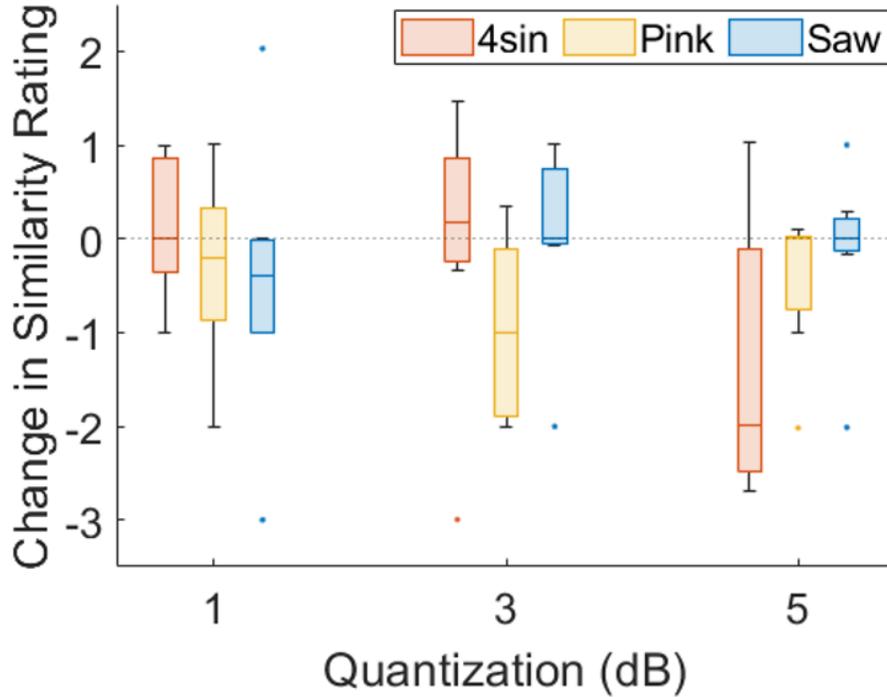


Figure 5.12. Change in similarity ratings (between reference textures and Texel-rendering matches made during the Free Exploration Test) for each amount of quantization as compared to the similarity rating made in the unquantized state. Superimposed dotted line highlights zero change.

5.5. Discussion

5.5.1. Parameter Tuning through Free Exploration

The collected texel parameters chosen during the Free Exploration Test, shown in 5.7 a) and 5.7 b), suggest that manually tuning Central Pitch (the mean of the underlying distribution) to match a reference represented an easier task than tuning Irregularity (the width of the underlying distribution). This is apparent in the higher variance in Irregularity choices (the height of the box plots in Figure 5.7 b), average standard deviation: 5.12 dB) as compared to the same in Central Pitch choices (the height of the box plots in

Figure 5.7 a), average standard deviation: 3.78 dB). Such behavior may be due to the fact that Central Pitch is naturally a more familiar perceptual dimension than Irregularity to a naive virtual texture user. While Central Pitch is analogous to the audible pitch of a sound, no such common connection exists for Irregularity. Additionally, it is likely that the just-noticeable difference (JND) for Irregularity exceeds that of Central Pitch when both are measured in cycles/mm, but the improved results from the Guided Exploration Test suggest that the variability in Irregularity choices in the Free Exploration Test is affected by factors beyond the threshold of human perception. For the Guided Exploration Test (Figures 5.8 a) and 5.8 b)), average standard deviations in Central Pitch and Irregularity choices (for all textures) are 2.42 dB and 2.90 dB, respectively. Additionally, the range of values selected during the Free Exploration test suggests that subjects were capable of tuning Central Pitch more effectively than Irregularity. While the mean value of Central Pitch varied significantly from texture to texture (standard deviation across means: 2.55 dB), the mean value of Irregularity did not (standard deviation across means: 1.54 dB). In short, Irregularity was not effectively used as a texture design parameter during the Free Exploration Test.

The Central Pitch choices made in the Free Exploration Test reveal some notable trends (see Figure 5.7 a)). In the Sums of Sinusoids texture family, the mean value of Central Pitch chosen by subjects follows a rising trend when additional (higher-frequency) sinusoids are added, closely matching the Central Pitch choices of the trained rater. This is in line with the results of [20], which demonstrated that the dominant sensation when exploring a texture composed of several discrete frequency components is a single dominant frequency located at the weighted mean of the original components. The Central

Pitch choices for this family of textures indeed follow a mean value near such a dominant frequency, increasing as additional frequency components are added to the high end of the spectrum. The sum of 2.5 and 6 cycles/mm (4.0 and 7.8 dB) ($2sin$) produces a mean Central Pitch around 3 cycles/mm (4.8 dB); the sum of 2.5, 6, and 8.5 cycles/mm (4.0, 7.8, and 9.3 dB) ($3sin$) produces a mean Central Pitch around 4 cycles/mm (6.0 dB); and the sum of 2.5, 6, 8.5, and 11 cycles/mm (4.0, 7.8, 9.3, and 10.4 dB) ($4sin$) produces a mean Central Pitch around 6 cycles/mm (7.8 dB).

In the Colors of Noise texture family, the mean Central Pitch choices loosely follow a trend related to the shape of the reference texture's spectrum: the mean Central Pitch was lower for Pink Noise (which is dominated by low-frequency components) than for Blue Noise (which is dominated by high-frequency components) and White Noise (which has a flat distribution across all frequencies). This suggests that the low-frequency content of a Shaped Noise texture is a more salient perceptual detail than the high-frequency content, dominating the sensation during exploration.

In the Shaped Waveform texture family, mean Central Pitch choices fall close to the fundamental frequency of the waveforms: 1 cycle/mm (0 dB) for all textures. As with the Sums of Sinusoids textures, this is to be expected: the dominant sensation for these textures should indeed be some weighted mean frequency of the components. Unlike the Sums of Sinusoids, however, the Shaped Waveforms are spectrally composed of a single high-amplitude fundamental frequency with many harmonic frequencies with steeply decreasing amplitude (see Figure 5.1). The weighted mean, lacking any significant contribution from the harmonic frequencies, remains close to the fundamental frequency, which closely matches the Central Pitch choices made by the trained rater.

Furthermore, the Free Exploration Test results demonstrate that variance in parameter choices, particularly Central Pitch, depend on the reference texture family. The height of the box plots in Figure 5.7 a) demonstrate different average variance in Central Pitch choices per texture family, with average standard deviation values: 3.88 dB for Sums of Sinusoids, 5.63 dB for Colors of Noise, and 3.98 dB for Shaped Waveforms. While the Central Pitch choices for the Sums of Sinusoids and Shaped Waveforms families have lower standard deviation (with means close to the Central Pitch choices made by the trained rater), the Colors of Noise family has a higher standard deviation in Central Pitch choices. This is an expected effect of the perceptual nature of the texture families: while the Sums of Sinusoids and Shaped Waveforms textures are dominated by a few frequency components and lack any wide-bandwidth noise signals, the Colors of Noise textures have no such dominant frequency components. (Compare the spectral signals of the reference textures in Figure 5.1.) Effectively, the choice of Central Pitch of the Texel Rendering for a texture in the Colors of Noise family has a smaller perceptual effect than it does for the other texture families. When matching to a high-bandwidth (noisy) texture, it seems to matter more that the Irregularity parameter is high and less that the Central Pitch parameter is located at some precise value. This is not to say that the Central Pitch parameter is unused when matching to noisy textures: note in Figure 5.7 a) that the average Central Pitch values were lower to match Pink Noise than those to match White or Blue Noise.

