

Frequency-based heuristics for material perception

Martin Giesel

Graduate Center for Vision Research,
SUNY College of Optometry, New York, NY, USA



Qasim Zaidi

Graduate Center for Vision Research,
SUNY College of Optometry, New York, NY, USA

People often make rapid visual judgments of the properties of surfaces they are going to walk on or touch. How do they do this when the interactions of illumination geometry with 3-D material structure and object shape result in images that inverse optics algorithms cannot resolve without externally imposed constraints? A possibly effective strategy would be to use heuristics based on information that can be gleaned rapidly from retinal images. By using perceptual scaling of a large sample of images, combined with correspondence and canonical correlation analyses, we discovered that material properties, such as roughness, thickness, and undulations, are characterized by specific scales of luminance variations. Using movies, we demonstrate that observers' percepts of these 3-D qualities vary continuously as a function of the relative energy in corresponding 2-D frequency bands. In addition, we show that judgments of roughness, thickness, and undulations are predictably altered by adaptation to dynamic noise at the corresponding scales. These results establish that the scale of local 3-D structure is critical in perceiving material properties, and that relative contrast at particular spatial frequencies is important for perceiving the critical 3-D structure from shading cues, so that cortical mechanisms for estimating material properties could be constructed by combining the parallel outputs of sets of frequency-selective neurons. These results also provide methods for remote sensing of material properties in machine vision, and rapid synthesis, editing and transfer of material properties for computer graphics and animation.

judge its warmth and softness. Estimating material properties is thus often at least as important as recognizing object classes (Adelson, 2001; Anderson, 2011; Zaidi, 2011).

Sometimes the best way to estimate the relevant property of a material is to feel, hear, or smell it, but we reliably ascribe properties to materials just by visual inspection (Bergmann-Tiest & Kappers, 2007; Binns, 1937). For example, when using sandpaper, we usually judge the roughness visually, and when buying a waterproof or water-absorbent material we rely on visual judgments of porousness rather than making direct tests. Visual inferences are especially important when judgments have to be made rapidly or at greater distance than arm's length.

In all the examples above, the intended use of the material specifies the relevant affordances (Gibson, 1986), i.e., properties that allow particular uses, which in turn determine a suitable observation distance for resolving diagnostic features. These judgments are based on the retinal images projected from interactions between illumination geometry, material structure, and object shape. Without external constraints, the physics of these interactions are too involved (Anderson, 2011; Koenderink & Doorn, 1996; Koenderink, Doorn, Dana, & Nayar, 1999) for the visual system to estimate material structure by inverse optics, yet human abilities to judge material properties such as roughness or porousness show that we can, to a certain degree, infer the underlying physical structure from the retinal image.

In images of materials, the 3-D structure of a surface is mainly conveyed by shape from shading. Many algorithms for the recovery of 3-D shape from shading cues have been proposed in the computational literature (Breton & Zucker, 1996; Horn, 1970; Ikeuchi & Horn, 1981; Pentland, 1982), and a few studies have investigated human performance in recovery of object shape (De Haan, Erens, & Noest, 1995; Erens, Kappers, & Koenderink, 1993), but none have studied

Introduction

When deciding where to walk, it is not sufficient to recognize a path, but also necessary to ascertain that it is not slippery (Lesch, Chang, & Chang, 2008), muddy, or flooded. When choosing a sweater, it is not enough to differentiate it from coats and jackets, but also to

Giesel, M., & Zaidi, Q. (2013). Frequency-based heuristics for material perception. *Journal of Vision*, 13(14):7, 1–19, <http://www.journalofvision.org/content/13/14/7>, doi:10.1167/13.14.7.

doi: 10.1167/13.14.7

Received June 14, 2013; published December 6, 2013

ISSN 1534-7362 © 2013 ARVO

the identification of material properties from shading cues.

