

A Quantitative Perceptual Model for Tactile Roughness

CHELSEA TYMMS, New York University, USA

ESTHER P. GARDNER, New York University School of Medicine, USA

DENIS ZORIN, New York University, USA

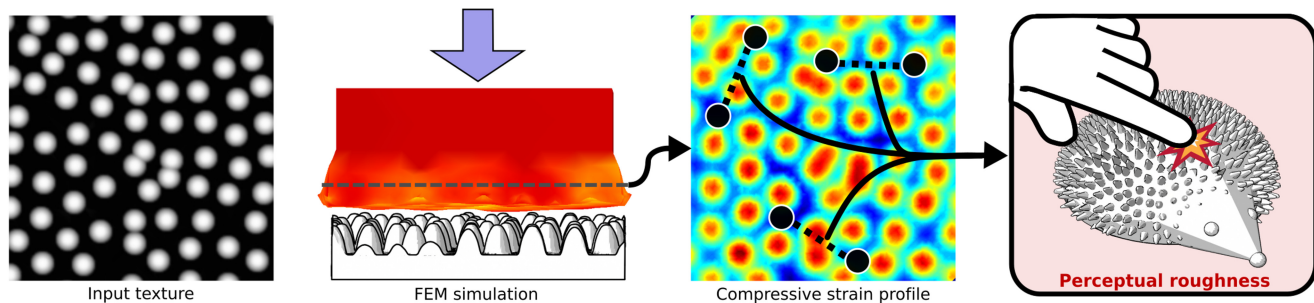


Fig. 1. Given an input surface geometry, our model computes the tactile roughness by computing skin deformation using elasticity simulations and sampling the resulting compressive strain field. Our model enables the fabrication of tactile objects with specified roughnesses.

Everyone uses the sense of touch to explore the world, and roughness is one of the most important qualities in tactile perception. Roughness is a major identifier for judgments of material composition, comfort, and friction, and it is tied closely to manual dexterity. The advent of high-resolution 3D printing technology provides the ability to fabricate arbitrary 3D textures with surface geometry that confers haptic properties. In this work, we address the problem of mapping object geometry to tactile roughness. We fabricate a set of carefully designed stimuli and use them in experiments with human subjects to build a perceptual space for roughness. We then match this space to a quantitative model obtained from strain fields derived from elasticity simulations of the human skin contacting the texture geometry, drawing from past research in neuroscience and psychophysics. We demonstrate how this model can be applied to predict and alter surface roughness, and we show several applications in the context of fabrication.

CCS Concepts: • **Human-centered computing** → *User studies*; • **Computing methodologies** → *Perception*;

Additional Key Words and Phrases: Roughness, fabrication, perception

ACM Reference format:

Chelsea Tymms, Esther P. Gardner, and Denis Zorin. 2018. A Quantitative Perceptual Model for Tactile Roughness. *ACM Trans. Graph.* 37, 5, Article 168 (October 2018), 14 pages.

<https://doi.org/10.1145/3186267>

This work was partially supported by the Rudin Foundation and by NSF awards IIS-1320635 and DMS-1436591.

Authors' addresses: C. Tymms, New York University, New York, NY, 10011; email: tymms@cs.nyu.edu; E. P. Gardner, New York University School of Medicine, New York, NY, 10016; email: esther.gardner@nyumc.org; D. Zorin, New York University, New York, NY, 10011; email: dzorin@cs.nyu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

0730-0301/2018/10-ART168 \$15.00

<https://doi.org/10.1145/3186267>

1 INTRODUCTION

Tactile perception consists of four “coordinates”: roughness, compliance, stickiness, and thermal sensitivity; of these, roughness is known to be the most significant (Tiest 2010). Tactile roughness is a major identifier for judging material composition; it informs the perception of important properties such as comfort and friction, and is essential for manual dexterity (Johansson and Westling 1984).

Varying surface roughness is common in conventional technologies, such as injection molding for plastic and, in a more restricted form, in *knurling* in CNC machining. Additive fabrication technologies drastically increase the flexibility of roughness control; on the other hand, as the range of materials used by current additive technologies is more limited, roughness and other tactile properties often must be controlled with geometry variation, rather than with the choice of material.

The relationship between surface geometry, material properties, and tactile perception is intricate, which makes the task of producing and designing tactile properties difficult. For many aspects of visual and auditory perception, a variety of models (e.g., perceptual color models and equal-loudness curves) are available, and these are broadly used for applications such as rendering and sound processing and synthesis. Our level of understanding of tactile perception is far less mature. Although many studies have been done over time, a widely accepted model is still absent: progress has been slow due to the complexity of the tactile system and the difficulty of producing controlled stimuli. An important goal would be to develop a model of tactile sensation that can be used for tactile “rendering,” i.e., producing objects with controlled tactile feel. Important steps in this direction were made recently in the context of additive fabrication applications by Piovarčič et al. (2016), who developed a model for compliance perception that can be used to control one dimension of tactile perception of 3D printed objects.

In this article, we focus on another dimension of tactile perception: roughness. Digital fabrication makes it more practical to control roughness directly and at the same time provides a way to study perception of roughness by producing diverse fine-scale stimuli for studies.

Generally, a rough surface is a surface that causes uneven pressure on the skin when touched statically and elicits vibrations when stroked (Tiest 2010). Human skin contains four different types of mechanoreceptors that mediate tactile sensation of spatial distributions, vibration, and skin stretch, and this low-level processing is followed by higher-level processing. Nevertheless, it is generally accepted (Hollins and Risner 2000) that roughness sensation is well described by at most two dimensions, one of which can be described as roughness proper (large-scale roughness) and the other as vibrational roughness corresponding to very small-scale effects.

In this article, we propose a computational model relating surface geometry to perceived (large-scale) roughness based on strain fluctuation variation in an elasticity model. We use human perception studies to compute a perceptual space of roughness by applying multidimensional scaling, and we demonstrate that our model matches experimental results significantly better than previously proposed models. We show how this model can be applied in the context 3D printing to achieve a range of effects.

As a part of our experiment design, we developed an approach to producing a range of controlled tactile stimuli, which is likely to be of value for future studies of roughness. The 3D stimuli models, the code for geometry generation, and the raw data of the experiments will be released to encourage further exploration of this topic.

2 RELATED WORK

Our work builds on sources in several areas, including psychophysical literature and work in computer graphics related to perceptually based rendering and digital fabrication.

2.1 Tactile Fabrication

Recently, several works have explored the fabrication of objects with tactile properties. Our work is closest to Piovarči et al. (2016) in aims and in the methods we use to construct and evaluate our model. This work develops a model for perceptual *compliance* using stimuli fabricated from materials with different compliance properties. In our work, we use similar experimental protocols and multidimensional scaling analysis techniques to find a perceptual space, but we develop a model for roughness. Compared to compliance, roughness is primarily related to small-scale surface geometry instead of overall bulk material properties; therefore, a much larger physical space of possible surface profiles is mapped to a single perceptual dimension.

Another work with similar aims is Elkharraz et al. (2014), which fabricated texture plates from a set of *visual* textures converted to shallow height maps, asked human subjects to rate the resulting textures according to a set of adjectives (including “rough”), and analyzed a set of computational texture features to find those highly correlated with the perceptual scales. Although this article reports success with prediction of tactile roughness, the dataset is quite limited compared to the dimension of the feature space. In




Fig. 2. Diagram of the structure of human skin (CNX 2017). Roughness perception is primarily mediated by Merkel cells, or SA1 receptors, located approximately 0.75mm deep in the skin.

Section 5, we compare this model to ours and demonstrate a significant improvement in performance.

Other recent work in the fabrication domain has aimed to facilitate the incorporation of tactile properties in 3D printed models. Torres et al. (2015) provide an interface to fabricate objects with a user-specified weight, compliant infill, and rough displacement map. However, their roughness metric relies on texture feature size, which is not always definable and does not describe a comprehensive model for all textures. Chen et al. (2013) also develop methods to fabricate objects with specified deformation behavior and textured surface displacement, but they similarly do not extend to tactual perceptual specifications. Our model offers an improvement to surface roughness specification in these types of interfaces. Relatedly, Lau et al. (2016) developed a method to analyze meshes for tactile saliency, identifying mesh regions most likely to be touched. Incorporating tactile saliency with our work could create a natural tool to suggest the application of appropriate textures on salient regions of objects prior to fabrication.

