HAPTICS SYMPOSIUM '16, APRIL 8–11, 2016, PHILADELPHIA, PENNSYLVANIA, USA
978-1-5090-0903-9/16/$31.00 ©2016 IEEE

Tactile Paintbrush: A Procedural Method for Generating Spatial Haptic Texture

David J. Meyer, Michael A. Peshkin, and J. Edward Colgate

***Abstract*— In this work, we aim to represent tactile textures in such a way that a given texture may be "painted" onto a selected spatial region of a tactile display. We recorded a series of fingertip swipes across eleven textures and stored the data as spatial friction maps – friction as a function of position. We analyzed these maps with a space-frequency transform, and observed stochasticity in our physical measurements. We modeled the randomness in spectral magnitude across space with three distributions: Rayleigh, Rice, and Weibull. We analyzed the quality of parameterizations using goodness of model fit as well as consistency across multiple swipes of the same texture. We found that a two-parameter Weibull model best represented the measured data, and propose to use this model in the Tactile Paintbrush for applying virtual textures to spatial regions.**

I. Introduction

As a fingertip scans across a textured surface, mechanoreceptors in the skin and subcutaneous tissue encode the the mechanical interaction for the brain to understand what is being felt. Although the mechanical interactions between a surface and fingertip create vibrations that are high-dimensional in nature, the human brain is nevertheless able to distill the tactile information to categorize and identify textures. In order to create virtual haptic stimuli that predictably elicit these same mental categorizations, both actuation technology and a model that describes the mechanics of texture perception are needed.

With regard to actuation technology, our research has focused on the development of variable friction technology for several years. Variable friction technology is currently being developed in two forms: ultrasonic vibration [1], [2] and electrostatic friction [3], [4]. Both technologies can produce wide-bandwidth vibrations in a scanning fingertip [5], and have been used to produce virtual textures [6], [7]. Even though these widely capable texture displays have been developed, the presentation of complex, realistic textures remains primitive. In our recent work, we have begun to lay framework for modeling *spatial* haptic textures – those which modify surface friction as a function of touch position – in an effort to create tools for building rich tactile scenes for digital display. To build upon this work, we seek a deeper understanding of textured surfaces, and focus on statistical analysis of measured friction data from real-world surfaces to inform methods for generating virtual texture.

The authors are with the Department of Mechanical Engineering, Northwestern University, 2145 N Sheridan Rd, Evanston, IL 60208, USA E-mail: meyerdj@u.northwestern.edu; peshkin@northwestern.edu; colgate@northwestern.edu

II. Related Work

A. Vision & Audio

Perhaps the most in-depth study of texture modeling is in the field of computer graphics. Since the advent of digital display technology, rendering realistic imagery has been a topic of great interest. Although many graphics texturing methods focus on mapping two-dimensional images to three-dimensional objects (texture mapping) or imposing altered surface normals (bump mapping), these techniques require explicit data storage for rendering. Because textures of this nature are explicitly defined, they are more likely to exhibit noticeable faults at boundaries of textured regions, for example in cases in which a textured region is larger than a source texture and tiling is apparent.

To address this issue, a variety of methods for generating synthetic textures have been developed, falling under the category of procedural texturing [8]. Procedural texturing models computationally generate texture from a relatively small number of parameters, and thus data storage is more compact and textures can be generated without discontinuities over large areas. These types of models typically render textures by modeling the randomness of texture found in nature. Research in statistical characterization of visual texture was initiated by the work of Julesz, who originally proposed that if two graphical scenes had the same second-order statistics of light intensity, they were indistinguishable [9]. More recent texture synthesis work has focused on the statistics of images processed by linear filters [10].

In the field of computational audio, texture is not such a widely adopted term, but the problem remains the same: digitally synthesizing realistic sounds is a challenge. In this area, similar methods of statistical modeling have emerged, including spectral filtering of noise signals [11] and controlling the joint statistics of the outputs of cochlea-inspired linear filters [12].

B. Haptic Textures

Texturing in haptics research falls short of that in vision and audition because haptic actuators have yet to be developed to the quality of visual displays and speakers. Minsky's Sandpaper system for generating haptic texture was one of the first examples of algorithmic haptic texturing, and is very similar to graphics bump mapping in that surface normals are used to compute haptic forces to be rendered with robotic linkages [13]. Because robotic linkages are difficult to control at the wide frequency bandwidths necessary for texture reproduction, more recent works in haptic textures have focused on vibrotactile actuation [14], [15].