The Free Exploration Test confidence ratings depicted in Figure 5.10 demonstrate that subjects tended to rate their confidence in a match to a reference texture higher when their chosen parameters were close to those chosen by the trained rater. This suggests

that the high variance in parameter choices (Figures 5.7 a) and 5.7 b)) and low average similarity ratings between the reference textures and their Free Exploration Test matches (Figure 5.11 a), mean similarity rating 3.4 out of 5) may not be an inherent weakness of the algorithm but rather a byproduct of the difficulty of the task of navigating the parameter space freely to find a perceptual match to a reference texture. After all, if confidence ratings tend to be higher when Texel renderings use parameters closer to those chosen by the trained rater, it is rational to believe that matches exist but are not easily found by all subjects.

5.5.2. Parameter Tuning through Guided Exploration

When depicted on a 2-D log-log plot as in Figures 5.9 a) and 5.9 b), the improvement provided by the Guided Exploration Test’s optimization algorithm is apparent. Compared to the parameter choices made in the Free Exploration Test, those made in the Guided Exploration Test exhibit much tighter distributions in both the Central Pitch and Irregularity dimensions, as visualized by the significantly smaller confidence ellipses. It is notable that for the Guided Exploration Test results, the variance in the Irregularity choices (average standard deviation across all textures: 2.90 dB) remains higher than that of the Central Pitch choices (average standard deviation: 2.42 dB), as it did for the Free Exploration Test. This persistent difference, despite the improved method of Guided Exploration, suggests that the limiting factor may be intrinsically perceptual (i.e., a higher JND) for Irregularity than for Central Pitch when both are measured in cycles/mm.

The success of the Guided Exploration Test is further demonstrated in the similarity ratings between reference textures and a Texel rendering utilizing the subject’s chosen

matching parameters. While the similarity ratings made during the Free Exploration test did not exhibit particularly high values for all textures (Figure 5.11 a), mean similarity rating: 3.4 out of 5), those made during the Guided Exploration Test demonstrated high similarity between reference textures and Texel renderings for all textures (Figure 5.11 b), mean similarity rating: 4.6 out of 5). Additionally, the matrix in Figure 5.11 c) demonstrates that Texel renderings exhibit high similarity ratings only with the reference textures they were tuned to match: all “mismatched” pairings exhibited low similarity ratings (mean similarity rating: 2.1 out of 5). This confirms that the increase in similarity ratings for matches between the Free Exploration Test and Guided Exploration Test were indeed effects of superior perceptual matching and not an effect related to the difference in test methodology (such as the significantly increased test duration in the Guided Exploration Test for each reference texture as compared to that off the Free Exploration Test).

5.5.3. Pitch Quantization

The effects of pitch quantization as shown in Figure 5.12 demonstrate several effects. First, it is evident that mean similarity ratings decrease with increasing level of quantization (mean change in similarity rating: -0.23, -0.33, and -0.64 for 1, 3, and 5 dB quantization, respectively). Using a repeated measures ANOVA test, the 4sin texture showed a significantly decreased similarity rating with quantization ($p = 0.03$) and the Pink texture approached significance ($p = 0.08$). The Saw texture, in contrast, was essentially unaffected by the level of quantization ($p = 0.80$).

Further analyses used one-sample t-tests comparing the mean change in similarity to zero. A null result is taken to indicate a negligible perceptual impact of quantization at the given level. On this basis it is unlikely that quantization to the 1 dB level has any perceptual effect for any of the textures tested, as the similarity ratings between Texel renderings and reference textures for the unquantized and 1 dB level quantization groups did not approach a statistically reliable difference ($p = 0.72, 0.42,$ and 0.43 for 4sin, Pink, and Saw textures, respectively). This suggests that 1 dB level quantization is a safe method for texture designers to reduce the required data to store a virtual texture (by greatly reducing the number of possible Texel pitch values) without perceptual effect to the end user. Additionally, it is evident that the effect of quantization is not homogeneous across texture families. Where the Saw texture saw no significant change in similarity rating even at the highest level of quantization, 5 dB, ($p = 0.72$) and the 4sin texture only showed a significant decrease in similarity rating between the unquantized and 5 dB level quantization groups ($p = 0.05$), the Pink texture was much more sensitive to quantization. At the 3 dB quantization level, a significant decrease in similarity ratings is already apparent ($p = 0.04$). Comparing these results to the Central Pitch / Irregularity choices for these textures in Figures 5.7 a) and 5.7 b), it can be inferred that quantization is less perceptible for textures with lower Irregularity values. This result can be explained by the nature of human vibrotactile perception. An unquantized low-Irregularity texture contains many pitch values close to a mean value. Owing to the lossy nature of human vibrotactile perception, the dominant sensation when exploring this texture would be some weighted mean frequency of excitation even when multiple frequencies are present, as demonstrated

in [20]. When this texture is quantized, it amounts to shifting the closely-distributed frequencies to a few, or even a single, value. As the texture is already being dominated by a single-frequency sensation, this shift is perceivable only when the quantized frequency is misaligned with the original dominant frequency by a sufficient margin. On the other hand, an unquantized high-Irregularity texture contains many pitch values across a large frequency band. Both the sensations of a dominant weighted mean frequency and texture *noisiness* are perceived, as demonstrated in [19]. In this case, quantization is more easily perceivable at a lower level: as frequency values are “rounded” to their nearest quantized value, the texture’s dominant frequency may remain stable (as with the low-Irregularity textures), but the *noisiness* sensation is more easily affected. This is likely a result of the wide-band nature of the high-Irregularity textures: as the widely-distributed frequency values are quantized, there is more opportunity for concentration in one frequency band that was not originally dominant, producing a perceivable effect. The result suggests that, to guarantee successful quantization, the spectral nature of the texture should first be analyzed: the wider the distribution of frequency values in the original texture, the lower the amount of quantization that can be safely used.

CHAPTER 6

Coarse Features in Texel Rendering

The Texel rendering algorithm has been shown to successfully reproduce the sensation of fine textural features using a statistical representation. A further test of the algorithm's efficacy would be determining its capacity to successfully convey the sensation of coarse textural features simultaneously. While the physical spacing of features is generally taken as the threshold between *fine* and *coarse* textures, general agreement has not been made on the specific length marking this threshold, with experimental results pointing to values ranging between 0.25 and 1.6 mm [56] [53] [15] [80]. An alternate means of differentiating fine and coarse textures lies in the perceptual formats in which they are represented during touch exploration. As demonstrated in Chapter 3, coarse textural features are stored in a spatio-temporal segment of working memory (the *visuospatial sketchbook*), while fine textural features are not. Indeed, it is the spatial dependency of coarse textures that prohibits a statistical representation: the physical position of coarse textural features is perceived during touch exploration, a fact that is not true for fine textural features. With this in mind, a representation of coarse features requires absolute spatial position to remain intact.

6.1. Coarse Feature Rendering via Amplitude Modulation

To achieve a position-locked coarse feature representation, the amplitude parameter of the Texel rendering algorithm is modified. In Chapters 4 and 5, the amplitude of friction

modulation was held constant for each texel; namely, each texel had maximum oscillation amplitude possible with the TPaD device, spanning the lowest to the highest value of friction reduction. The result was a series of texels with a “flat” envelope.

By adjusting the amplitude of friction modulation in each texel, coarse textural features can be superimposed onto the fine texture of a texel rendering. Given a sufficiently small texel length, an approximation of any pattern of coarse textural features can be achieved via this amplitude modulation. Furthermore, this coarse pattern can be stored as an amplitude map vector (with length equal to the total number of texels) and is thus resilient to the stochastic fine texture rendering process. This is visualized in Figure 6.1.

While this scheme allows for the simultaneous rendering of spatial coarse features and non-spatial (statistical) fine features, it must be demonstrated that the combination is also perceived during tactile exploration. To this end, the following experiment was performed.