2.2 Psychophysics of Roughness Perception

Perception. Research in the sense of touch, while extensive, is still limited when compared to vision and hearing. Various works have aimed to quantify different aspects of the sense of touch. Research suggests that tactile perception falls in a multidimensional space varying along three (Hollins et al. 1993) or four to five (Tiest 2010) dimensions; of these, roughness has been found as the most significant dimension. Roughness can further be divided into two aspects: small-scale roughness caused by features less than 0.1mm in size, which elicit vibrational cues perceived by Pacinian corpuscles deep in the skin, and large-scale roughness, which is mediated primarily by Merkel cells at the base of the epidermis that respond to strain (a third mechanoreceptor, the Meissner’s corpuscle, also plays a part in lower-frequency vibration detection) (Figure 2). The remaining three dimensions are compliance, stickiness, and surface warmth or heat conductivity.

A major issue that has prevented the comprehensive study of tactile perception, and specifically roughness perception, is the lack of suitable stimuli: in comparison to the generation of visual or auditory stimuli, tactile stimuli are relatively difficult to manufacture with the desired accuracy. Most early studies in tactile perception used different types of common manufactured surfaces, such

as sandpaper (Hollins and Risner 2000), paper, and fabric (Chen et al. 2009; Manfredi et al. 2014; Weber et al. 2013), and other everyday items. However, these surfaces are not ideal as stimuli because they vary across multiple dimensions, and their surface features cannot be easily quantified; even graded sandpaper varies widely by manufacturer.

To overcome these problems, some studies in roughness perception have manufactured artificial stimuli. The two common types of custom stimuli are square gratings, which are manufactured using machine engraving (Lederman 1974; Yoshioka et al. 2001), and photo-embossed dots. The studies suggest that perceptual roughness has a positive correlation with groove width or spacing and a slightly negative correlation with ridge or dot width. Yoshioka et al. (2001) and Connor et al. (1990) developed neural models using these stimuli; they found that the mechanoreceptor firing rates elicited by touching these surfaces were correlated to perceived roughness via firing rate *spatial variation*. The spatial variation hypothesis was also explored recently by Goodman and Bensaïa (2017), who showed that it could predict perceptual roughness for several raised-dot stimuli used in past experiments. We use this hypothesis as a base for our model in Section 5.

The artificial stimuli used in these past experiments have significant limitations. While engraved gratings are precise, they present a limited and unnatural stimulus. The photo-embossing process most often used for dot stimuli also has several limitations; it effectively supports only two levels of height (unblocked areas are solidified, and blocked areas are not) with little control over the intermediate areas, and the material is quite flexible. In contrast, recent developments in fabrication, such as the stereolithography 3D printing we use to manufacture our stimuli, allow the creation of solid stimuli with any surface geometry at a high resolution, covering the complete range for (large-scale) roughness. In fact, to our knowledge, no psychophysics study to date has used stimuli with raised dots as small or finely spaced as those used here. Additionally, no study to our knowledge has tested the roughness of textured curved surfaces.

Physical Modeling. Research in perception has also worked to model the lower-level mechanoreceptor responses in humans and related primates. Several works have aimed to model the skin mechanics and receptor responses: Phillips and Johnson (1981) created a continuum mechanics model for the skin contacting bars and gratings, and Dandekar et al. (2003) developed a 3D FEM model to accurately match skin displacement under line loads and predict Merkel cell responses to these indentations. Both studies found that Merkel cells have a firing rate closely matched to the maximum compressive strain elicited in the skin. Indeed, recent strides in molecular biology (Woo et al. 2014) have shown that Merkel cells express the mechanically activated ion channel Piezo2 as a means to mediate mechanosensation. This relation between strain and firing rates is used as a basis for our computational model in Section 5.

2.3 Haptic Rendering

Other work exploring tactile perception has been motivated by haptic applications. Early haptics research focused on virtual touch using rigid styluses or probes as feedback devices. For example,

Otaduy et al. (2005) developed a perceptual model for haptic rendering of forces and vibrations elicited by the contact of textures in this manner. Several studies have measured perceived surface roughness of virtual textures using a stylus (e.g., Klatzky and Lederman (1999), Yoshioka et al. (2007), and Okamoto et al. (2012)). Importantly, these works are based on a vibratory perceptual model of texture and roughness perception, as transmitted through a rigid object. In contrast, we aim to manufacture objects to be touched with bare skin, which provides more natural and improved tactile information.

Other recent developments have enabled the rendering of tactile features on various types of touch displays. Many types of tactile touchscreens have been proposed (e.g., Kim et al. (2013) and Meyer et al. (2013)), which use electrostatic attraction to alter friction and simulate texture or gradient. Bau and Poupyrev (2012) developed AR technology using reverse electrovibration, which injects electric signals into the user to elicit smaller-scale tactile sensations. However, this is still fundamentally vibratory signaling. Iwata et al. (2001) developed an early method of rendering haptic height surfaces directly by projecting an image on a flexible screen over an actuator array. More recently, Hashizume et al. (2016) developed a tactile rendering display using magnetorheological fluid (MRF), a liquid that changes viscosity in response to magnetic fields. These methods represent promising advancements to on-the-fly rendering of real textures and could be used in conjunction with our model to further explore roughness perception and texture rendering.

3 OVERVIEW AND MAIN RESULTS

The main goal of this work is to develop a computational model describing the perceptual roughness of a given geometric surface texture, which can be applied to produce a range of tactile surface behaviors.

Model Summary. Given an input 2D height field describing a surface texture $H(x, y)$, our model is a function $f(H)$ producing a scalar value estimating the perceptual roughness of the texture. Our pipeline is depicted in Figure 1. The perceptual roughness estimate $f(H)$ is obtained by first running a simulation to find the contact area and resulting deformation of a skin-like elastic medium in contact with the surface texture. Then the field of maximum compressive strain is computed at a fixed depth corresponding to the average depth of Merkel receptors. The strain field is sampled at pairs of points separated by a fixed distance, extracted from experimental data, to find the *mean strain fluctuation variation*, defined as the mean absolute difference between strain magnitudes of the sampled pairs. This value describes the perceptual roughness of the input geometry.

We use psychophysical experiments to validate that the perceptual space of roughness can be described by a single dimension and to construct a metric on this space. We demonstrate that our model describes this space with a linear relationship ($r = 0.911$). Experiments are also used to place a standard scale of reference textures in the perceptual space, which can be used as a reference set to match against other textures to predict roughness equivalence.

Our model tuning and validation process consists of the following steps:

- We synthesize a set of textures comprising a variety of parametric and natural textures. We fabricate these textures as flat plates using a high-resolution ($30\mu\text{m}$) stereolithography 3D printer.
- We perform a series of psychophysical experiments with human subjects, comparing the differences between textures.
- We use nonmetric multidimensional scaling (NMDS) with the experimental results to establish that the roughness perceptual space is one-dimensional (Section 4) and to place the stimuli in this space.
- We acquire the actual geometry of the printed plates using micro-CT scanning.
- For each texture, we perform FEM simulation to compute the deformation of a skin-like elastic medium when pressed in contact with the texture. An iterative process is used to resolve the contact area.
- From the simulation, we compute the maximum compressive strain field at a depth of 0.75mm , and we compute $f(H)$ as the strain fluctuation variation of this field.
- We test the linear dependence between the coordinates assigned by the NMDS to the stimuli H_i and the roughness estimates $f(H_i)$ produced by our model, and tune the sampling parameters of the model to match the perceptual data as closely as possible (Section 5).