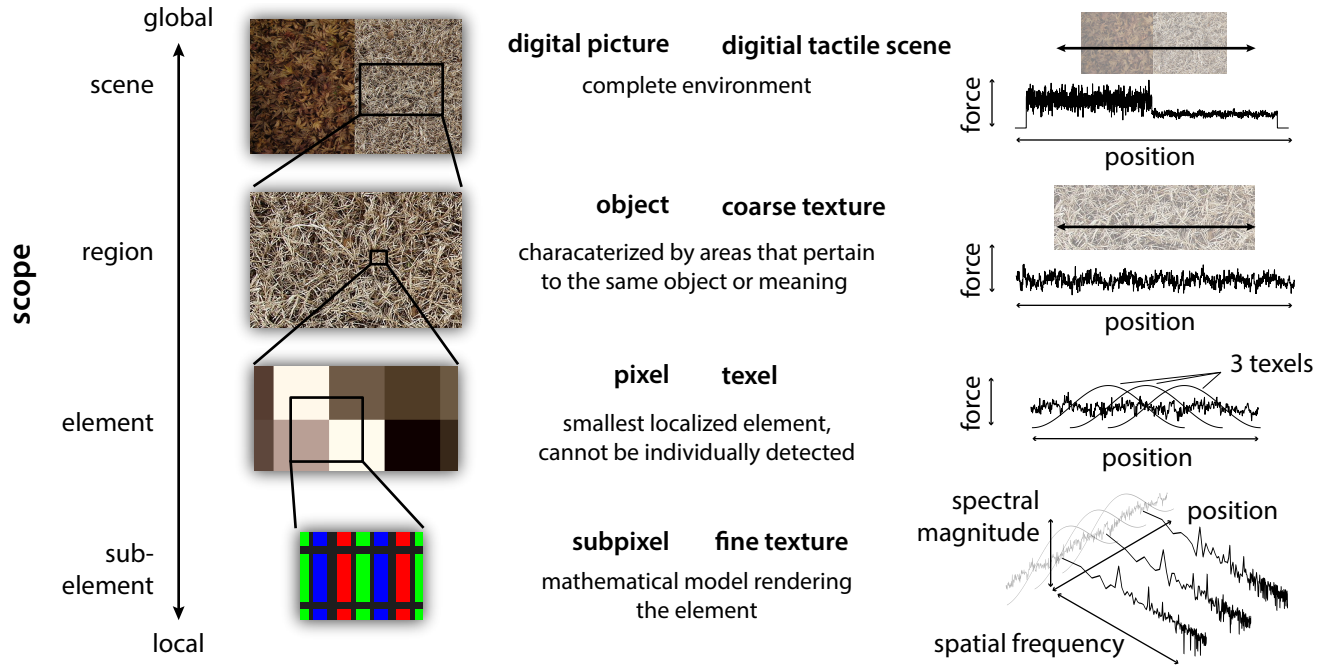

Fig. 1: Depiction of the space-frequency framework for tactile texture rendering and its analogue in the digital image space.

As in graphics texturing, procedural methods for synthetic texture generation have emerged. Early work in this area noted that surface height profiles often were Gaussian in nature, and synthesized haptic texture from Gaussian white noise parameterized by a mean and variance only [16]. Gaussian mixture modeling has also been used to generate haptic textures [17]. Studies of tactile perception have since shed light on the importance of spectral modeling for haptic textures [18]. Texture synthesis has shifted towards spectral modeling of measurements of real-world textured surfaces for informing the design of filters for generating texture from noise sources [19], [20].

III. Methods

In this paper, we present a new procedural texturing method for creating haptic textures that are designed spatially. Because the textures exist spatially, our texture models do not take into account scanning velocity. This method was developed for display of tactile texture on variable friction displays, and relies on the core principle of a space-frequency transform for representation. We call it the Tactile Paintbrush because the texture model is one that can be applied to a specific location, and from a designer's perspective it allows for spectrally designed textures to be "painted" onto a virtual tactile canvas.