An attractive possibility is that we use heuristics based on rapidly extracted image properties to infer those attributes that are relevant to our interaction with a material. Such heuristics have been proposed for various material properties, for example, X junctions and contrast relations for transparency (Beck, Prazdny, & Ivry, 1984; Metelli, 1985; Robilotto, Khang, & Zaidi, 2002), highlights, contrast, blur, etc. for the perception of translucency (Fleming & Bühlhoff, 2005), and statistics based on the luminance distribution of an image for estimating gloss properties (Motoyoshi, Nishida, Sharan, & Adelson, 2007; Sharan, Li, Motoyoshi, Nishida, & Adelson, 2008).

An objection against the use of image statistics for the recovery of material properties is that they are useful only under limited conditions mainly because they disregard spatial structure and scale (Anderson & Kim, 2009). In the case of gloss, it has been shown that the perception of surface gloss is affected by the 3-D structure (Ho, Landy, & Maloney, 2008; Nishida & Shinya, 1998; Olkkonen & Brainard, 2011; Vangorp, Laurijssen, & Dutré, 2007; Wijntjes & Pont, 2010), and velocity flows (Doerschner et al., 2011). Both aspects are unaccounted for by simple image statistics. However, this shortcoming could be alleviated by combining image measures that cover different aspects of an object's appearance. Along this line, linear combinations of image cues have recently been proposed as predictors for perceived glossiness (Marlow, Kim, & Anderson, 2012) and perceived viscosity (Fleming & Paulun, 2012).

An alternative or complement to this approach could be image measures that by themselves are reliably related to 3-D configurations. Pentland (1984a), for example, reported a correlation between the perceived roughness of textures and their fractal dimension. He showed that the fractal dimension of 3-D scenes and the fractal dimension of their images are identical, thus providing a direct link between image statistics and underlying physical 3-D structure.

In this paper, we propose a midlevel perceptual mechanism for the identification of the material properties undulation, thickness, and roughness. This mechanism estimates 3-D surface properties from the 2-D frequency representation of images. We will show that spatial frequency analysis is sufficient for rapid identification of material properties. This success relies on the repetitive structure of materials, and we discuss what aspects of the analysis generalize to the recovery of 3-D object shapes.

To investigate the perception of material properties, we chose images of fabrics as stimuli. Fabrics are especially suited for the investigation of material

properties since they are familiar materials that come in a wide variety and have diverse uses. Fabric properties are a function of the nature of the fiber and the structure of the knit or weave. Since these structures vary within a restricted spatial scale, an additional advantage is that most fabrics are examined within a narrow range of distances, similar to what we used in the experiments.

Material property classification

We started our investigation of the perception of material properties with an exploratory experiment in which we asked observers to classify images of materials on four material property dimensions.

Methods

Stimuli

We cropped 256 color images of fabrics to a size of 150×150 pixels (Figure 1). They were presented on a monitor against a black background. The viewing distance was 70 cm. At this distance the images subtended 3.5° of visual angle.

Procedure

We asked observers to rate the materials on four opponent affordance dimensions: soft–rough, flexible–stiff, warm–cool, and water-absorbent–water-repellent, using five-point scales. The scale was presented on the monitor below the images. The computer mouse was used to make a choice. To guide observers, we gave them questions aimed at potential uses: “If you felt this material on your skin, would it feel soft or rough?”; “If you folded or draped this material, would it be stiff or flexible?”; “Would clothes made of this material keep you warm or cool?”; “Would you use this material to repel water or would you use it to absorb water?” The relevant affordance related question was also displayed on the monitor in each trial.

The scale had two levels for each of the two properties described as “very” and “mostly.” The middle scale level could be chosen if neither one of the two property poles was applicable. For example, for the property pair soft–rough, the scale levels were: very soft, mostly soft, neither soft or rough, mostly rough, very rough. The property dimensions were run in a blocked design. All property ratings were done in one session. The image and property sequences were randomized between observers. Each observer repeated the experiment once.