4 PERCEPTUAL SPACE

In this section, we describe how we establish that the *perceptual space* for (coarse) roughness is one-dimensional and construct a map from stimuli to points in this space. The overall idea, following from previous work, is to use experiments to obtain an estimate of relative perceptual proximity for triplets of stimuli (i.e., which of a pair of stimuli feels closer to a reference stimulus). A multidimensional scaling algorithm is then used to assign n -dimensional points to stimuli so that Euclidean distances between these points have the same ordering as obtained from the experiments: stimuli corresponding to closer points feel more similar than stimuli corresponding to points at a larger distance.

4.1 Stimuli

A fundamental issue in the study of perceptual roughness is the large variety of possible surface textures. Inherently, conducting experiments with human subjects limits the number of textures we can study: using three-way comparison experiments, with N textures we need to obtain a reasonably large sample of $\binom{N}{3}$ combinations.

We used 46 textures in our experiments. We used two distinct types of textures as stimuli: artificially synthesized textures, and real-life “natural” textures described below.

Synthetic Texture Geometry. Synthetic textures serve two purposes: (1) to provide a one-dimensional family that can be used as a reference scale for roughness and (2) to explore the effects of some of the geometry parameters that previous studies have shown to be important for tactile perception.

Specifically, we choose bump textures that allow us full control over texton spacing, shape, and distribution (in the present study, all textures had elements with random isotropic distribution).




Fig. 3. Modeling process for isotropic textures. The modeling process allows for variation of texton arrangement, spacing λ , and shape (here we control the tip diameter d).

Bump spacing allows control of the spatial frequency of contact points between the texture and the skin, which is known to be correlated with roughness; bump shape, on the other hand, controls the size of the contact zones and the sharpness of the transition between contact and noncontact zones, which has a significant effect on strain distribution, a key component of our model. Unlike, e.g., ridge patterns, our bump textures are isotropic, so we can factor out the aspect of strong regularity.

The 3D modeling procedure used to specify isotropic textures is shown in Figure 3. We begin with a magnitude spectrum, representing the magnitude component of the Fourier transform of the image. For random, isotropic textures, with one dominant wavelength over all directions, this magnitude spectrum is visualized as the outline of a circle with a radius given by s/λ , where s is the pixel size of the image and λ is the dominant wavelength (Figure 3(a)).

We then take the 2D inverse Fourier transform using the magnitude and a random phase matrix, resulting in a random texture with a dominant wavelength of λ (Figure 3(b)). To produce the bump texture, we take the relative maxima of the inverse Fourier transform to obtain a dot pattern yielding the centers of the bumps. The resulting textured surface has textons arranged with an average spacing at the specified wavelength (Figures 3(c) and 3(d)). In this way, we have complete control over the local bump shape: in our experiments we use truncated cones with or without a rounded cap.

We fabricated a reference set of 13 standard isotropic bump samples varying in a single parameter, average distance between bumps, which ranges from 0.625 to 1.625mm . These textures have standard texture elements shaped as 1mm -tall truncated cones with spherical caps 0.3mm in diameter. Additionally, we fabricated five isotropic bump samples with texture elements shaped as truncated cones with flat tops of varying diameters from 0.1 to 0.5mm , to explore the dependence on the contact area.




Fig. 4. Psychometric curves fit to percentage data, from our previous study. Errors bars show SEM across subjects.

We also fabricated a set of 28 real-world texture height maps; of these, 24 were acquired from high-resolution scans on the SurfaceMimic 3D scan repository and three from the Brodatz texture database, and one was procedural. The texture set was composed of many common natural and manufactured materials, such as burlap, leather, and stucco; textures varied in height, pattern, and isotropy. We fabricated these textures out of the same material, instead of, e.g., using the natural samples, to eliminate the effects of small-scale roughness and compliance.

All texture samples were fabricated as 25-by-25mm square plates using a B9Creator stereolithography 3D printer with $50\mu\text{m}$ resolution in xy and $30\mu\text{m}$ resolution in z . Textures ranged from having little surface texture to having greatly varying surface texture. A selection of stimuli is shown in Figure 7, and a table of all 46 textures is provided as supplementary material.

Threshold of Discrimination. To select the reference set of bump textures, we refer to data gathered in our earlier study (described in detail in Tymms et al. (2017); all other studies presented in this article have not been described elsewhere).

In this previous experiment, 16 subjects participated in two-alternative-choice trials to identify which of two presented bump textures felt smoother. To this data, we fit psychometric curves (Figure 4) and extracted the absolute thresholds σ . The Weber fraction, or relative threshold of discrimination, describes the proportional difference in texture wavelength required in order to distinguish (with 75% accuracy) between two textures. The Weber fraction was found to be 0.19, which is in good agreement with the existing literature (Figure 5).

This value serves two important purposes: first, to act as a threshold for selecting substantially different textures, and second, to demonstrate that our choices for reference textures sample the perceptual space densely enough. At the smoother end of the scale, we fabricated textures with 0.0625mm wavelength differences; at the rougher end we used an interval of 0.125mm. These reference textures are shown in Figure 6 along with the threshold σ . In most cases, our textures sampled the space with a density around 2σ .




Fig. 5. The threshold of discrimination is proportional to the tested wavelength by a factor of 0.19, known as the Weber fraction.




Fig. 6. The standard scale of reference textures (numbers 1 through 13) depicted along a logarithmic line scaled to the threshold of discrimination σ . The textures sample the perceptual space at a greater density than the threshold of discrimination.

4.2 Experiments

Our experiments serve three purposes in this study:

- We validate the hypothesis that perceptual roughness is reasonably measured by a one-dimensional parameter, more specifically, that any set of textures can be ordered in a way so that if one texture comes before the other, it is always perceived as rougher.
- We show that the perceptual roughness ranking inferred from the experiments is closely matched to the one given by our choice of the model for roughness perception, the strain fluctuation variation model.
- By including a one-parametric reference set of textures in the set of stimuli, we establish a correspondence between a simple parameter of the family (spacing of the bumps) and the physical roughness measure.

The two most common types of psychophysical experiments used with human subjects involve either pairwise comparisons of roughness of stimuli (which stimulus feels rougher) or estimating subjective roughness on a scale, e.g., 0 to 10. The former provides a simpler and more intuitive task for the subjects, and as such the data is more reliable. However, the results are harder to use for perceptual space dimension analysis, as they do not provide direct information on the perceived distances between stimuli. The second type of experiments is more suitable for dimension analysis, but the task is more complex and leads to less reliable results.

In line with Piovarcı et al. (2016), we use a variation of the first type of the experiment, where instead of a single pairwise comparison between two textures per trial, two test textures are compared in their perceptual proximity to a third, reference texture. This type of experiment, as we explain below, can be used with




Fig. 7. Seven examples of the 46 textures used in our experiments. Three bump textures are shown with a green background, and four natural textures are shown with a blue background. From left to right: standard bump texture with wavelength 0.625mm; standard bump texture with wavelength 1.625mm; flat bump texture with wavelength 1.0 and size 0.4mm; foam; lizard skin; stucco; knit wool. A list of all 46 textures is provided in the supplementary material.




Fig. 8. In a trial, participants were presented with three texture samples. They touched each texture using the index finger and answered whether the leftmost or rightmost sample felt more similar in roughness to the middle sample.

multidimensional scaling to determine the perceptual space and dimensionality *without* relying on subjects assigning a subjective rating to each stimulus.

The experiment setup is shown in Figure 8. During each trial, the subject was given three textures, arranged in a row in a fitted case. They were asked to feel each texture by pressing with the index finger and to answer whether the leftmost or rightmost texture felt more similar in roughness to the center texture. Subjects were free to feel each texture multiple times before making the decision. Subjects initiated a trial by pressing down the right and left arrow keys on a provided keyboard with their nondominant hand before making contact with the textures with the dominant hand. They indicated their choice (the leftmost or rightmost texture) by releasing the left or right arrow key, respectively. This method allowed us to record both the subjects' responses and the amount of time taken for each response.