A. Space-Frequency Framework

When a fingertip scans across a surface, small surface features induce high-frequency vibrations that are mediated by rapidly adapting mechanoreceptors. At the same time, larger features cause slowly-varying changes in these vibrations that are perceived in time and located by the proprioceptive and visual systems. A space-frequency analysis is used to represent data in two dimensions simultaneously: space and spatial frequency. The high-frequency vibrations are represented by spatial frequency spectra, and the slowly varying changes caused by larger surface features are represented explicitly in space. In effect, the spatial haptic texture is stored as a series of Fourier magnitude spectra that vary across the virtual surface. In our work, each Fourier analysis is conducted within a window of 1mm, and the texture is sampled in space with a resolution of 0.25mm, so that each analysis window overlaps with three of its neighbors. Synthesis of a purely spatial map for rendering is computed via an iterative overlap-add method [21].

An illustration of our framework for describing digital tactile scenes is shown in Figure 1. Each of the sampled points in space, characterized by its local Fourier spectrum, is a fundamental unit of texture, what we call a texel. In previous work, we have shown that a texel of about 0.25mm is the smallest unit that is uniquely detected on a tactile surface [21]. A texel is very similar to a pixel in an image; a tactile scene is made up of an array of texels, each of which has its own fine texture, and an image is made up of a large array of pixels, each of which has its own color. Although true color is a complex combination of wavelengths of light entering the retina, an RGB color model takes advantage of human perception and replicates color space with a simple three parameter representation. Each texel is a complex combination of wavelengths of spatial features, but may in fact only need to be described by a smaller subset of parameters. However, we do not investigate this possibility in this paper.

We focus on studying the distributions of texel information

TABLE I: Descriptions of the eleven scanned textures

1	plastic 2mm scales	7	leather
2	plastic 2mm square grating	8	cardboard
3	patterned paper	9	wood veneer
4	hard fine-textured plastic	10	carpet
5	rubbery vinyl	11	burlap
6	machined nylon		

across space for uniform textures, those which feel perceptually the same across space. It is clear that the explicit friction profile of a uniform texture will change across space, i.e. not all texels will be exactly the same. In fact, since the texel stores a local spatial signal, if all texels were identical, the texture would exhibit 0.25mm periodicity. Therefore, a stochastic representation of texel data is crucial for creating realistic textures. A surface that is uniform should consist of texels that do not change their underlying statistics.

B. Stochasticity in Uniform Textures

In order to assess the stochasticity of real textures, we measured the friction coefficient due to sliding across eleven physical uniform fine textures, described in Table I. A diagram of our experimental setup is shown in figure 2. We mounted these textures to a lightweight aluminum plate instrumented with a high-bandwidth piezo force sensor (Kistler Type 9203) for measuring lateral force, and two strain gauge load cells for measuring normal force. A 3D-printed plastic piece clamped to the top of a sliding fingertip moved a rotary encoder as the finger scanned across the surface of the texture. The resolution of the encoder was $3.86\,\mu$m, and a data acquisition system collected position and lateral force data at a rate of 125 kHz. The data was filtered with non-causal, zero-phase, second-order Butterworth filters. The lateral force cutoff frequency was 2 kHz, and because the normal force load cells are not high-bandwidth, the normal force cutoff was 100 Hz. The normal force remained around 1 N for the duration of the experiment, and the friction coefficient was calculated as the lateral force divided by the normal force.

We recorded ten cycles of back-and-forth scanning along each texture. The velocity was kept roughly constant using a metronome for timing the cycles. After recording both friction coefficient and finger position in time, the friction signal was re-sampled and interpolated to produce a friction vs position signal, with a resolution of 5 microns. This spatial friction map was truncated to include only the constant-velocity sections of the swipe. We analyzed this signal with a space-frequency transform to represent the data as a series of texels spaced 0.25 mm apart. Each texel has 101 parameters, one for each wavelength and corresponding wavenumber, representing the magnitude of the spectrum within the texel window, which is 1mm wide. The zeroth wavenumber coefficient corresponds to the average friction within the texel, the first wavenumber has a wavelength of 1 mm, and the 101$^{\text{st}}$ wavenumber has a wavelength of 10 microns. We analyzed the rightward-moving swipes on eleven textures, for a total of 110 swipes.