6.2. Experimental Design

An adaptive threshold test was chosen to compare spatial coarse texture discrimination performance with and without statistical fine texture content. The forced-choice discrimination task asked subjects to sort textures into two categories based solely on their coarse features. These categories were defined by two reference textures, both of which consisted solely of coarse features (feature separation above 5 mm). Recognizing that coarse texture perception relies on a spatio-temporal representation, the two reference textures were chosen to be mirror images of one another. In this way, the two reference

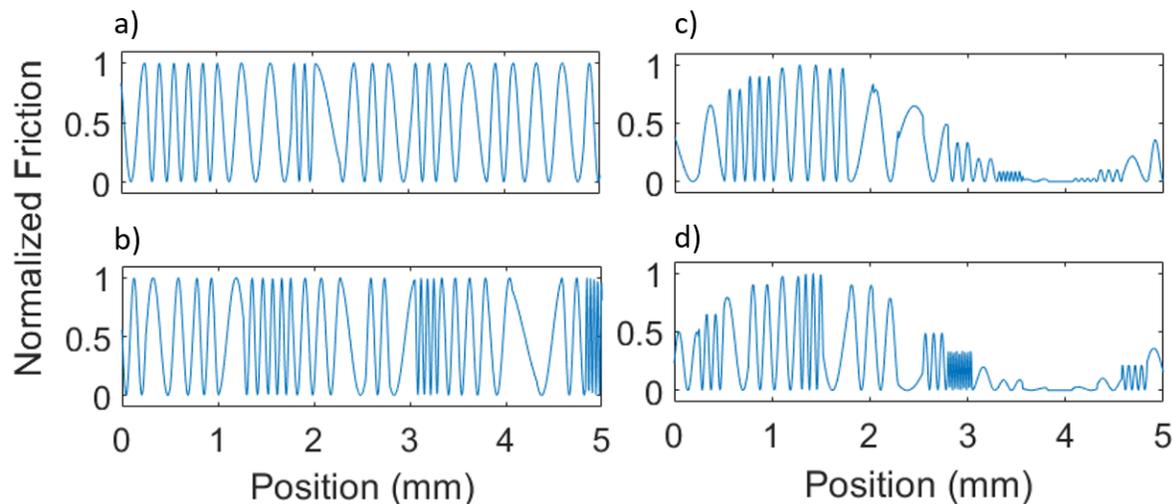


Figure 6.1. Texel algorithm used to render multiple textures using identical fine texture parameters (Central Pitch: 10 dB; Irregularity: -5 dB; Texel length: 0.25 mm). a) Rendering with amplitude held constant across texels (purely fine texture). b) Another rendering with amplitude held constant across texels. The stochastic nature of texel pitch selection results in another texture with similar statistical characteristics but different spatial characteristics. c) Rendering with amplitude modulated to include coarse sinusoidal features with period 5 mm (spatial frequency: -7 dB). d) Another rendering with identical amplitude modulation. Stochastic texel pitch selection does not affect coarse textural feature location.

textures had identical spectral magnitudes and could only be discriminated through perception involving spatial awareness of coarse features. Furthermore, as in Chapter 3, both reference textures were composed of sequences of identical coarse features (in this case, sinusoidal pulses with 5 mm period) with varying space between features (see Figure 6.2).

To independently scale both coarse and fine feature representation on a friction modulation device, two new parameters were developed: coarse strength and fine strength. Coarse strength measures the fraction of the maximum friction reduction amplitude utilized to render coarse features, while fine strength measures the fraction of the remaining

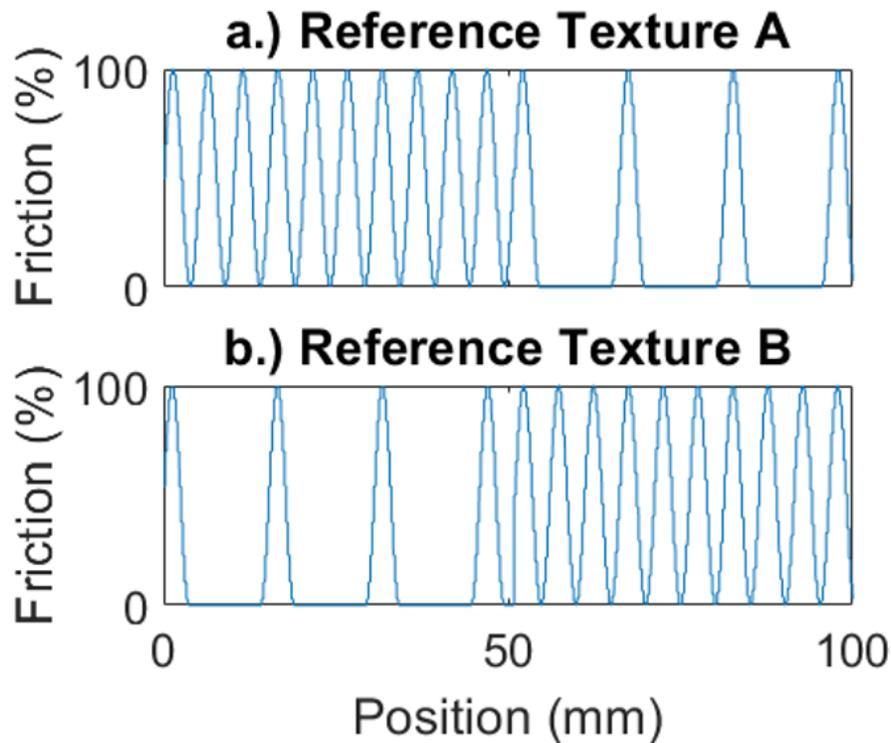


Figure 6.2. Reference textures provided for the coarse texture discrimination test. Values given for "Friction" correspond to the friction range that the TPaD is capable of producing.

available friction reduction amplitude utilized to render fine features. Note that the particular scheme of the rendering algorithm (fine features achieved via stochastic texel pitch, coarse features achieved via amplitude modulation texel-to-texel) requires that the maximum amplitude of fine feature oscillation (i.e., 100% fine strength) for any given texel is bounded by the texel's amplitude (defined by the coarse feature rendering). These two parameters allow for independent scaling of the sensation intensity of coarse and of fine features, as visualized in Figure 6.3.

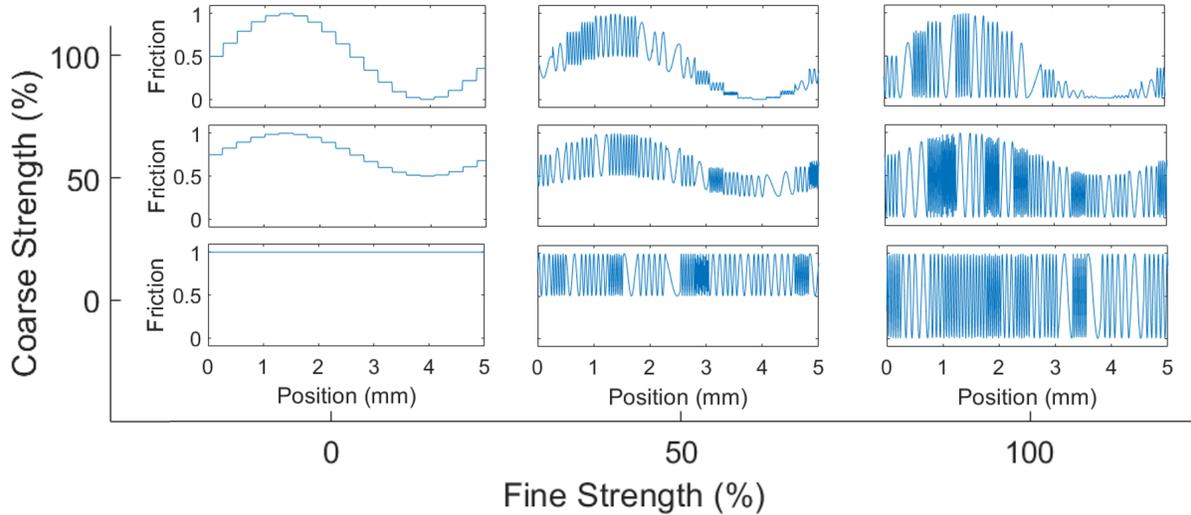


Figure 6.3. Examples of virtual textures demonstrating independent scaling of fine strength (horizontal axis) and coarse strength (vertical axis) for a given set of Texel rendering parameters (Central Pitch: 10 dB; Irregularity: -5 dB; Texel length: 0.25 mm; sinusoidal coarse features with spatial frequency: -7 dB). Values given for "Friction" are normalized.