We included a set of practice trials at the beginning of each session to familiarize the subjects with the mechanics of the task and to ensure that they understood the task protocol. Visual feedback was avoided by using a cut box placed over the texture case so that textures could be touched but not seen from the subjects' viewpoint. Fourteen subjects (21 to 35 years old) participated in these experiments. All self-identified as strongly right-handed and reported normal sensory and motor functionality of the hands and fingers. Each subject performed trials for 1 to 2 hours, resulting in between 100 and 250 triplet comparisons. The order of trial comparisons was randomized.

4.3 Data Analysis

MDS. We briefly review NMDS, which we use to find a Euclidean embedding of the roughness space of our samples.

This method takes as input a set of triples $\{(i, j, k)\}$, where i, j, k are stimuli used in a trial (with j as the reference sample), and the ordering $D_{ij} < D_{jk}$, where D_{ij} is the perceptual distance between samples i and j . In other words, the input consists of the ordering of the pair comparisons, without magnitudes of the dissimilarities.

As in classical MDS, the problem is transformed into one that can be stated in terms of the Gram matrix K . Because all distances are assumed to be nonnegative, the set can be transformed to

$$S = \{(i, j, k) | D_{ij}^2 < D_{jk}^2\}.$$

Now

$$\begin{aligned} D_{ij}^2 &= \|x_i - x_j\|_2^2 = x_i^\top x_i - 2x_i^\top x_j + x_j^\top x_j \\ &= K_{ii} - 2K_{ij} + K_{jj}. \end{aligned}$$

Therefore, we can write the triplet conditions as

$$= K_{ii} - 2K_{ij} + K_{jj} < K_{jj} - 2K_{jk} + K_{kk}.$$

Thus, the goal is to find a Gram matrix K that satisfies the inequality constraints of this form for every triplet (i, j, k) in S . Since it is a Gramian matrix, K is the inner product for some set of points.

The set of constraints does not define a unique matrix K : it is still subject to scale, translation, and rotation. Therefore, we simply constrain K so that it is centered at the origin and has a range of $(-1, 1)$.

Finally, we want to minimize the dimensionality or rank of K while also ensuring that it adequately satisfies the inequality constraints. As the rank of a matrix is a nonconvex term and minimizing it while optimizing for inequality constraints is an NP-hard problem, Wills et al. (2009) formulate the problem as instead minimizing the trace of the matrix. The full problem is formulated as

$$\begin{aligned} \arg \min_{K, \xi} \sum_{(i, j, k) \in S} \xi_{ijk} + \lambda \text{Tr}(K) \\ \forall (i, j, k) \in S: SK_{kk} - K_{ii} + 2K_{ij} - 2K_{jk} \geq 1 - \xi_{ijk}, \\ \xi_{ijk} \geq 0, \sum_{ab} K_{ab} = 0, K \geq 0. \end{aligned}$$

ξ_{ijk} are slack variables introduced to allow for violations in the inequalities, and the variable λ is a regularization parameter that balances the complexity (i.e., dimensionality) of the model with the fit. The SeDuMi solver for Matlab is used to solve this optimization problem.

Cross-Validation. The goal of cross-validation is to find the optimal λ value. We use 10-fold cross-validation as suggested in Wills et al. (2009). In each iteration, the data is divided into 10 sets, with one set treated as the testing set and the rest as the training set. The optimization is run on the training set for a given λ , and the average testing and training error is reported as the mean across




Fig. 9. The placement of our 46 texture samples in one dimension using the NMSD algorithm. The number labels correspond to those in Figure 7 and in the table in the supplementary material.




Fig. 10. Left: Cross-validation determines the optimal λ value as 0.7. Right: Most of the variance is confined to the first dimension, indicating a good fit in one-dimensional space.

these trials. We aim to optimize the tradeoff between testing and training error. When λ is small, the algorithm favors reducing the violations for all conditions and results in low training error but a high testing error, implying that the embedding does not generalize and may overfit to noise in the data.

We obtained the optimal embedding with $\lambda = 0.7$ as shown in Figure 10. The spread for this embedding shows the variance across each dimension; most of the variance is contained in the first coordinate, indicating that the space fits well in one dimension.

Errors. To find margins of error, we use a standard bootstrapping technique. We recomputed the embedding 500 times by randomly resampling the data points with replacement, and we compute confidence intervals as the 5th and 95th percentile of the resulting embeddings. The final embedding with these error bars is shown in Figure 9. Samples are numbered according to the labels in Figure 7. Note that because of the large number of stimuli, only a small proportion of possible combinations was tested; therefore, the error bars represent upper bounds for error and are likely larger than the subjects' actual margin of error.

4.4 Curved Surfaces

Many natural applications for rough textured surfaces involve 3D objects with nonzero curvature, rather than flat objects. An obvious question is whether the perceived roughness of textured curved objects is the same as that for textured flat plates. To answer this question, we performed an additional psychophysical experiment comparing a new set of curved textures to the set of flat textured plates. The stimulus set was manufactured in the same way




Fig. 11. Our curvature experiments tested the roughness of textured curved surfaces with one of four curvatures (20 or 40mm, concave or convex) against the standard set of bump texture plates. Curved surfaces were textured with the standard bump texture in one of three wavelengths.

and consisted of 12 singly curved surfaces, textured with one of three bump texture wavelengths of different roughness (0.75mm, 1.0mm, 1.25mm). Textures had one of four different curvatures: concave or convex, with radius of curvature 20mm or 40mm, as shown in Figure 11. Here we define a concave object as having a negative radius of curvature, and a convex object as having a positive radius of curvature.

In experiments, subjects were presented with one curved bump texture and one flat bump texture. They were instructed to touch the surfaces with the index finger as in the previous experiment, and they answered which of the two textures felt smoother. Subjects performed between two and six comparisons per pair. Data was aggregated over all subjects to find the proportion for each pair, and the proportions were used to find the Point of Subjective Equality for each curved texture, i.e., the *flat* texture wavelength judged as equivalent to each *curved* texture. Here, we assume that if one textured surface is judged smoother than another in 50% of trials, the two are equivalent in roughness.

Figure 12 shows the point of subjective equality of the four curved textures for each wavelength. The overlapping shaded areas, as measured by the double arrows, show the interval for 19% threshold of the given reference wavelength, which represents the threshold of equivalence for that wavelength. For all wavelengths, all of the points fall well within this threshold area, indicating that the curved textures feel equivalent to their corresponding flat textures. While some data points (e.g., wavelength 0.75mm, radius -20) appear a small distance from the base wavelength indicated by the dashed line, most are close; furthermore, no trend exists over all of the textures to suggest that curvature has a consistent




Fig. 12. The points of subjective equality for all curved textures are plotted, with color indicating texture wavelength; the 19% threshold interval for each wavelength is indicated by the corresponding shaded intervals measured by the double arrows. All of the points of subjective equality fall within their corresponding interval, indicating that the curved textures are judged equivalent in roughness to the corresponding flat texture of the same wavelength.

effect on perceived roughness. The small and seemingly random differences likely result from variations caused by printing.

From these results, we posit that moderate curvature (with radius of curvature many times larger than the characteristic texture scale) does not have any significant effect on perceived roughness, and that the human tactile system integrates shape and texture independently. In regions of extreme curvature on the scale of texture itself, the separation may become unclear and curvature might have an effect on perceived roughness. Characterizing perception of roughness in these areas is more complicated, as it becomes much more dependent on the complex details of contact and highly dependent, e.g., on the exact finger location or orientation of the applied force. We propose a method of separating underlying smooth shape and texture in our model in Section 5, and we show that simulations match our experimental data in cases of moderate curvature.

5 MODEL FOR ROUGHNESS

5.1 Biological Basis

Our computational model for roughness is built on several previous works in neuroscience and engineering. It is derived from a combination of two parts: first, the relationship between the physics of skin deformation and neural responses, and second, the relationship between neural responses and high-level perceived roughness. The former connection has been studied on nonhuman primates with direct measuring of skin deformation and neuron firing rates. The latter was studied using firing rate data from nonhuman primates and perceived roughness ratings from human subjects using the same stimuli.