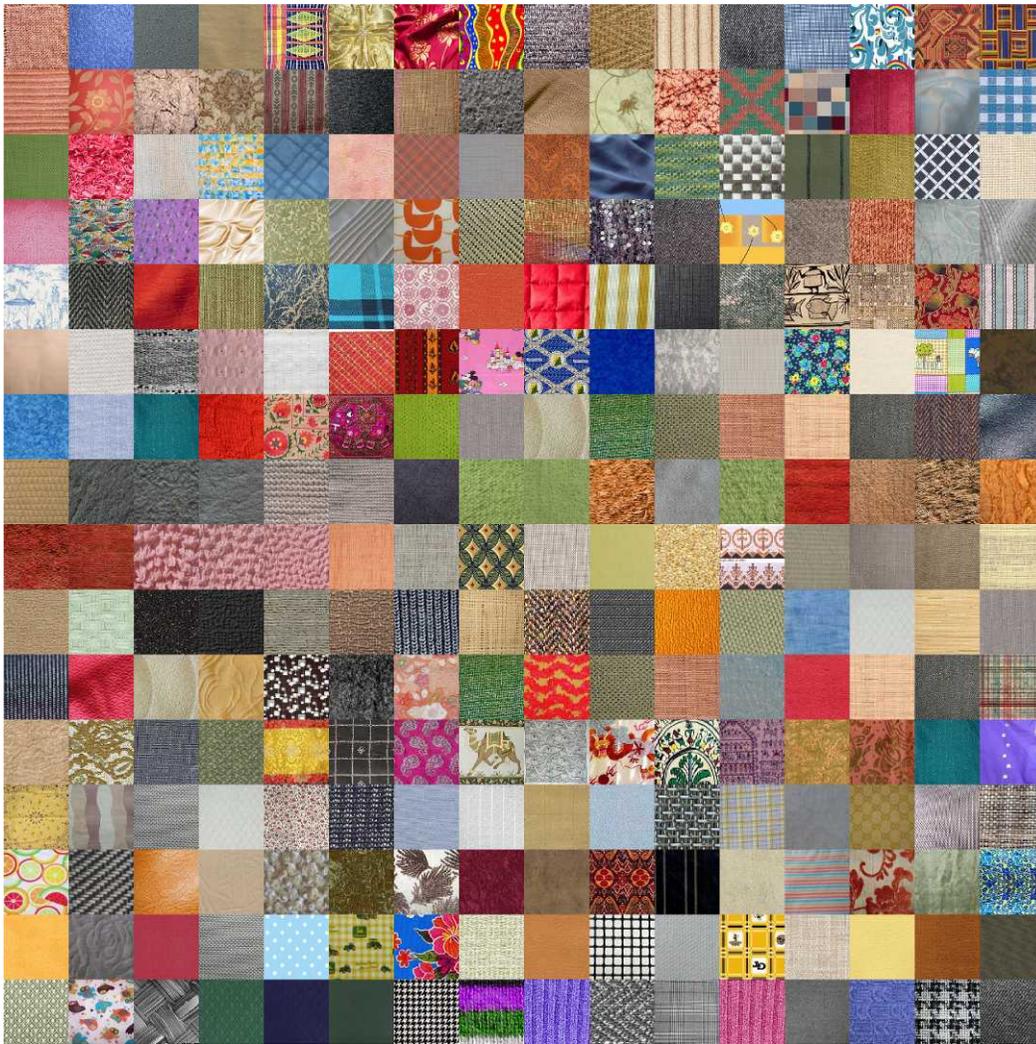


Figure 1. Images of fabrics used in the classification experiment.

Observers

Nine paid naïve observers (five female, four male) participated in the experiment. All experiments presented in this paper were conducted in compliance with the protocol approved by the institutional review board at SUNY College of Optometry and the Declaration of Helsinki.

Data analysis

Since only few images were assigned to the “very” categories, we collapsed the “mostly” and “very” ratings for each property pole. For the subsequent analysis, we used only images that had been rated consistently by an observer in both sessions. Furthermore, we excluded the “neither” category, so that the final data table consisted of the frequencies with which an image had been rated to belong to each of the eight properties. The contingency table containing

these frequencies was then analyzed using correspondence analysis (Hirschfeld, 1935). The analysis was done in R (R Foundation for Statistical Computing, 2013) using the package *ca* (Nenadic & Greenacre, 2007).