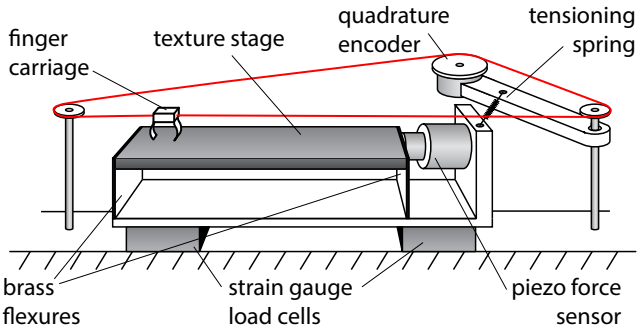

Fig. 2: Diagram of the tribometer used to measure friction coefficients on a fingertip scanning across textures. The encoder tracks the position of the finger, the piezo force sensor records the lateral force at each position, and the strain gauge load cells record the normal force of contact.

C. Fitting Statistical Models

To assess the spatial stocasticity of the textures, we generated histograms of the texel values for each wavelength, k, over all texels for all scans of the fingertip. Each swipe yields 101 histograms. With 110 swipes, we have a total of over eleven thousand histograms. Each histogram is a distribution of the spectral magnitudes resulting from an FFT analysis inside every texel in a swipe. To model the randomness in the spectral magnitudes, we tested three non-negative distributions on each histogram: Rayleigh, Rice, and Weibull.

1) Rayleigh Distribution: The Rayleigh distribution is derived from the magnitude of a vector generated using two zero-mean Gaussian distributions with equal variance. If a spatial friction map was modeled as pure Gaussian noise, the spectral magnitudes would be distributed according to a Rayleigh distribution. The probability density function is a one-parameter model described by equation 1 in which the parameter σ models the variance of Gaussian distributions in the real-imaginary plane. This model assumes that the Fourier coefficients of a given wavelength are distributed about the origin of the real-imaginary plane.

$$f(x;\sigma) = \frac{x}{\sigma^2} e^{-x^2/2\sigma^2}, \quad x \geq 0 \tag{1}$$

When modeled with a Rayleigh distribution, each texture swipe is reduced to a 101-dimensional set, one number for each wavenumber, that represents the scale of the Rayleigh distribution across space.

2) Rice Distribution: The Rice distribution is derived from the magnitude of a vector generated using two Gaussian distributions that are perhaps not centered at zero. This model extends the Rayleigh model, and accounts for surfaces that may have deterministic spectral components in addition to Gaussian noise. The probability density function is a two-parameter model given in equation 2 in which ν is the distance of the mean from the origin of the real-imaginary

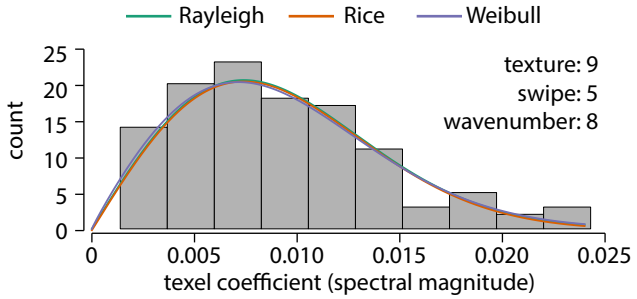

(a) Typical distribtion of texel data for one swipe and wavenumber.

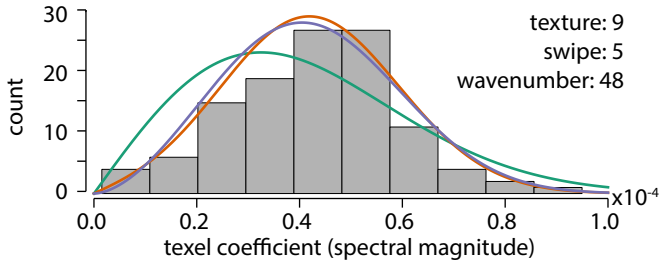

(b) An example distribution that shows evidence of complex Fourier coefficients that were non-zero mean.

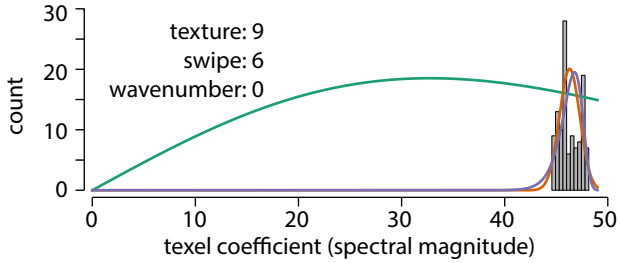

(c) The zeroth wavenumber is the average friction in the texel, which is clearly non-zero mean. A Rayleigh model is not suitable for these data.