In the context of the first question, Phillips and Johnson (1981) created a continuum mechanics model for skin deformation and tested several models for predicting the SA1 firing rate data under

bars and gratings; they found that firing rates were most closely correlated to maximum compressive strain and strain energy density in the skin. Dandekar et al. (2003) developed an FEM model of the skin and examined this model in relation to firing rate data under line loads, finding again that SA1 firing rates were well predicted by the maximum compressive strain at the receptor depth in the skin.

Connor and Johnson (1992) and Yoshioka et al. (2001) examined the second question of the relationship between firing rates and perceptual roughness assessments, using as stimuli embossed dots or square wave gratings with different spacings. After examining several hypotheses, they both found that perceptual roughness best matched to the spatial variation in SA1 firing rates, meaning the average absolute difference between SA1 firing rates, at intervals of approximately 2mm. They also verified that Merkel receptors were the major predictors of surface roughness; other mechanoreceptors (RA and PC) did not have strong correlation with roughness judgments.

We combine observations from these studies to hypothesize that *strain fluctuation variation* at the depth of mechanoreceptors in skin is a physical quantity well correlated with subjective roughness perception. We define strain fluctuation variation as the mean absolute difference of maximum compressive strain fluctuation between pairs of sample points in the skin, after the effects of bulk shape are eliminated. The difficulty of using this quantity is estimating the strain of the skin when in contact with specific geometry, which we discuss now.

5.2 Simulation

In general, accurate simulation of skin deformation is a difficult problem, primarily because of the complex nonlinear and inhomogeneous mechanical properties of biological tissues and potentially high variability of mechanical properties among subjects.

However, we observed that the specific quantity of interest, variation of strain, is relatively insensitive to these properties: e.g., increasing or decreasing Young's modulus of the material results in relatively small changes in the contact with the surface (Figure 18).

We first simulate the skin deformation when in contact with particular texture geometry of the texture defined by a height field $H(x, y)$.

To describe our computation of the roughness value $f(H)$, we introduce the following notation:

- $\mathbf{u}_H(u, v, w)$ is the field of displacements obtained by solving an elastic contact problem between a model for skin and subcutaneous tissue and the height field $H(x, y)$. (u, v, w) are the coordinates in the material domain representing the tissue, with $w = 0$ corresponding to the lower surface.
- $\epsilon(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + \nabla^T\mathbf{u})$ is the small-deformation approximation of the strain tensor.
- $\lambda_3(\epsilon)$ is the largest-magnitude negative eigenvalue of a symmetric tensor.
- $V(f)$ is the variation functional described in detail in the next section.
- w_0 is the depth of the receptors, which we take to be 0.75mm (Dandekar et al. 2003).

Then our roughness estimate $f(H)$ can be written as a composition:

$$f(H) = \mathcal{V}(\lambda_3(\epsilon(\mathbf{u}_H(\cdot, \cdot, w_0)))).$$

The critical and most expensive step is the computation of \mathbf{u}_H for a given H . Although we use simple linear elasticity as a model, this function is *nonlinear* due to the need to resolve contact. This function also depends on the choice of the model for tissue and applied pressure.

Skin Model and Geometry Data. To approximate the skin, we use a two-layer block model created with Tetgen (Si 2015), adaptively refined with about 24,000 nodes, which represents 1cm^2 in surface area and 0.5cm in height. The lower half of the model has a Young's modulus of 0.042MPa , representing all tissue, and the upper half, representing the bone, has a Young's modulus of $18,000\text{MPa}$, i.e., essentially rigid (Agache et al. 1980). The force used for this model was 10N , which we found in pilot experiments where subjects touched textured surfaces placed on a force sensor. As seen in Figure 18, however, the model is relatively insensitive to small changes in force.

While this model is rather crude, given the level of noise in the experimental data, we found it to be adequate. We did not fit the parameters of the simulation to the data, but this would certainly be possible once larger datasets with less variance are collected.

We used as input $H(x, y)$ scanned height maps of the textures, rather than the original digital models from which the textures were printed, to take into account the effects of the printing process (inflation, rounding of edges which results in decreased strain, etc.). We acquired high-resolution 3D scans of all of the textures using a Bruker Skyscan 1172 micro-CT scanner.

Resolving Contact. To build an accurate simulation, we must find the contact between the skin and the rigid surface. For most surfaces, this problem is nontrivial, as the skin may or may not come into contact with various parts of the surface geometry depending on the heights, sizes, and distribution of features. Because of the complexity of the surface geometries, we were unable to obtain reliable results from commercial FEM software supporting contact mechanics. Instead, we have implemented our own accurate simulation software based on a simple and robust algorithm, illustrated in Figure 13.

First, the undeformed geometry is positioned so that it touches the top points of the texture; i.e., a map $p(u, v, w) \rightarrow \mathbb{R}^3$ is defined, so that $p_z(u, v, 0) = \max_{x, y} H(x, y)$. The points in contact with the texture are included in the initial contact zone Ω_c .

The vertices constrained by this boundary represent the parts of skin that touch the upper crests in the contacted texture. Then, an iterative process is run with two steps to build the correct boundary condition. The simulation is run for each step. First, we add to the fixed boundary set Ω_c any vertices that fall below the texture height $H(x, y)$. However, this process may overestimate the vertices in contact with the substrate. Thus, a second step of the iterative process releases vertices for which the elastic forces are pointing away from the contact surface. This iterative process terminates after it converges or after a number of iterations to approximate the contact condition. We found that the process typically converges in three of these two-step iterations.

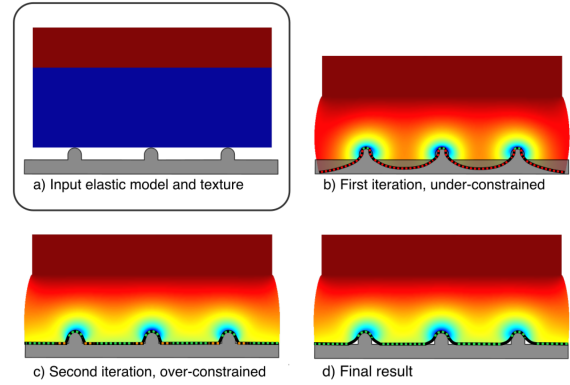


Fig. 13. (a) The iterative process simulates the contact between the skin (blue) attached to a rigid bone (red) when the skin contacts a textured surface (gray). (b) The simulation begins with an arbitrary contact area set using a threshold (constrained vertices are shown in green). The resulting boundary condition may be underconstrained, as many vertices fall below the contact area (shown in red). (c) All vertices that fall beneath the boundary are constrained. However, some vertices (as shown in orange) may be overconstrained, or stretched. (d) Overconstrained vertices are released from the boundary. After a number of iterations, the result converges to a stable contact boundary condition.

The pseudocode below summarizes the algorithm more formally; σ denotes the stress and \mathbf{n} the normal to the surface:

```

 $\Omega_c \leftarrow \{(u, v) | p_z(u, v, 0) = H(p_x(u, v, 0), p_y(u, v, 0))\}$ 
repeat
  Solve for  $\mathbf{u}$  with  $\Omega_c$  fixed
   $\mathbf{p}' \leftarrow \mathbf{p} + \mathbf{u}$ 
   $\Omega_c^{incl} \leftarrow \{(u, v) | p'_z(u, v, 0) < H(p'_{x'}(u, v, 0), p'_{y'}(u, v, 0))\}$ 
   $(u, v) \notin \Omega_c$ 
   $\Omega_c^{excl} \leftarrow \{(u, v) | \mathbf{n}^T \sigma(u, v, 0) < 0, (u, v) \in \Omega_c\}$ 
   $\Omega_c \leftarrow \Omega_c \cup \Omega_c^{incl} \setminus \Omega_c^{excl}$ 
until  $\Omega_c$  does not change

```

5.3 Computing Variation

Strain Fluctuation. In 3D objects with a nonzero curvature, the simulated maximum compressive strain field indicates two independent properties: the global object shape, indicated by the large-scale, global changes in strain across the surface, and the texture, which results in small-scale local *fluctuation*. An object with zero curvature has a negligible change in global strain, resulting in a globally flat strain field; however, curved surfaces elicit a varying strain field that reflects the global shape of the curved surface, as shown by the blue curve in Figure 15. If our model is to be used directly on curved surfaces, we must separate the qualities of shape and texture and eliminate effects from shape when computing the texture roughness. Thus, we propose a strain fluctuation variation model, which eliminates effects in the maximum compressive strain field caused by changes in large-scale shape.