The test was divided into three segments. In the first segment, Test texture fine strength was held constant at 0% while coarse strength was varied based on subject performance. The two following segments were identical other than Test texture fine strength being held constant at 50% and 100% for the second and third segments, respectively.

In each segment, the test procedure was as follows. Using a GUI (see Figure 6.4) accessed through a laptop PC running MATLAB r2022a, subjects could freely select either Reference texture (unchanging through the duration of the test) or the current Test texture (which changed every trial). When a texture was selected, it was displayed on the TPAD and could be freely explored. During each trial, the subject was tasked with indicating the Reference texture that felt most like the Test texture and submit this

choice via the GUI. This action ended the trial. Trials were repeated in this way until the test's conclusion.

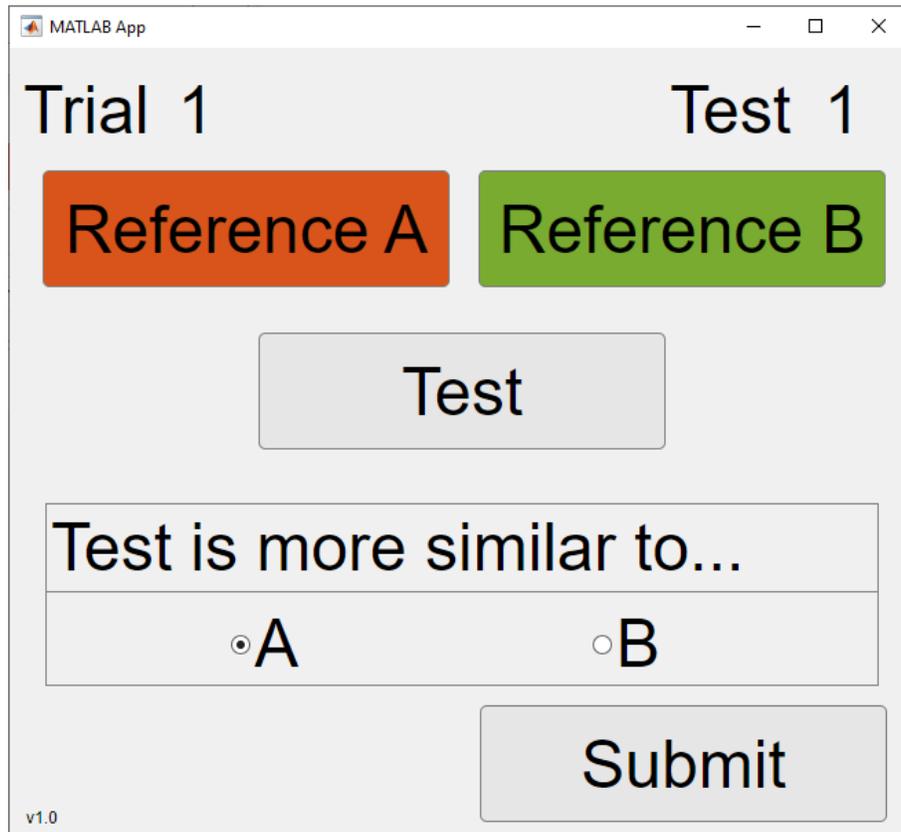


Figure 6.4. Graphical User Interface (GUI) for the coarse feature discrimination test.

During the first trial of each of the three test segments, the Test texture was rendered with 100% coarse strength. The coarse strength was scaled using a one-up three-down (1U3D) adaptive staircase technique, decreasing coarse strength with correct responses. Initially, the coarse strength was decreased by a step size of 10%. Following the first reversal in direction, this step size was decreased to 5%. Each test segment concluded after three reversals. The final coarse strength value of each test segment was taken as

the *convergent* value and is used as a rough estimate of the threshold coarse strength at which coarse feature discrimination cannot be reliably performed above chance.

All textures used a texel length of 0.25 mm. For the second and third test segments, the fine features were rendered with texel parameters of: Central Pitch: 10 dB; Irregularity: -5 dB. (The calculation of dB values for spatial frequency matches that of Chapter 5.)

6.3. Results

There were 7 subjects for this test, 4 female. Subjects were instructed to explore textures with one hand and control the GUI with the opposite hand. All subjects chose to explore textures with the right hand. The protocol was approved by the Northwestern University Institutional Review Board, all subjects gave informed consent, and all subjects were paid for participation.

Figure 6.5 summarizes subject performance for all three test segments, displaying the coarse strength setting for each subject during each trial. In this figure, data has been offset by less than 1% for visibility.

Figure 6.6 visualizes convergent coarse strength values by displaying individual subject performance (6.6 a)) and 95% confidence intervals around mean convergent coarse strength values across all subjects (6.6 b)). In 6.6 a.), data has been offset by less than 1% for visibility.

6.4. Discussion

The subject performance in the coarse feature discrimination test suggests two conclusions. First, that the rendering algorithm enables subjects to successfully distinguish coarse textural patterns consistently, even among naive users of virtual texture display

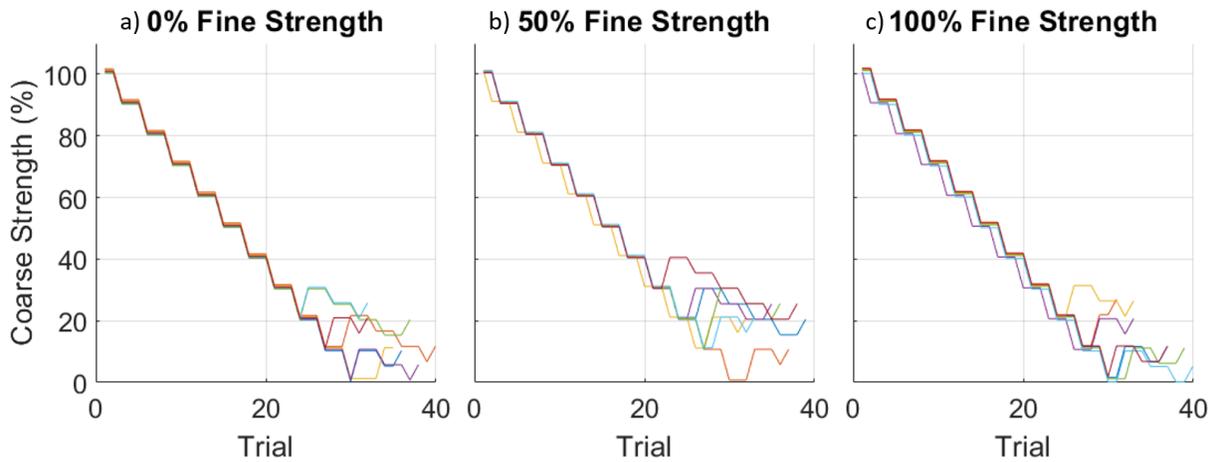


Figure 6.5. Subject performance in the coarse texture discrimination test for test segments with fine strength held constant at a.) 0%, b.) 50%, and c.) 100%. Note that all lines have been offset by less than 1% for visibility.