Fig. 3: Selected histograms of texel data with best fit model shapes plotted for reference

plane, and σ is the variance of the underlying Gaussian distributions. I_0 is the modified Bessel function of the first kind with order zero.

$$f(x;\nu,\sigma) = \frac{x}{\sigma^2}\exp\left(\frac{-(x^2+\nu^2)}{2\sigma^2}\right) I_0\left(\frac{x\nu}{\sigma^2}\right), \quad x \geq 0 \quad (2)$$

When modeled with a Rice distribution, each texture swipe is reduced to a 202-dimensional set, two numbers for each wavenumber, that represents the scale and location of the Rice distribution across space.

3) Weibull Distribution: The Weibull is another generalization of the Rayleigh distribution, with an additional parameter that allows for shape. It offers similar flexibility to the Rice distribution, but is much simpler mathematically. The probability distribution function is given in equation 3, where λ is the scale factor, a rough estimate of the size of the parameters. In addition to the scale parameter, a shape parameter, k, interpolates the shape between an exponential at $k = 1$, Rayleigh at $k = 2$, and Dirac delta function at $k = \infty$.

For $k = 2$, the distribution is exactly the Rayleigh distribution where $\lambda = \sqrt{2}\sigma$.

$$f(x;\lambda,k) = \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k}, \quad x \geq 0 \quad (3)$$

When modeled with a Weibull distribution, each texture swipe is reduced to a 202-dimensional set, two numbers for each wavenumber, that represents the scale and location of the Weibull distribution across space.

We calculated the maximum likelihood estimates for all three distributions for each of the 11,110 texel distributions in our dataset. Figure 3 shows a few example histograms and their distributions. For the most part, all three distributions approximate the data with reasonable accuracy, as in the example given in Figure 3a. However, some of the data exhibits non-zero mean, and in these cases, a Rayleigh model does not fit well, as shown in Figure 3b. The top two wavelengths are both heavily influenced by the overall average texture in the texel, and therefore have significant non-zero components in their Fourier coefficients. An example of this data is shown in Figure 3c, and clearly a Rayleigh model is inappropriate.

D. Model Validation

To test the validity of these parameterizations, we calculated the goodness of fit for each histogram and distribution. We used a Kolmogorov-Smirnov test (K-S test) on each fit to determine the maximum distance between the fitted cumulative distribution function (CDF) and the empirical CDF. The result of this test is called the K-S statistic; a better fit gives a lower number, and a perfect fit results in a K-S statistic that approaches 0. The results are shown in Figure 4. Looking at the image on the left, it is clear that the first two coefficients in each texel are poorly fit with a Rayleigh distribution. Overall, comparing the three distributions, we observe that the Weibull distribution fits the data most accurately.

We also wanted to test the ability of these distributions to capture the important textural parameters. We assume that if a model is capturing a texture perfectly, subsequent swipes of the texture yield identical measures. For each surface, the ten swipes we recorded should result in representations that are clustered together in the parameter space. Because the similarities may depend on the type of distance metric used, we tried four different distance metrics for evaluating the similarity of the textures: Euclidean, standardized Euclidean, Canberra, and Bray-Curtis. The latter two metrics are both normalized metrics that have been shown to perform well in image texture similarity ratings [22].

With each distance metric, we measured the quality of texture swipe clustering with a silhouette coefficient. The coefficient is calculated for each point i by equation 4, in which $a(i)$ is the average distance to each other point in the same cluster (swipes from the same texture), and $b(i)$ is the average distance to each point in the nearest neighboring cluster (swipes from a different texture) [23].