Results

The average percentages of consistently rated images for the different properties were (the percentage of images consistently assigned to the “neither” category are given in parentheses.): soft: 29%, rough: 26% (14%); stiff: 22%, flexible: 32% (5%); warm: 33%, cool: 21% (9%); water repellent: 19%, water absorbent: 35% (10%). The agreement between the raters as measured by Krippendorff’s alpha (Gamer, Lemon, Fellows, & Singh, 2012; Krippendorff, 1980) was modest (soft–rough: 0.40; stiff–flexible: 0.36; warm–

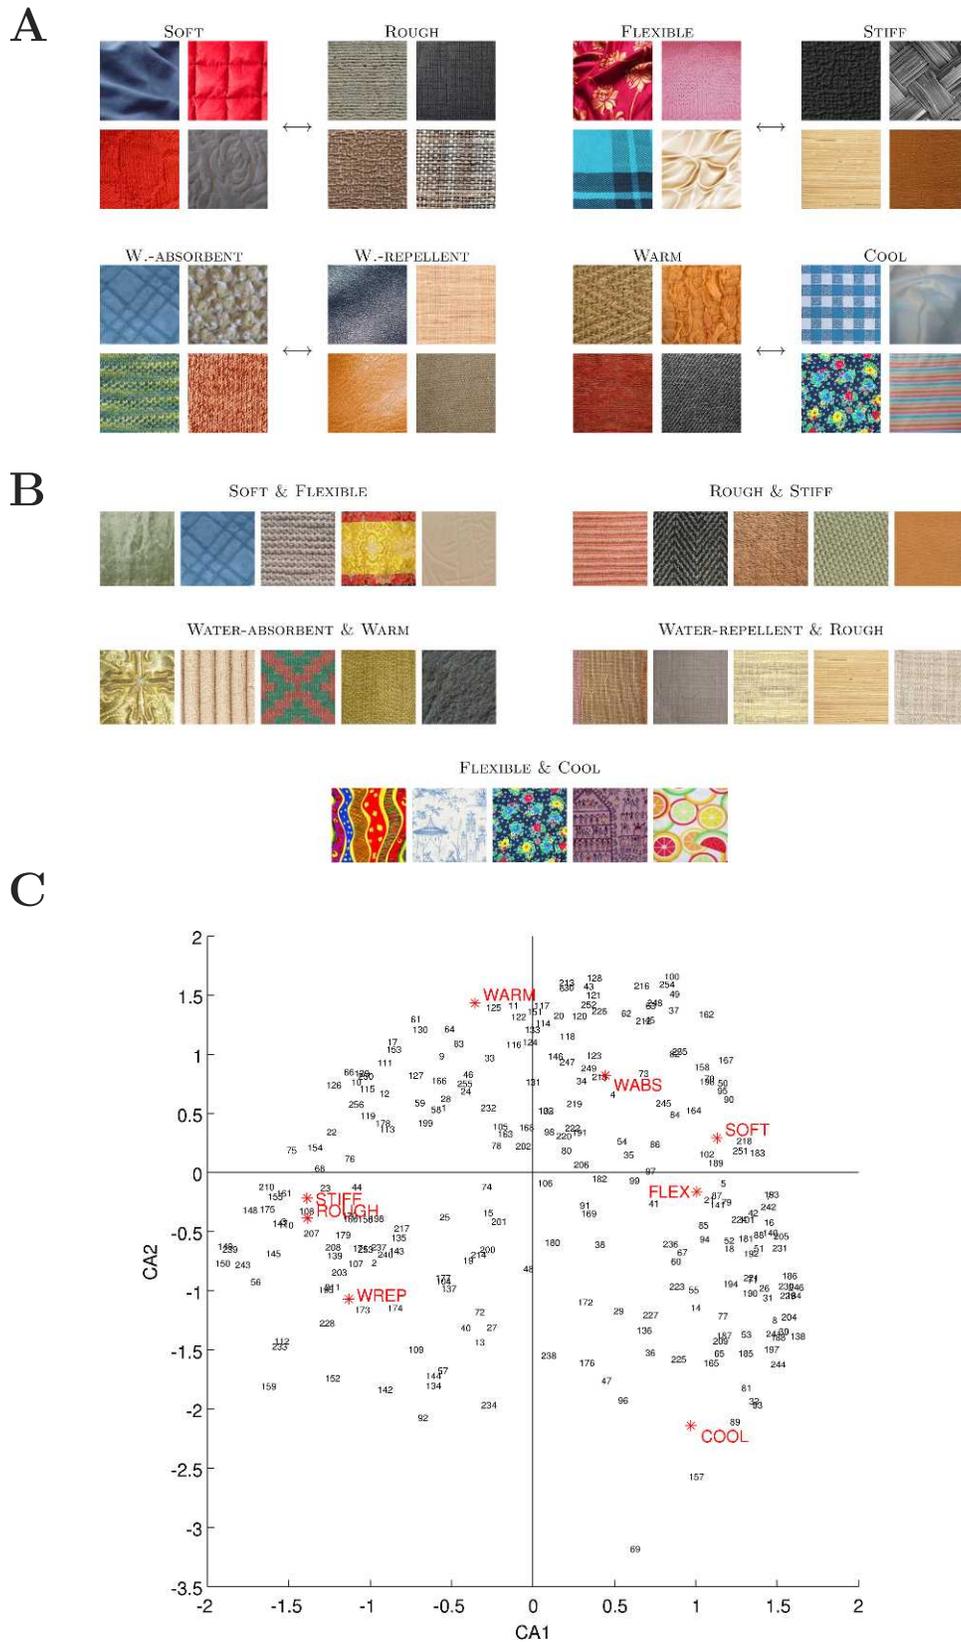


Figure 2. Results of the rating experiment. (A) Examples of materials with highest observer consensus for four opponent material property pairs. (B) Strongest associations across material properties. (C) Results of the correspondence analysis. The two axes are the two orthogonal dimension determined by the correspondence analysis. The locations of the properties on these axes are shown in red (FLEX = flexible, WABS = water absorbent, WREP = water repellent). The numbers refer to the positions of the images in Figure 1 numbered row wise starting from the top left.

cool: 0.33; water-repellent–water-absorbent: 0.25). Figure 2A shows examples of the classifications for each property dimension. The examples illustrate that even images that belong to the same property dimension vary on multiple perceptual dimensions. In addition, some properties were strongly associated with others (Figures 2B, 2C). We tabulated the contingency table of frequencies with which an image had been rated in the two highest scale values for each of the eight properties, and used correspondence analysis to reduce the dimensionality of the problem. Correspondence analysis applies singular value