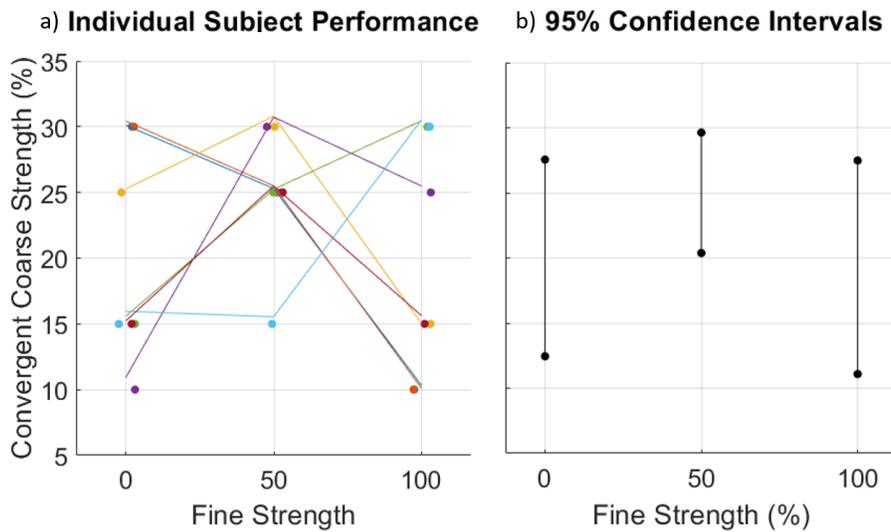


Figure 6.6. Convergent coarse strength results for each test segment. a.) Individual subject performance given as points with lines connecting performance across segments for each subject. All points and lines have been offset by less than 1% for visibility. b.) 95% confidence intervals around mean convergent coarse strength values for each test segment.

devices. Second, that the threshold in coarse strength at which subjects can no longer reliably distinguish coarse features is independent of the appearance or strength of simultaneous fine texture sensation.

Figure 6.5 demonstrates that coarse feature discrimination is possible using the Texel rendering algorithm, even for individuals naive to virtual texture display technology. In all three test segments, all subjects performed with perfect performance when coarse strength exceeded 40%. This threshold is likely specific to the friction force and bandwidth capacity of the particular texture display technology being used, but it indicates that the Texel rendering algorithm is successful at reproducing the sensation of coarse textural features discriminable through spatial differences (and without differences in spectral magnitude).

Figure 6.6 suggests that subject performance in coarse feature discrimination, as measured by convergent coarse strength, is not affected by the appearance nor strength of simultaneous fine texture sensation. In 6.6 a.), it is apparent that subjects exhibit no consistent trend relating convergent coarse strength to simultaneous fine strength. While some subjects converged at a higher coarse strength for the 0% fine strength case, others reversed this trend. Further, 6.6 b.) shows extensive overlap, across all levels of fine strength, of the 95% confidence intervals around the mean convergent coarse strength value. Should there exist any underlying effect of fine strength level on convergent coarse strength, these results indicate that the size of the effect would be trivial, and the unreliability across individuals would necessitate statistical power well beyond the range of conventional experimentation.

Taken together, these two conclusions suggest that the Texel rendering algorithm is capable of simultaneous and independent rendering of fine and coarse features. Spatial

discrimination of coarse features is possible using this rendering algorithm and is unaffected by the algorithm's statistical fine texture representation.

CHAPTER 7

Conclusions**7.1. Texel Rendering Algorithm Summary**

In this thesis, I have introduced the Texel rendering algorithm, designed to produce a rich, multi-scale virtual texture with very little input data. The motivation for such an algorithm is twofold. First, the rapid transmission of haptic and tactile signals via telecommunications networks as proposed for the Tactile Internet necessitates a compression scheme for these signals [73]. Second, the development and testing of such an algorithm offers to elucidate several aspects of human tactile texture perception.

In chapter 3, I described a test demonstrating that the component of working memory [66] utilized to recall the tactile sensation of a virtual texture depends on the length-scale of the features in this texture. Specifically, textures composed of features with average distance greater than 4 mm utilized the spatio-temporal working memory segment known as the visuospatial sketchpad, while textures with smaller average feature distances did not. This was confirmed using a distractor task known to preferentially attack the visuospatial sketchpad. These results suggest a novel means by which to classify coarse and fine textures. Since the *duplex theory of texture perception* first described by Katz [38], the conventional boundary between coarse and fine textures was defined by the average distance between features. In realistic textures, it is logical to expect stimuli that exist in both fine and coarse length scales. This study suggests that the ability for humans

to discriminate textural features based on spatial cues alone represents the deciding factor as to whether these features are *fine* or *coarse*. From the viewpoint of building virtual textures, this study confirms the requirement for retaining spatial information in coarse, but not fine, textural features.

In chapters 4 and 5, I described the Texel rendering algorithm in detail and provided test results confirming various aspects of the algorithm’s utility in rendering a diverse set of fine textures. The development of this algorithm took inspiration on previous work involving statistical representations of sound textures [48] [50], spatially-located texture building blocks (and the source of the term *texel*) [54], and the representation of high-bandwidth fine textures using a relatively few statistical parameters [19]. The Texel rendering algorithm is capable of producing perceptually identical fine textures (textural *surrogates*) from a stochastic process, taking as input only three parameters: Central Pitch, Irregularity, and Texel Length. Confirmed via psychophysical studies, these parameters could be tuned to produce close perceptual matches to a remarkably diverse set of fine textures, including wide-band noise signals and waveforms of specific shape. By leveraging a novel methodology drawing on optimization techniques, even users unfamiliar with virtual texture display technology could tune these parameters to match reference textures successfully. These results suggest that the algorithm is successful in achieving fine texture rendering with very few required inputs, essentially compressing the entire tactile sensation into three data points. Furthermore, these studies contribute towards illuminating the mechanism by which humans experience fine textures. In addition to the previously-established phenomenon of perceiving a multi-frequency tactile signal as single-frequency [20], these results suggest that this perceptual frequency-averaging phenomenon

exists even when the stimuli consisting of individual frequency components are separate in time, a mechanism already known to occur in sound perception [51].

In chapter 6, I described the extension of the Texel rendering algorithm to include coarse textural features: modulation of the maximum friction reduction amplitude texel-to-texel. By providing the algorithm with a map of coarse features, the stochastically-produced fine texture representation can be overlaid with coarse textural features, allowing both to be rendered simultaneously. I described a test performed to confirm that coarse texture discrimination is possible using this rendering scheme, finding that the addition of fine textural features had no effect on coarse texture discrimination. This extension of the algorithm allows rendering of textures with any length scale, bounded only by the capability of the display technology. Drawing on the results of chapter 3, this rendering algorithm gives virtual texture designers a framework to segregate those textural aspects that require spatial information to remain intact (e.g., coarse features) from those that could gain an efficiency benefit from a statistical representation (e.g., fine features).

7.2. Future Work

Following the research of this thesis, three major areas of open questions remain.

First, the tie between coarse texture recall and the visuospatial sketchbook as demonstrated in chapter 3 suggests that some other mechanism within the working memory framework is responsible for the recall of fine textural features. This mechanism is as-yet unknown, and distractor tasks analogous to those described in chapter 3 may be capable of discerning the specific component or components of working memory utilized in fine

texture perception. Elucidating this may further clarify the stimuli required to fully reproduce the sensation of a fine texture, allowing for more efficient renderings of virtual textures.