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}} \quad (4)$$

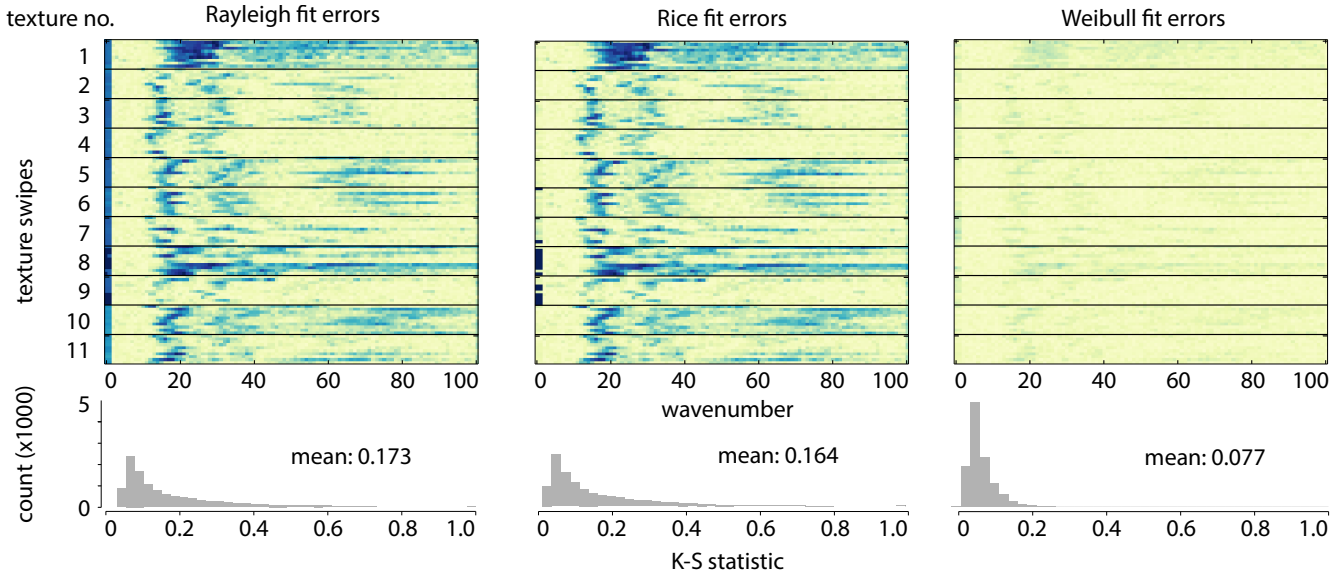

Fig. 4: Images (top) of the K-S statistics for each histogram fit across all swipes and wavenumbers. Dark blue corresponds to a high K-S statistic, meaning large error in the fit. Light yellow means a low K-S statistic and good fit. The histograms of K-S statistics for each distribution model are shown below their respective image representation.

The silhouette coefficient ranges from -1 to 1, higher values indicate better clustering. A point that is near the edge of two clusters gives a score around 0, and points that have a negative are closer to neighboring clusters than their own. To assess the clustering of our texture swipes, we use the silhouette score: an average of the coefficients for all data points in a set. We computed this score for parameter spaces computed by all three statistical fits, using the four distance metrics stated above. The results of these clusterings are shown in Table II.

To aid in visualizing clustering, we computed MDS analyses for the parameter spaces using the same four distance metrics to create a dissimilarity matrix. Four of the twelve MDS results are shown in Figure 5. We also place the average of all swipes for each texture in the MDS space. The Weibull parameterization is the only model that produces a positive silhouette score for every distance metric we tested.

IV. Discussion

It is evident that multiple swipes of a single texture will yield different patterns of friction, even though perceptually they are indistinguishable. The goal of texture modeling is to capture the important properties of the texture in a compact parameter space. The distances between textures in this parameter space should be measured by a distance metric that correlates with human perceptual dissimilarity of textures. In future modeling work, the perceptual dissimilarity should be tested to additionally validate mathematical distance metrics.

TABLE II: Silhouette coefficients for three distributions and four distance metrics. Values in bold correspond to MDS plots in Figure 5.

	Rayleigh	Rice	Weibull
Euclidean	**-0.074**	-0.157	**0.098**
standardized Euclidean	0.107	**0.009**	0.033
Canberra	0.069	0.017	0.055
Bray-Curtis	-0.067	-0.065	**0.154**

The Tactile Paintbrush provides a method for creating stochastic friction patterns by drawing samples from a distribution for each fine-texture wavelength. In this work, each texture is represented by a 202-parameter model, or "brush." A Weibull model is particularly well-suited for practical use in this case because it is computationally simple to generate Weibull-distributed random variables and because the two-parameter model is intuitive. By selecting the shape of the distribution curve to be $k = 2$, a designer is deciding that any power in that spectrum comes from a zero-mean distribution. However, if a designer decides that a specific wavelength should have more definitive non-zero content, a higher value of $k \to \infty$ can be chosen to reduce the randomness of the texture at a given wavelength. By modifying the scale parameter, λ, in the distribution model, the rough spectral shape of the texture can be designed. Once a texture brush has been designed, it is painted onto a virtual canvas by populating the space-frequency transformed texture with random samples drawn from the designed distributions.