decomposition to the chi-square statistics of the contingency table, and as in principal component analysis, the Eigen functions provide a reduced set of orthogonal dimensions on which the eight properties and 256 images can be represented (Figure 2C). The first dimension (CA1), which explained 57.1% of the total variance in the data matrix, largely coincided with the properties soft and flexible on the positive side, and their opponents, rough and stiff, on the negative side. The second dimension (CA2), which explained 24.8% of the total variance, was closest to the material properties cool on one end and warm on the other. As would be expected, the perceived warmth of a fabric is often a function of its perceived thickness. Absorbent sensibly is closer to warm and soft, while repellent is closer to stiff and rough. Based on the dimensionality reduction, we focused the subsequent image analyses on materials classified as soft or rough and on thin and thick appearing materials. A visual inspection of the soft and rough images in Figure 2A suggests that the size of the dominant structure or pattern is a distinguishing cue between them. Rough materials have a fine grained structure with sharp transitions, whereas soft materials have larger structures with smooth transitions. Within the soft materials there seems to be a further subdivision into a group of thicker looking fabrics, and a group with thinner fabrics that contain broad undulations, probably due to the suppleness of the fabrics. To determine the dominant scale of a material's structure, we analyzed the images' amplitude spectra. Since we found no obvious effect of color in the classification experiment, we used gray-scale versions of the images for the frequency analysis.