Second, recent studies have demonstrated that vibrational stimuli distributed across the finger-pad are perceptually salient, even when confined to frequencies historically associated with fine textures [26], so it would benefit the rendering algorithm to be extended into a framework that can apply spatially-distributed stimuli simultaneously. Although the TPaD technology used for the studies in this thesis is capable only of modulating the friction experienced by the entire fingerpad, additional research could confirm the efficacy of a spatially-distributed version of the Texel rendering algorithm. One open question is whether Irregularity, measured as the width of the pitch-generating distribution, is perceptually identical when the individual pitches are played simultaneously in a spatially-distributed set of stimuli as compared to sequentially in a time-segregated series of stimuli (as it is in the current state of the algorithm). Furthermore, while rendering of two-dimensional spatial coarse features like edges could be achieved using the existing algorithm's amplitude modulation scheme, it is unclear whether two-dimensional spatial maps of fine features will remain perceptually realistic. A two-dimensional swipe direction necessitates consideration of how a direction-independent spatial fine texture frequency can be achieved, which remains an open question.

Finally, it is apparent that the frictional (and vibrational) stimuli provided by this rendering algorithm are only one facet of a rich touch experience. Beyond the perceptual cues provided by frictional forces and bulk vibration, the human somatosensory system also perceives many other surface aspects during contact with a texture, such as fabric

pile, surface compliance, and thermal conductivity [40]. With the goal of providing a fully realistic touch sensation during virtual texture exploration, the Texel rendering algorithm could be used as a single component in a set of diverse algorithms operating simultaneously. One open question is whether this algorithm can be modified to include some other textural dimension (such as thermal effects) by adding one or more input parameters (and operating with a suitably capable rendering device). Further, research will be required to determine the extent to which the Texel rendering algorithm can be adapted to other force-feedback technologies (outside of the ultrasonic friction reduction devices used in the research for this thesis). The flexibility and ease of use demonstrated here in the Texel rendering algorithm can certainly be a starting point for future algorithms with the ultimate goal of providing an efficient representation of a realistic texture in a digital format.

References

- [1] Haptic codec task group. http://grouper.ieee.org/groups/1918/1/haptic_codecs/index.html. Accessed: 2020-03-01.
- [2] Victoria E. Abraira and David D. Ginty. The sensory neurons of touch. *Neuron (Cambridge, Mass.)*, 79(4):618–639, 2013.
- [3] Merav Ahissar and Shaul Hochstein. Task difficulty and the specificity of perceptual learning. *Nature*, 387:401–406, 1997.
- [4] Mehmet Ayyildiz, Michele Scaraggi, Omer Sirin, Cagatay Basdogan, and Bo N J Persson. Contact mechanics between the human finger and a touchscreen under electroadhesion. *Proceedings of the National Academy of Sciences - PNAS*, 115(50):12668–12673, 2018.
- [5] Alan D. Baddeley and Graham Hitch. Working memory. volume 8 of *Psychology of Learning and Motivation*, pages 47–89. Academic Press, 1974.
- [6] Alan D. Baddeley, Neil Thomson, and Mary Buchanan. Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 14(6):575 – 589, 1975.
- [7] Cagatay Basdogan, Frederic Giraud, Vincent Levesque, and Seungmoon Choi. A review of surface haptics: Enabling tactile effects on touch surfaces. *IEEE transactions on haptics*, 13(3):450–470, 2020.
- [8] Edward Batschelet. *Circular statistics in biology*. Mathematics in biology. Academic Press, London ; New York ; Sydney, 1981.
- [9] Wael Ben Messaoud, Eric Vezzoli, Giraud Frédéric, and Betty Lemaire-Semail. Pressure dependence of friction modulation in ultrasonic devices. In *Proceedings of the 2015 World Haptics Conference*, 06 2015.
- [10] L. R. Brooks. The suppression of visualization by reading. *Quarterly Journal of Experimental Psychology*, 19(4):289–299, 1967.

- [11] David Arthur Burns, Roberta L. Klatzky, Michael A. Peshkin, and J. Edward Colgate. Spatial perception of textures depends on length-scale. In *2021 IEEE World Haptics Conference (WHC)*, pages 415–420, 2021.
- [12] David Arthur Burns, Roberta L Klatzky, Michael A Peshkin, and J. Edward Colgate. A low-parameter rendering algorithm for fine textures. *IEEE transactions on haptics*, 15(1):57–61, 2022.
- [13] Augustin-Louis Cauchy. Méthode générale pour la résolution des systèmes d'équations simultanées. *C. R. Acad. Sci. Paris*, 25:536–538, 1847.
- [14] R. Chaudhari, C. Schuwerk, M. Danaei, and E. Steinbach. Perceptual and bitrate-scalable coding of haptic surface texture signals. *IEEE Journal of Selected Topics in Signal Processing*, 9(3):462–473, April 2015.
- [15] S. A. Cholewiak, K. Kim, H. Z. Tan, and B. D. Adelstein. A frequency-domain analysis of haptic gratings. *IEEE Transactions on Haptics*, 3(1):3–14, Jan 2010.
- [16] Nelson Cowan. Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological bulletin*, 104 2:163–91, 1988.
- [17] Heather Culbertson, Samuel B Schorr, and Allison M Okamura. Haptics: The present and future of artificial touch sensation. *Annual review of control, robotics, and autonomous systems*, 1(1):385–409, 2018.
- [18] G. P. Fettweis. The tactile internet: Applications and challenges. *IEEE Vehicular Technology Magazine*, 9(1):64–70, March 2014.
- [19] Rebecca F. Friesen, Roberta L. Klatzky, Michael A. Peshkin, and Edward Colgate. Building a navigable fine texture design space. *IEEE Transactions on Haptics*, pages 1–1, 2021.
- [20] Rebecca Fenton Friesen, Roberta L. Klatzky, Michael A. Peshkin, and J. Edward Colgate. Single pitch perception of multi-frequency textures. In Katherine J. Kuchenbecker, Gregory J. Gerling, and Yon Visell, editors, *2018 IEEE Haptics Symposium, HAPTICS 2018, San Francisco, CA, USA, March 25-28, 2018*, pages 290–295. IEEE, 2018.
- [21] Rebecca Fenton Friesen, Michaël Wiertlewski, and J. Edward Colgate. The role of damping in ultrasonic friction reduction. In *2016 IEEE Haptics Symposium (HAPTICS)*, pages 167–172, 2016.

- [22] F Giraud, M Amberg, B Lemaire-Semail, and G Casiez. Design of a transparent tactile stimulator. In *2012 IEEE Haptics Symposium (HAPTICS)*, pages 485–489. IEEE, 2012.
- [23] Gowri Shankar Giri, Yaser Maddahi, and Kouros Zareinia. An application-based review of haptics technology. *Robotics (Basel)*, 10(1):29–, 2021.
- [24] Sebastjan Glinsek, Mohamed Aymen Mahjoub, Matthieu Rupin, Tony Schenk, Nicolas Godard, Stéphanie Girod, Jean-Baptiste Chemin, Renaud Leturcq, Nathalie Valle, Sébastien Klein, Cédric Chappaz, and Emmanuel Defay. Fully transparent friction-modulation haptic device based on piezoelectric thin film. *Advanced functional materials*, 30(36):2003539–n/a, 2020.
- [25] Sebastjan Glinsek, Longfei Song, Veronika Kovacova, Mohamed A. Mahjoub, Nicolas Godard, Stéphanie Girod, Jean-Luc Biagi, Robert Quintana, Thomas Schlee, Matthieu Guedra, Matthieu Rupin, and Emmanuel Defay. Inkjet-printed piezoelectric thin films for transparent haptics. *Advanced materials technologies*, 7(8):2200147–n/a, 2022.
- [26] Roman V Grigorii, J. Edward Colgate, and Roberta Klatzky. The spatial profile of skin indentation shapes tactile perception across stimulus frequencies. *Scientific reports*, 12(1):13185–13185, 2022.
- [27] M. Grunwald. *Human Haptic Perception: Basics and Applications*. Birkhäuser Basel, 2008.
- [28] Jianglong Guo, Jinsong Leng, and Jonathan Rossiter. Electro-adhesion technologies for robotics: A comprehensive review. *IEEE transactions on robotics*, 36(2):313–327, 2020.
- [29] Waseem Hassan, Arsen Abdulali, and Seokhee Jeon. Authoring new haptic textures based on interpolation of real textures in affective space. *IEEE transactions on industrial electronics (1982)*, 67(1):667–676, 2020.
- [30] Rania Hassen, Basak Guecyuez, and Eckehard G. Steinbach. Pvc-slp: Perceptual vibrotactile-signal compression based-on sparse linear prediction. *IEEE Transactions on Multimedia*, pages 1–1, 2020.
- [31] M Hollins, S.J Bensmaia, and E.A Roy. Vibrotaction and texture perception. *Behavioural Brain Research*, 135(1):51–56, 2002.