V. Conclusion

We measured the friction patterns of eleven textures, and analyzed them using a space-frequency transform. We observed the distributions of the local Fourier coefficients in space, and modeled them with three different distributions. We found that a Weibull distribution was the most suitable for both fitting measured data and categorizing textures. We presented the Tactile Paintbrush, a method for generating stochastic haptic textures from a Weibull distribution.

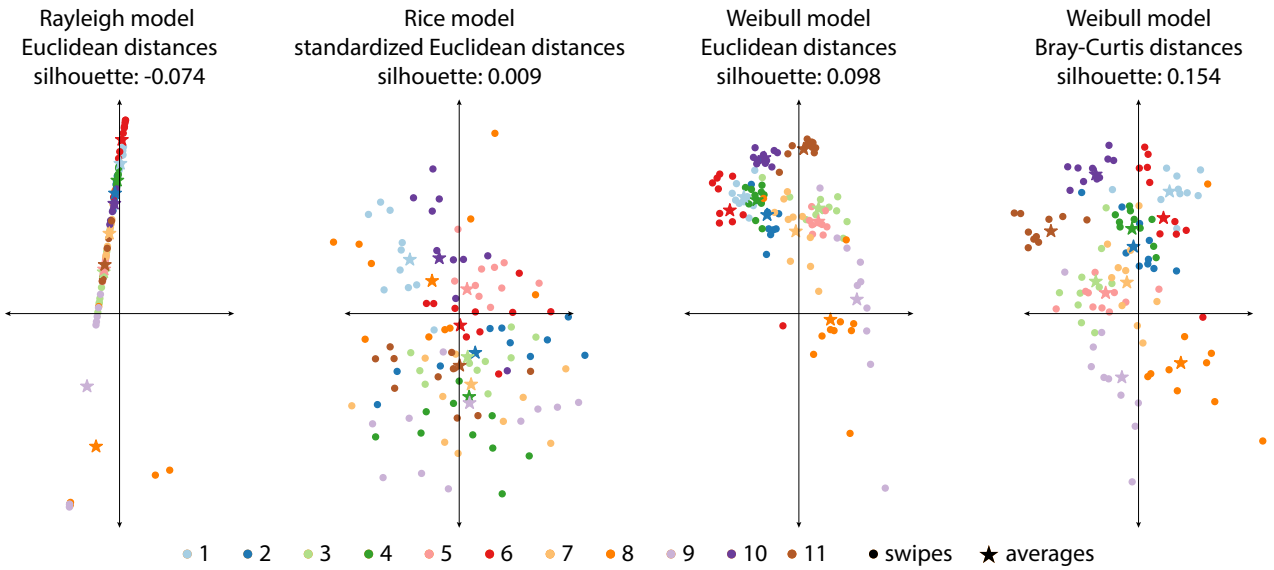

Fig. 5: Multi-dimensional scaling analyses for four different combinations of parameterization and distance metric. Each color represents all swipes from the same texture, which should be ideally clustered. The Weibull model most stably produces well-clustered texture representations.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1518602.

REFERENCES

[1] M. Biet, F. Giraud, and B. Lemaire-Semail, "Squeeze film effect for the design of an ultrasonic tactile plate," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 54, no. 12, pp. 2678–2688, 2007.

[2] L. Winfield, J. Glassmire, J. E. Colgate, and M. Peshkin, "T-PaD: Tactile Pattern Display through Variable Friction Reduction," in *EuroHaptics Conference, 2007 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2007. Second Joint*, Mar. 2007, pp. 421 –426.

[3] O. Bau, I. Poupyrev, A. Israr, and C. Harrison, "TeslaTouch: electrovibration for touch surfaces," in *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, ser. UIST '10. New York, NY, USA: ACM, 2010, pp. 283–292.

[4] Z. Radivojevic, P. Beecher, C. Bower, S. Haque, P. Andrew, T. Hasan, F. Bonaccorso, A. C. Ferrari, and B. Henson, "Electrotactile touch surface by using transparent graphene," in *Proceedings of the 2012 Virtual Reality International Conference*, ser. VRIC '12. New York, NY, USA: ACM, 2012, pp. 16:1–16:3.