Figure 3 shows amplitude spectra of pairs of fabrics chosen to be exemplary of the opposite ends of the undulation (Figure 3A), thickness (Figure 3B), and roughness (Figure 3C) properties. The histograms in the middle column of Figure 3 show the relative energy in various bands of spatial frequencies. The colored parts of the bars indicate the amount of energy by which one of the fabrics exceeds the other one in a given frequency band. The spectra of undulated fabrics contained more energy at low frequencies as compared

to spectra of flat fabrics. Spectra of thick and thin fabrics differed in a frequency band slightly higher than the first band, and spectra of rough fabrics contained more energy at middle frequencies than spectra of soft fabrics.

Image manipulations

To determine whether the amount of energy at certain spatial scales systematically influences the perception of the material properties undulation, thickness, and roughness, we chose three bands of spatial frequencies based on the image analysis: A low-frequency band corresponding to undulations in fabrics covering 2–8 cycles per image (cpi) or 0.57–2.29 cycles/degree (cpd), a frequency band corresponding to the thickness of the weave or knit (8–15 cpi or 2.29–4.28 cpd), and a middle-frequency band corresponding to the roughness of the fabric (23–53 cpi or 6.57–15.14 cpd).

To verify that these bands are related to their corresponding material properties, we assessed the appearance of images as a function of relative energy in the three bands. All image transformations were done in MATLAB® (R2012a, The MathWorks, Natick, MA). All filtering was done by using ideal band-pass filters or notch filters, respectively. To increase or decrease the energy in a frequency band, the frequency band was multiplicatively scaled. To keep the sum of the energy across the amplitude spectrum constant, the remainder of the amplitude spectrum was scaled accordingly. The zero-frequency component was excluded from the scaling procedure. The manipulated images had the same mean as the original images. We constricted scaling to values that did not result in out-of-range pixel values after the inverse Fourier transform.

Figure 4 shows the results of multiplicatively scaling the energy in each of the three bands while keeping the overall energy constant. The icons in the left-most column of Figure 4 indicate the spatial frequency bands. Increasing the energy in the low-frequency band (Figure 4A, Movie 1) inflates the quilt, whereas decreasing the energy deflates it. The three-dimensional appearance of the quilt is largely due to shading variations that are generally gradual, so the energy is concentrated at low spatial frequencies. Increasing the energy in the second frequency band (Figure 4B, Movie 2) increases the thickness of the weave, whereas decreasing the energy results in a flatter, thinner appearance. Increasing the middle to high-frequency energy (Figure 4C, Movie 3) leads to a coarser or rougher texture, while decreasing the energy in this frequency range results in a smoother texture. Varying