- [32] Mark Hollins, Sliman Bensmaia, and Sean Washburn. Vibrotactile adaptation impairs discrimination of fine, but not coarse, textures. *Somatosensory & Motor Research*, 18:253–62, 02 2001.
- [33] Mark Hollins and S. Risner. Evidence for the duplex theory of tactile texture perception. *Perception & Psychophysics*, 62:695–705, 2000.
- [34] R S Johansson and A B Vallbo. Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *The Journal of physiology*, 286(1):283–300, 1979.
- [35] Roland S. Johansson and Åke B. Vallbo. Tactile sensory coding in the glabrous skin of the human hand. *Trends in Neurosciences*, 6(1):27–32, 1983.
- [36] R.S. Johansson, U. Landström, and R. Lundström. Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements. *Brain Research*, 244(1):17–25, 1982.
- [37] Alfred Johnsen and Knud Lyne Rahbek. A physical phenomenon and its applications to telegraphy, telephony, etc. *Journal of the Institution of Electrical Engineers*, 61:713–725, 1923.
- [38] D. Katz and L.E. Krueger. *The World of Touch*. L. Erlbaum, 1989.
- [39] Pinar Kirci and Tara Ali-Yahiya. *6G for Tactile Internet*, chapter 4, pages 35–64. John Wiley & Sons, Ltd, 2021.
- [40] Qiang Li, Oliver Kroemer, Zhe Su, Filipe Fernandes Veiga, Mohsen Kaboli, and Helge Joachim Ritter. A review of tactile information: Perception and action through touch. *IEEE Transactions on Robotics*, 36(6):1619–1634, 2020.
- [41] J Lieberman and C Breazeal. Tikl: Development of a wearable vibrotactile feedback suit for improved human motor learning. *IEEE transactions on robotics*, 23(5):919–926, 2007.
- [42] Huaping. Liu. *Robotic Tactile Perception and Understanding : A Sparse Coding Method*. Springer Singapore, Singapore, 1st ed. 2018. edition, 2018.
- [43] Shihan Lu, Mianlun Zheng, Matthew C. Fontaine, Stefanos Nikolaidis, and Heather Marie Culbertson. Preference-driven texture modeling through interactive generation and search. *IEEE transactions on haptics*, PP:1–1, 2022.

- [44] D. A. Mahns, N. M. Perkins, V. Sahai, L. Robinson, and M. J. Rowe. Vibrotactile frequency discrimination in human hairy skin. *Journal of Neurophysiology*, 95(3):1442–1450, 2006.
- [45] Nicholas D Marchuk, J Edward Colgate, and Michael A Peshkin. Friction measurements on a large area tpad. In *2010 IEEE Haptics Symposium*, pages 317–320. IEEE, 2010.
- [46] Juan S. Martinez, Hong Z. Tan, and Roger W. Cholewiak. Psychophysical validation of interleaving narrowband tactile stimuli to achieve broadband effects. In *2021 IEEE World Haptics Conference (WHC)*, pages 709–714, July 2021.
- [47] Taylor Mason, Jeong-Hoi Koo, Jae-Ik Kim, Young-Min Kim, and Tae-Heon Yang. A feasibility study of a vibrotactile system based on electrostatic actuators for touch bar interfaces: Experimental evaluations. *Applied Sciences*, 11:7084, 07 2021.
- [48] J.H McDermott, A.J Oxenham, and E.P Simoncelli. Sound texture synthesis via filter statistics. In *2009 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 297–300. IEEE, 2009.
- [49] Josh H McDermott, Michael Schemitsch, and Eero P Simoncelli. Summary statistics in auditory perception. *Nature neuroscience*, 16(4):493–498, 2013.
- [50] Josh H. McDermott and Eero P. Simoncelli. Sound texture perception via statistics of the auditory periphery: Evidence from sound synthesis. *Neuron*, 71(5):926 – 940, 2011.
- [51] Richard McWalter and Josh H. McDermott. Adaptive and selective time averaging of auditory scenes. *Current biology*, 28(9):1405–1418.e10, 2018.
- [52] Richard McWalter and Josh H McDermott. Illusory sound texture reveals multi-second statistical completion in auditory scene analysis. *Nature communications*, 10(1):5096–18, 2019.
- [53] 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 *2015 IEEE World Haptics Conference (WHC)*, pages 125–130, June 2015.
- [54] David J. Meyer, Michael A. Peshkin, and J. Edward Colgate. Tactile paintbrush: A procedural method for generating spatial haptic texture. In Seungmoon Choi, Katherine J. Kuchenbecker, and Greg Gerling, editors, *IEEE Haptics Symposium 2016, HAPTICS 2016 - Proceedings*, IEEE Haptics Symposium, HAPTICS, pages 259–264, United States, April 2016. IEEE Computer Society.

- [55] David J. Meyer, Michaël Wiertlewski, Michael A. Peshkin, and J. Edward Colgate. Dynamics of ultrasonic and electrostatic friction modulation for rendering texture on haptic surfaces. In *2014 IEEE Haptics Symposium (HAPTICS)*, pages 63–67, 2014.
- [56] Rabih Moshourab, Henning Frenzel, Stefan Lechner, Julia Haseleu, Valérie Bégay, Damir Omerbašić, and Gary Lewin. Measurement of vibration detection threshold and tactile spatial acuity in human subjects. *Journal of visualized experiments : JoVE*, 2016, 09 2016.
- [57] Joe Mullenbach, Craig Shultz, Anne Marie Piper, Michael Peshkin, and J Colgate. Surface haptic interactions with a tpad tablet. In *Proceedings of the adjunct publication of the 26th annual ACM symposium on user interface software and technology, UIST '13 Adjunct*, pages 7–8. ACM, 2013.
- [58] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 01 1965.
- [59] Andreas Noll, Basak Gülecüyüz, Alexander Hofmann, and Ekehard Steinbach. A rate-scalable perceptual wavelet-based vibrotactile codec. In *2020 IEEE Haptics Symposium (HAPTICS)*, pages 854–859, 2020.
- [60] Andreas Noll, Lars Nockenberger, Basak Gülecüyüz, and Ekehard Steinbach. Vc-pwq: Vibrotactile signal compression based on perceptual wavelet quantization. In *2021 IEEE World Haptics Conference (WHC)*, pages 427–432, 2021.
- [61] S. Okamoto and Y. Yamada. Perceptual properties of vibrotactile material texture: Effects of amplitude changes and stimuli beneath detection thresholds. In *2010 IEEE/SICE International Symposium on System Integration*, pages 384–389, Dec 2010.
- [62] Shogo Okamoto and Yoji Yamada. Lossy data compression of vibrotactile material-like textures. *IEEE Transactions on Haptics*, 6(1):69–80, 2013.
- [63] Erik Pescara, Florian Dreschner, Karola Marky, Kai Kunze, and Michael Beigl. Gen-vibe: Exploration of interactive generation of personal vibrotactile patterns. In *Proceedings of the Augmented Humans International Conference, AHs '20*, pages 1–9. ACM, 2020.
- [64] Erik Pescara, Michael Hefenbrock, Thomas Rzepka, and Michael Beigl. Vibration personalization with evolutionary algorithms. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on pervasive and ubiquitous computing and wearable computers, UbiComp '18*, pages 436–439. ACM, 2018.