Given a textured surface, we first compute the maximum compressive strain field as described in the previous section. We also compute the maximum compressive strain field for a *nontextured* surface with the same global curvature. We subtract this nontextured strain field from the original strain field to find an




Fig. 14. Examples of simulation input height maps (left) and output compressive strain fields (right), visualized as grayscale height maps.

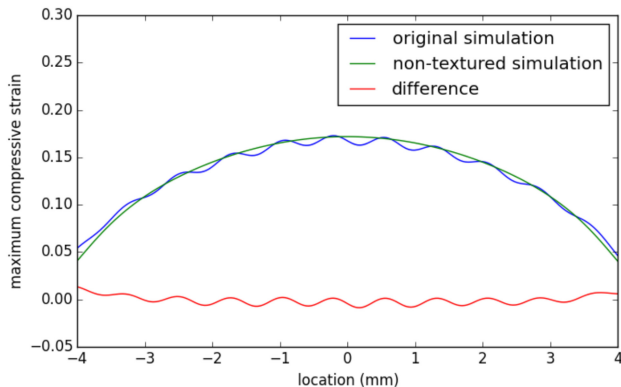


Fig. 15. For curved objects, we must find an updated strain field that eliminates effects from large-scale shape. As shown in blue, the original strain field changes over the surface, according to the convex object shape. However, when we compute and subtract the strain field of the same nontextured convex shape (green), we acquire the texture strain *fluctuation* field without any effects from shape (red).

updated field that eliminates effects from large-scale shape. A two-dimensional simulation example is shown in Figure 15. Note that as described below, we take the variation, or absolute difference, between points on this field; therefore, the change in magnitude across the whole field does not have an effect on the final value.

Roughness may be measured on input geometry (e.g., acquired by scanning) that is not initially separated into a smooth base surface and a displacement field. The spectral domain provides a natural way to perform this separation by extracting the low-frequency part below the texture frequency. Typically, however, the input samples can be obtained in a flat form; for tasks like texture synthesis and application of an existing texture to a surface, no




Fig. 16. For each texture simulation, the maximum compressive strain profile was obtained for a depth of 0.75mm. To find the strain variation, pairs of sample points (shown in red) were chosen from inside sampling areas (shown in yellow) with radius r separated by distance d . The mean absolute difference of all such pairs is defined to be the mean strain variation. Mean strain variation was found to match perceptual estimates of roughness.

adjustment of roughness is needed as long as the curvature is not extreme.

Strain Variation. From the maximum compressive strain fluctuation field (Figure 14), we take a set of random samples. Previous works have suggested taking an interval of 2.2mm (Connor et al. 1990) or anywhere from 1 mm to 3 mm (Yoshioka et al. 2001). In reality, mechanoreceptors occupy random locations in the skin according to an average density, and a single neuron may be informed by multiple branching receptors. Thus, our sample pairs do not have a single exact spacing; instead, we take samples randomly from radii r separated by a given distance d , as shown in Figure 16. We take 8,000 random sample pairs across a 8×8 mm area in the center of the simulation area (to avoid unwanted effects from the edges). From these pairs we find the mean strain variation, the average absolute difference in maximum compressive strain at those point pairs. We found the optimal sampling for our perceptual data to be $d = 0.8$, $r = 0.8$, which resulted in a coefficient of correlation $R = 0.91$. As a note, many other samplings also produced good fits, with correlation coefficients greater than 0.85 (see Figure 18).

5.4 Results

The results of our method and a comparison to several measures suggested in previous works are shown in Figure 17. R_a is the arithmetic average of absolute values of heights, a simple roughness parameter commonly used in profilometry that has also been used to estimate perceptual roughness in many works (e.g., Chen et al. (2009)). The 90th percentile height, GLCM sum of squares, and GLCM variance were suggested by Elkharraz et al. (2014) as close estimates to perceptual roughness. Our strain variation roughness measure performs significantly better than any of the other measures.

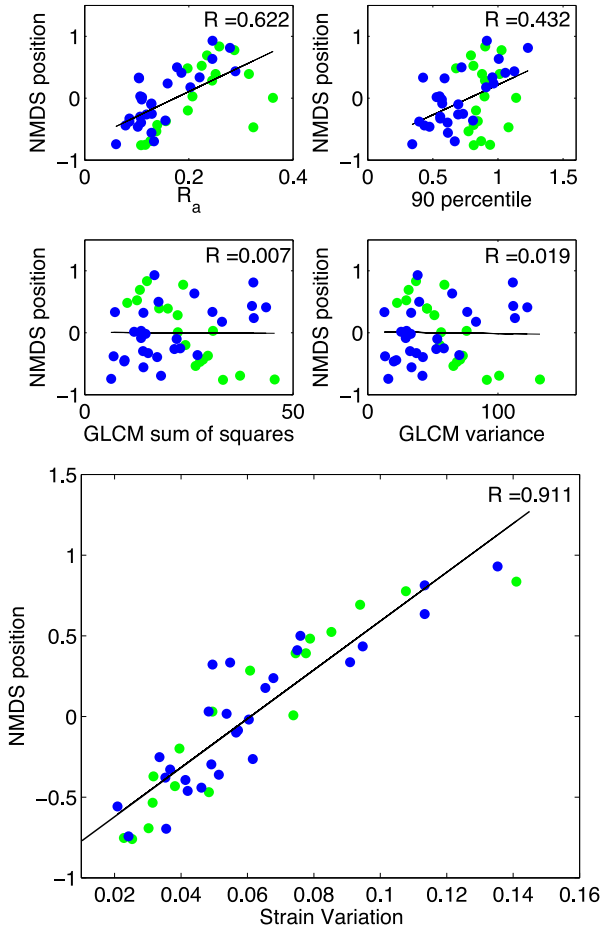


Fig. 17. Comparison of roughness values found by four different methods suggested in previous work (R_a , 90th percentile, GLCM sum of squares, and GLCM variance) and our method (strain variation). Our method produces the best fit, with an R value of 0.91. The bump textures are shown in green dots, and the natural textures are shown in blue dots.

Model Sensitivity. To test whether this model is stable with respect to the force and sampling parameters, we recomputed the fit with small changes to these variables (Figure 18). As shown on the left, the model is relatively insensitive to changes in the simulated force. Good fits are still produced within a range of 50% to 200% change in force from our choice of 10N. The model may fail if the force is very high or low, likely because the contact area changes significantly. As shown on the right, our model was also found to be insensitive to changes in the sampling parameters; we tested a large range of sampling parameters, and all resulted in a correlation of $R > 0.85$.

Texture Homogeneity. In this article, we define the word *texture* as a somewhat homogeneous surface, which has a consistent roughness across all areas. The textures we used as stimuli seemed homogeneous in appearance, but we also wanted to determine whether the textures were homogeneous in tactile roughness. To this end, we computed the roughness of each surface on fifty 5mm square sliding windows across the span of the surface. We




Fig. 18. Left: The model is relatively insensitive to small changes in applied force, but the fit worsens for large changes. Right: The model fit is relatively insensitive to change in the choices of sampling distance and radius; the fit changes slightly but still has a correlation over 0.85 for all values tested.




Fig. 19. Fit for the data excluding the five textures shown in red, which had significantly higher variance in calculated roughness across the surface.

then computed the *variance* of the computed roughnesses in proportion to the average roughness of the surface. Most textures had a very small variance, on the order of 10^{-4} . Five natural textures had significantly higher variance, greater than 10^{-3} . These textures are highlighted in red in Figure 19. The calculated R correlation value excluding these textures is slightly improved, at $R = 0.965$. We posit that the inhomogeneity of these textures may have caused subjects difficulty in determining the roughness consistently across the surface; when touching different regions of the textures, subjects may have perceived contradictory results.