[5] D. Meyer, M. Wiertlewski, M. Peshkin, and J. Colgate, "Dynamics of ultrasonic and electrostatic friction modulation for rendering texture on haptic surfaces," in *2014 IEEE Haptics Symposium (HAPTICS)*, Feb. 2014, pp. 63–67.

[6] M. Wiertlewski, D. Leonardis, D. J. Meyer, M. A. Peshkin, and J. E. Colgate, "A High-Fidelity Surface-Haptic Device for Texture Rendering on Bare Finger," in *Haptics: Neuroscience, Devices, Modeling, and Applications*, ser. Lecture Notes in Computer Science, M. Auvray and C. Duriez, Eds. Springer Berlin Heidelberg, June 2014, pp. 241–248.

[7] S.-C. Kim, A. Israr, and I. Poupyrev, "Tactile Rendering of 3d Features on Touch Surfaces," in *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '13. New York, NY, USA: ACM, 2013, pp. 531–538.

[8] D. S. Ebert, *Texturing & Modeling: A Procedural Approach*. Morgan Kaufmann, 2003.

[9] B. Julesz, "Visual Pattern Discrimination," *IRE Transactions on Information Theory*, vol. 8, no. 2, pp. 84–92, Feb. 1962.

[10] J. Portilla and E. P. Simoncelli, "A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients," *International Journal of Computer Vision*, vol. 40, no. 1, pp. 49–70, Oct. 2000.

[11] X. Serra, "Spectral Modeling Synthesis: A Sound Analysis/Synthesis Based on a Deterministic plus Stochastic Decomposition," *Computer Music Journal*, vol. 14, pp. 12–24, 1990.

[12] J. H. McDermott and E. P. Simoncelli, "Sound Texture Perception via Statistics of the Auditory Periphery: Evidence from Sound Synthesis," *Neuron*, vol. 71, no. 5, pp. 926–940, Sept. 2011.

[13] M. Minsky, O.-y. Ming, O. Steele, F. P. Brooks, Jr., and M. Behensky, "Feeling and Seeing: Issues in Force Display," in *Proceedings of the 1990 Symposium on Interactive 3D Graphics*, ser. I3D '90. New York, NY, USA: ACM, 1990, pp. 235–241.

[14] S. Okamoto, M. Konyo, S. Saga, and S. Tadokoro, "Detectability and Perceptual Consequences of Delayed Feedback in a Vibrotactile Texture Display," *IEEE Transactions on Haptics*, vol. 2, no. 2, pp. 73–84, Apr. 2009.

[15] H. Culbertson, J. M. Romano, P. Castillo, M. Mintz, and K. J. Kuchenbecker, "Refined methods for creating realistic haptic virtual textures from tool-mediated contact acceleration data," in *Haptics Symposium (HAPTICS), 2012 IEEE*, 2012, pp. 385–391.

[16] J. Siira and D. Pai, "Haptic texturing-a stochastic approach," in *, 1996 IEEE International Conference on Robotics and Automation, 1996. Proceedings*, vol. 1, Apr. 1996, pp. 557–562 vol.1.

[17] J. P. Fritz and K. E. Barner, "Stochastic models for haptic texture," vol. 2901, 1996, pp. 34–44.

[18] S. Bensmaïa and M. Hollins, "The vibrations of texture," *Somatosensory & motor research*, vol. 20, no. 1, pp. 33–43, 2003.

[19] J. Romano, T. Yoshioka, and K. Kuchenbecker, "Automatic filter design for synthesis of haptic textures from recorded acceleration data," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*, May 2010, pp. 1815–1821.

[20] H. Culbertson, J. Unwin, B. Goodman, and K. Kuchenbecker, "Generating haptic texture models from unconstrained tool-surface interactions," in *World Haptics Conference (WHC), 2013*, Apr. 2013, pp. 295–300.

[21] D. Meyer, M. Peshkin, and J. Colgate, "Modeling and synthesis of tactile texture with spatial spectrograms for display on variable friction surfaces," in *2015 IEEE World Haptics Conference (WHC)*, June 2015, pp. 125–130.

[22] M. Kokare, B. Chatterji, and P. Biswas, "Comparison of similarity metrics for texture image retrieval," in *Conference on Convergent Technologies for the Asia-Pacific Region TENCON 2003*, vol. 2, Oct. 2003, pp. 571–575 Vol.2.

[23] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, Nov. 1987.