- [65] Ivan Poupyrev, Shigeaki Maruyama, and Jun Rekimoto. Ambient touch: designing tactile interfaces for handheld devices. In *Proceedings of the 15th annual ACM symposium on user interface software and technology*, UIST '02, pages 51–60. ACM, 2002.
- [66] G. Repovš and A. Baddeley. The multi-component model of working memory: Explorations in experimental cognitive psychology. *Neuroscience*, 139(1):5 – 21, 2006.
- [67] Thomas Sednaoui, Eric Vezzoli, Brygida Dzidek, Betty Lemaire-Semail, Cedrick Chappaz, and Michael Adams. Friction reduction through ultrasonic vibration part 2: Experimental evaluation of intermittent contact and squeeze film levitation. *IEEE Transactions on Haptics*, 10(2):208–216, 2017.
- [68] Aaron Raymond See, Jose Antonio G Choco, and Kohila Chandramohan. Touch, texture and haptic feedback: A review on how we feel the world around us. *Applied sciences*, 12(9):4686–, 2022.
- [69] Craig D. Shultz, Michael A. Peshkin, and J. Edward Colgate. Surface haptics via electroadhesion: Expanding electrovibration with johnsen and rahbek. In *2015 IEEE World Haptics Conference (WHC)*, pages 57–62, 2015.
- [70] M. Simsek, A. Aijaz, M. Dohler, J. Sachs, and G. Fettweis. 5g-enabled tactile internet. *IEEE Journal on Selected Areas in Communications*, 34(3):460–473, March 2016.
- [71] A. Spanias, T. Painter, and V. Atti. *Audio Signal Processing and Coding*, chapter 5, pages 113–144. John Wiley & Sons, Ltd, 2005.
- [72] A. Spanias, T. Painter, and V. Atti. *Audio Signal Processing and Coding*, chapter 3, pages 51–89. John Wiley & Sons, Ltd, 2005.
- [73] E. Steinbach, M. Strese, M. Eid, X. Liu, A. Bhardwaj, Q. Liu, M. Al-Ja’afreh, T. Mahmoodi, R. Hassen, A. El Saddik, and O. Holland. Haptic codecs for the tactile internet. *Proceedings of the IEEE*, 107(2):447–470, Feb 2019.
- [74] Robert M Strong and Donald E Troxel. An electrotactile display. *IEEE transactions on man-machine systems*, 11(1):72–79, 1970.
- [75] Diana Angelica Torres, Anis Kaci, Frederic Giraud, Christophe Giraud-Audine, Michel Amberg, Stephane Clenet, and Betty Lemaire-Semail. Pca model of fundamental acoustic finger force for out-of-plane ultrasonic vibration and its correlation with friction reduction. *IEEE transactions on haptics*, 14(3):551–563, 2021.

- [76] J van der Linden, E Schoonderwaldt, J Bird, and R Johnson. Musicjacket-combining motion capture and vibrotactile feedback to teach violin bowing. *IEEE transactions on instrumentation and measurement*, 60(1):104–113, 2011.
- [77] Eric Vezzoli, Wael Ben Messaoud, Michel Amberg, Frederic Giraud, Betty Lemaire-Semail, and Marie-Ange Bueno. Physical and perceptual independence of ultrasonic vibration and electrovibration for friction modulation. *IEEE transactions on haptics*, 8(2):235–239, 2015.
- [78] Eric Vezzoli, Zlatko Vidrih, Vincenzo Giamundo, Betty Lemaire-Semail, Frederic Giraud, Tomaz Rodic, Djordje Peric, and Michael Adams. Friction reduction through ultrasonic vibration part 1: Modelling intermittent contact. *IEEE Transactions on Haptics*, 10(2):196–207, 2017.
- [79] T Watanabe and S Fukui. A method for controlling tactile sensation of surface roughness using ultrasonic vibration. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 1, pages 1134–1139 vol.1. IEEE, 1995.
- [80] Alison Weber, Hannes Saal, Justin Lieber, Ju-Wen Cheng, Louise R Manfredi, John Dammann, and Sliman Bensmaia. Spatial and temporal codes mediate the tactile perception of natural textures. *Proceedings of the National Academy of Sciences of the United States of America*, 110, 09 2013.
- [81] Michael Wiertlewski and J. Edward Colgate. Power optimization of ultrasonic friction-modulation tactile interfaces. *IEEE transactions on haptics*, 8(1):43–53, 2015.
- [82] Michael Wiertlewski, Rebecca Fenton Friesen, and J. Edward Colgate. Partial squeeze film levitation modulates fingertip friction. *Proceedings of the National Academy of Sciences - PNAS*, 113(33):9210–9215, 2016.
- [83] Michaël Wiertlewski, Daniele Leonardis, David J. Meyer, Michael A. Peshkin, and J. Edward Colgate. A high-fidelity surface-haptic device for texture rendering on bare finger. In *Haptics: Neuroscience, Devices, Modeling, and Applications*, Lecture Notes in Computer Science, pages 241–248, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [84] L Winfield, J Glassmire, J.E Colgate, and M Peshkin. T-pad: Tactile pattern display through variable friction reduction. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, pages 421–426. IEEE, 2007.
- [85] Jung-Han Woo and Jeong-Guon Ih. Vibration rendering on a thin plate with actuator array at the periphery. *Journal of sound and vibration*, 349:150–162, 2015.

- [86] Heng Xu, Michael A Peshkin, and J. Edward Colgate. How the mechanical properties and thickness of glass affect tpad performance. *IEEE transactions on haptics*, 13(3):483–492, 2020.
- [87] Xun Liu, M. Dohler, T. Mahmoodi, and Hongbin Liu. Challenges and opportunities for designing tactile codecs from audio codecs. In *2017 European Conference on Networks and Communications (EuCNC)*, pages 1–5, June 2017.
- [88] Zhe Yuan, Bin Kang, Xin Wei, and Liang Zhou. Exploring the benefits of cross-modal coding. *IEEE transactions on circuits and systems for video technology*, 32(12):1–1, 2022.

APPENDIX A

Parameter List

The following tabulation collects the input parameters of the Single-Pitch Texel Rendering Algorithm. Documented parameter interactions are noted.

Parameter	Description	Notes
Central Pitch	Location of frequency distribution for fine features	Bounded on the low side by Texel length (although continuous Texel to Texel oscillation phase may allow violation of this bound).
Irregularity	Spread of frequency distribution for fine features	May have inverse perceptual mapping with Central Pitch: subjects report generating low-frequency content by increasing Irregularity.
Texel Length	Physical length of a single Texel	Texel length must fall below a perceptual fine/coarse texture boundary to limit spatial information being conveyed by fine features.
Coarse Texture Profile	Amplitude map used to draw coarse features	Resolution of coarse features is limited by Texel length.
Coarse Strength	Percentage of friction modulation amplitude used to draw coarse features	
Fine Strength	Percentage of available friction modulation amplitude used to draw fine features	Limited by coarse strength, as the maximum amplitude of fine features is equal to the Texel amplitude defined by the coarse texture profile.
Distribution Shape	Shape of frequency distribution for fine features	Lognormal distribution was used for all testing.
Quantization Level	Spacing between available fine feature frequency choices	Significant perceptual effect on Central Pitch and Irregularity due to quantization.