Curved Surfaces. Figure 20 shows results of the strain fluctuation variation model for curved surfaces with the three textures tested in Section 4. For each of the three textures, the predicted roughness of 10 different curvatures was compared against the roughness of the uncurved texture, and the percentage difference from the uncurved texture was calculated. As expected, the percentage difference between the curved and flat texture was typically below the 19% threshold, except in the case of high curvatures.

6 APPLICATIONS

Applying particular tactile textures to fabricated objects is useful for both functional and aesthetic reasons. Many everyday objects have a tactile texture, including figurines, grips and handles, buttons, and other parts of objects that are used for interaction. While past work (e.g., Torres et al. (2015)) has been interested in applying tactile textures to objects, they did not use a specific perceptual model. Our model can be used to confer more precise and more diverse roughness characteristics to fabricated objects.




Fig. 20. For the three textures used in Section 4, the percent difference in simulated roughness was calculated between the flat texture and curved textures of different radii. The percent difference is typically below the 19% threshold (the black dashed line), with the exception of very high curvatures (less than 10mm for concave surfaces, and less than 3mm for convex surfaces).




Fig. 22. (a) Hedgehog with a rough bump texture that feels spiky. (b) Screwdriver handle with a moderately rough bump texture for easy and comfortable gripping. (c) Spinning “fidget ring” with rough cloud texture.




Fig. 23. From an original image (a), a depiction of a lion was created using two different roughnesses to correspond to the colors (b). The fabricated model is shown in (c). (The black image outlines are represented as smooth raised lines.)




Fig. 21. Two models, the gecko (left) and key cap (right), were fabricated with two textures each. Fabricated models have height textures scaled to have different tactile roughnesses, as matched against the standard set.

In this section, we describe several applications illustrating how our model can be used in fabrication. We first show that it can aid in the process of choosing textures to fabricate a 3D object with a desired roughness. Second, we show how our scaled perceptual model can be used to fabricate representations of visual images using distinct tactile properties in the place of distinct visual properties, such as color. Lastly, we demonstrate how our general model can be used to manufacture unique combinations of visual and tactual properties.




Fig. 24. A representation of De Stijl-style geometric artwork (a) was created using five different easily distinguishable roughnesses (b) to represent the five colors. The fabricated model is shown in (c).

Height Maps and Roughness. Often, users desire to apply a given height map to a texture to confer both visual and tactual properties; however, until now, it has been difficult to predict or alter the tactile properties of a visual height map. Using our model, we can make small changes to the scale or height of a height map to preserve the appearance while drastically altering the tactile roughness. We show two examples in Figure 21. First, a gecko model (WebmasterZero 2016) is given different displacements of the same “lizard skin” texture and three of the “burlap” texture. By altering the scale of these textures, we produce three different geckos, each with distinct tactile roughnesses.

We also produce a second model, a textured key cap of the type used for key identification; the texture enables the key to be identifiable in the dark. The key caps are textured with two scaled versions of the “leather” height map and three of the “foam” height map, and each of these occupies a different perceptual roughness.

Additional examples of fabricated tactile objects are shown in Figure 22: a hedgehog model with a rough bump texture, a screwdriver handle with a moderately rough texture, and a spinning “fidget ring” with a rough cloud texture.

Tactual Images. Producing tactual images is of interest not only for aesthetic reasons but also for practical reasons relating to




Fig. 25. Textured plates with visual (left) or tactile (right) renderings. The top row shows the rendered models; the middle row shows the 3D prints; the bottom row shows a visual representation of the tactile strain map. The left panel shows four texture plates with visible features that cannot be perceived tactually. Conversely, the right panel shows the same four shapes, which cannot be seen but whose outlines can be felt, as seen in their strain profiles.




Fig. 26. Textures can be produced that contain hidden visual or tactual messages.

accessibility: people who are blind can greatly benefit from textured representations of images, e.g., picture books, educational materials, or works of art (Theurel et al. 2013). We use the roughness scale generated by our model to identify sets of textures that have significantly different tactile roughnesses. In this way, we can replace the colors in images with distinct textures, so that the colors can effectively be perceived tactually. Figures 23 and 24 show examples of images fabricated in this way. As shown in the center panels, the chosen textures are distinct from one another by an interval larger than the threshold of discrimination.

Controlling Visual and Tactile Properties Independently. To demonstrate the control over tactile properties our model affords, we consider two examples: visible variation in a geometric texture that is perceived as tactually homogeneous and, conversely, an invisible variation in texture producing a variation in tactile perception.

For the first example, we manufacture several bump texture plates that feel like uniform textures; however, the noncontacted parts of the model are altered to produce a visual feature on the model, i.e., a feature that can be perceived visually but not tactually.

Similarly, we show that we can manufacture an object that has a consistent visual appearance but contains a hidden tactual characteristic. We manufacture the same plates, using the variation in texton height that produces strain differences, so that the difference cannot be clearly seen but the shape outlines can distinctly be felt.

These plates, fabricated with visible or tactile shapes, are shown Figure 25 and Figure 26.

In an informal experiment, eight subjects were presented with the tactile plates in Figure 25 either in a visual or tactile context and asked to identify whether a shape was present on the plate. On the tactile plates (Figure 25, right), all subjects were able to tactually identify the presence of shapes, and in most cases could identify the specific shape (the cross shape was most difficult to discern); in the visual plates (Figure 25, left), subjects did not tactually perceive any shape. Similarly, subjects could not visually identify specific shapes in the tactile plates, but they were able to clearly see the shapes on the visual plates.

7 CONCLUSION

We have presented a perceptually based model for evaluating the tactile roughness of surface textures. Using perceptual experiments with 46 different surface textures, we found that the perceptual space for roughness fits well in one dimension, and we used this space to create a physically based perceptual model for evaluating tactile roughness. We demonstrated how this model can be used to predict or compare the tactile roughness of fabricated objects, to select surface textures with different roughnesses, and to create textures with unique combinations of visual and tactile features.

7.1 Limitations and Future Work

Our model only applies to large-scale roughness. Our stimulus set was limited to a single material and a minimum feature resolution, as our 3D printer tends to smooth small features, so fine features ($<0.2\text{mm}$) were not possible to test. As mentioned in Section 5, our model only applies to homogeneous textures, which feel the same over the surface; in practice, it may give incorrect results for non-homogeneous textures. However, our model also naturally provides a method of testing whether a texture is homogeneous. While our provided model relies on our experimental evidence that curvature does not have a significant effect on roughness perception, we note that results may be different for extreme high-curvature objects where the contact area is small.

Our model represents a simplification of the mechanosensory system. In reality, the skin is not a uniform substrate; its structure is nonisotropic and layered, and it has complex, nonuniform geometry, e.g., fingerprint ridges. While our model appears to work well, it may be improved by more closely simulating the skin structure and mechanosensory mechanics.

Our model requires FEM simulations to evaluate a roughness estimate for a texture. It would benefit, at minimum, from precomputing the roughnesses for the desired set of textures. However, we believe that one can extract an adequate approximation for the map $f(H)$ from running multiple simulations and using machine-learning techniques to predict the roughness value directly from height map data.

To aid in the fabrication process, our model and reference stimuli could be integrated into an existing tool such as that of Torres et al. (2015) to provide more precise predictions of tactual roughness when modeling a fabricable object. Fabricating objects with different roughnesses and other tactual properties, such as weight or compliance, would create interesting stimuli for future studies of perception.

Finally, as our model performs significantly better than existing methods to obtain tactual roughness, it also has broad implications for other fields of work: in addition to use in psychophysics and neuroscience, it may also be of use for virtual haptics, robotics tactile perception, and many other fields with an interest in perception or tactual properties.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their helpful comments.

REFERENCES

- Pierre G. Agache, C. Monneur, Jean Luc Leveque, and Jean De Rigal. 1980. Mechanical properties and Young's modulus of human skin in vivo. *Archives of Dermatological Research* 269, 3 (1980), 221–232.
- Olivier Bau and Ivan Poupyrev. 2012. REVEL: Tactile feedback technology for augmented reality. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 89.
- Desai Chen, David I. W. Levin, Piotr Didyk, Pitchaya Sitthi-Amorn, and Wojciech Matusik. 2013. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 135.
- Xiaojuan Chen, Fei Shao, Cathy Barnes, Tom Childs, and Brian Henson. 2009. Exploring relationships between touch perception and surface physical properties. *International Journal of Design* 3, 2 (2009), 67–76.
- OpenStax CNX. 2017. Somatosensation. Retrieved from <https://archive.cnx.org/contents/b32f61fc-5fab-4b07-bcc8-e455aa4a903d@6/somatosensation>. Modified.
- Charles E. Connor, Steven S. Hsiao, John R. Phillips, and Kenneth O. Johnson. 1990. Tactile roughness: Neural codes that account for psychophysical magnitude estimates. *Journal of Neuroscience* 10, 12 (1990), 3823–3836.
- Charles E. Connor and Kenneth O. Johnson. 1992. Neural coding of tactile texture: Comparison of spatial and temporal mechanisms for roughness perception. *Journal of Neuroscience* 12, 9 (1992), 3414–3426.
- Kiran Dandekar, Balasundar I. Raju, and Mandayam A. Srinivasan. 2003. 3-D finite-element models of human and monkey fingertips to investigate the mechanics of tactile sense. *Journal of Biomechanical Engineering* 125, 5 (2003), 682–691.
- Galal Elkharraz, Stefan Thumfart, Diyar Akay, Christian Eitzinger, and Benjamin Henson. 2014. Making tactile textures with predefined affective properties. *IEEE Transactions on Affective Computing* 5, 1 (2014), 57–70.
- James M. Goodman and Sliman J. Bensmaia. 2017. A variation code accounts for the perceived roughness of coarsely textured surfaces. *Scientific Reports* 7 (2017), 46699.
- Satoshi Hashizume, Kazuki Takazawa, Amy Koike, and Yoichi Ochiai. 2016. Cross-field haptics: Push-pull haptics combined with magnetic and electrostatic fields. In *ACM SIGGRAPH 2016 Posters*. ACM, 30.
- Mark Hollins, Richard Faldowski, Suman Rao, and Forrest Young. 1993. Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis. *Perception & Psychophysics* 54, 6 (1993), 697–705.
- Mark Hollins and S. Ryan Risner. 2000. Evidence for the duplex theory of tactile texture perception. *Perception & Psychophysics* 62, 4 (2000), 695–705.
- Hiroo Iwata, Hiroaki Yano, Fumitaka Nakaizumi, and Ryo Kawamura. 2001. Project FEELEX: Adding haptic surface to graphics. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 469–476.
- Roland S. Johansson and Goran Westling. 1984. Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Experimental Brain Research* 56, 3 (1984), 550–564.
- Seung-Chan Kim, Ali Israr, and Ivan Poupyrev. 2013. Tactile rendering of 3D features on touch surfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. ACM, 531–538.
- Roberta L. Klatzky and Susan J. Lederman. 1999. Tactile roughness perception with a rigid link interposed between skin and surface. *Perception & Psychophysics* 61, 4 (1999), 591–607.
- Manfred Lau, Kapil Dev, Weiqi Shi, Julie Dorsey, and Holly Rushmeier. 2016. Tactile mesh saliency. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 52.
- Susan J. Lederman. 1974. Tactile roughness of grooved surfaces: The touching process and effects of macro- and microsurface structure. *Perception & Psychophysics* 16, 2 (1974), 385–395.
- Louise R. Manfredi, Hannes P. Saal, Kyle J. Brown, Mark C. Zielinski, John F. Dammann, Vicky S. Polashock, and Sliman J. Bensmaia. 2014. Natural scenes in tactile texture. *Journal of Neurophysiology* 111, 9 (2014), 1792–1802.
- David J. Meyer, Michael A. Peshkin, and J. Edward Colgate. 2013. Fingertip friction modulation due to electrostatic attraction. In *World Haptics Conference (WHC'13)*. IEEE, 43–48.
- Shogo Okamoto, Takahiro Yamauchi, Masashi Konyo, and Satoshi Tadokoro. 2012. Virtual active touch: Perception of virtual gratings wavelength through pointing-stick interface. *IEEE Transactions on Haptics* 5, 1 (2012), 85–93.
- Miguel A. Otaduy, Nitin Jain, Avneesh Sud, and Ming C. Lin. 2005. Haptic display of interaction between textured models. In *ACM SIGGRAPH 2005 Courses*. ACM, 133.
- John R. Phillips and Kenneth O. Johnson. 1981. Tactile spatial resolution. II. Neural representation of bars, edges, and gratings in monkey primary afferents. *Journal of Neurophysiology* 46, 6 (1981), 1192–1203.
- Michal Piovarči, David I. W. Levin, Jason Rebello, Desai Chen, Roman Đurkovič, Hanspeter Pfister, Wojciech Matusik, and Piotr Didyk. 2016. An interaction-aware, perceptual model for non-linear elastic objects. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 35, 4 (2016), 55.
- Hang Si. 2015. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software (TOMS)* 41, 2 (2015), 11.
- Anne Theurel, Arnaud Witt, Philippe Claudet, Yvette Hatwell, and Edouard Gentaz. 2013. Tactile picture recognition by early blind children: The effect of illustration technique. *Journal of Experimental Psychology: Applied* 19, 3 (2013), 233.
- Wouter M. Bergmann Tiest. 2010. Tactual perception of material properties. *Vision Research* 50, 24 (2010), 2775–2782.
- Cesar Torres, Tim Campbell, Neil Kumar, and Eric Paulos. 2015. HapticPrint: Designing feel aesthetics for digital fabrication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. ACM, 583–591.
- Chelsea Tymms, Denis Zorin, and Esther P. Gardner. 2017. Tactile perception of the roughness of 3D-printed textures. *Journal of Neurophysiology* 119, 3 (2017), 862–876.
- Alison I. Weber, Hannes P. Saal, Justin D. Lieber, Ju-Wen Cheng, Louise R. Manfredi, John F. Dammann, and Sliman J. Bensmaia. 2013. Spatial and temporal codes mediate the tactile perception of natural textures. *Proceedings of the National Academy of Sciences* 110, 42 (2013), 17107–17112.
- WebmasterZero. 2016. Gecko. Retrieved from <https://www.thingiverse.com/thing:1363148>. Modified.
- Josh Wills, Sameer Agarwal, David Kriegman, and Serge Belongie. 2009. Toward a perceptual space for gloss. *ACM Transactions on Graphics (TOG)* 28, 4 (2009), 103.
- Seung-Hyun Woo, Sanjeev Ranade, Andy D. Weyer, Adrienne E. Dubin, Yoshichika Baba, Zhaozhu Qiu, Matt Petrus, Takashi Miyamoto, Kritika Reddy, Ellen A. Lumpkin, Cheryl L. Stucky, and Ardem Patapoutian. 2014. Piezo2 is required for merkel-cell mechanotransduction. *Nature* 509, 7502 (2014), 622–626.
- Takashi Yoshioka, Sliman J. Bensmaia, Jim C. Craig, and Steven S. Hsiao. 2007. Texture perception through direct and indirect touch: An analysis of perceptual space for tactile textures in two modes of exploration. *Somatosensory & Motor Research* 24, 1–2 (2007), 53–70.
- Takashi Yoshioka, Barbara Gibb, Andrew K. Dorsch, Steven S. Hsiao, and Kenneth O. Johnson. 2001. Neural coding mechanisms underlying perceived roughness of finely textured surfaces. *Journal of Neuroscience* 21, 17 (2001), 6905–6916.

Received November 2017; revised March 2018; accepted March 2018