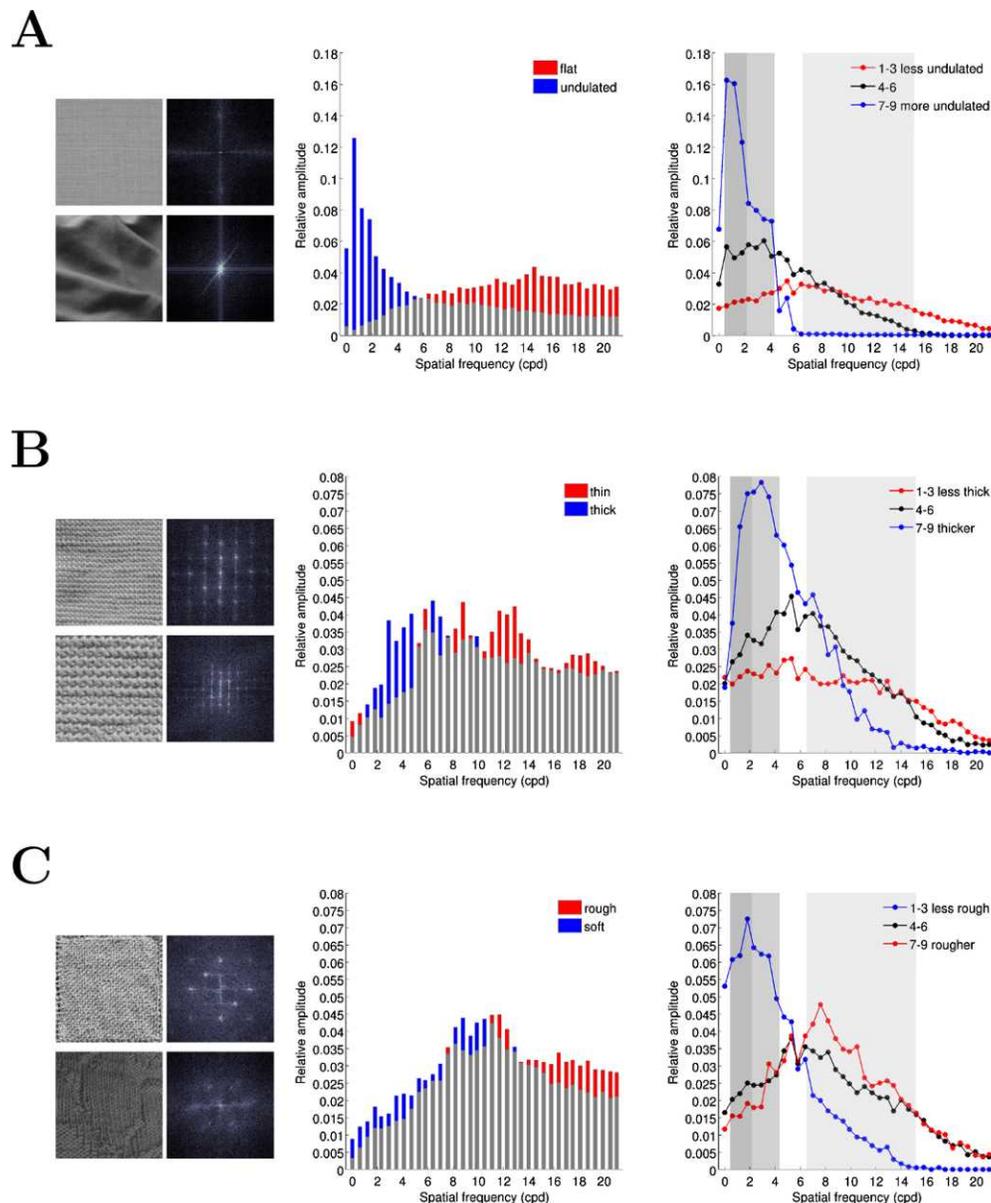


Figure 3. Comparisons of amplitude spectra for opponent material properties summarized by spatial frequency histograms, and results of nine-level property ranking task for three observers. (Left column) Fabric images with their amplitude spectra. (Center column) Histograms of amplitude distributions across spatial frequencies. The colored parts of the bars indicate the amount by which one image exceeds the other. (Right column) Curves show the median relative energy at different frequencies for images sorted to nine levels, collapsed into three categories (see Appendix A, Figures S1A1–C3 for detailed results). (A) Flat (top) versus undulated (bottom), (B) Thin (top) versus thick (bottom), and (C) Rough (top) versus soft (bottom) fabrics.

the relative energy of a frequency band influences how much structures at a certain spatial scale contribute to the overall appearance of the material. It does not alter existing structures or creates new structure. If a material's original spectrum has no structure in a band, multiplying the energy in this band will not lead to the desired appearance changes, manipulating, e.g., the amplitude spectrum of white noise will in general not result in an appearance change that is related to a change in material property (Appendix D, Figure S4B).

However, the frequency-band analysis suggests a method to transfer qualities across materials. Figure 4D shows how the soft and flexible appearance conveyed by folds can be transferred to a material originally rated as rough and stiff. For the folded material, comparing the randomized phase image (Φ_R) to the whitened amplitude image (A_W) demonstrates that the amplitude spectrum determines the volume of the undulations, while the phase component determines the shapes and locations of the folds. A similar

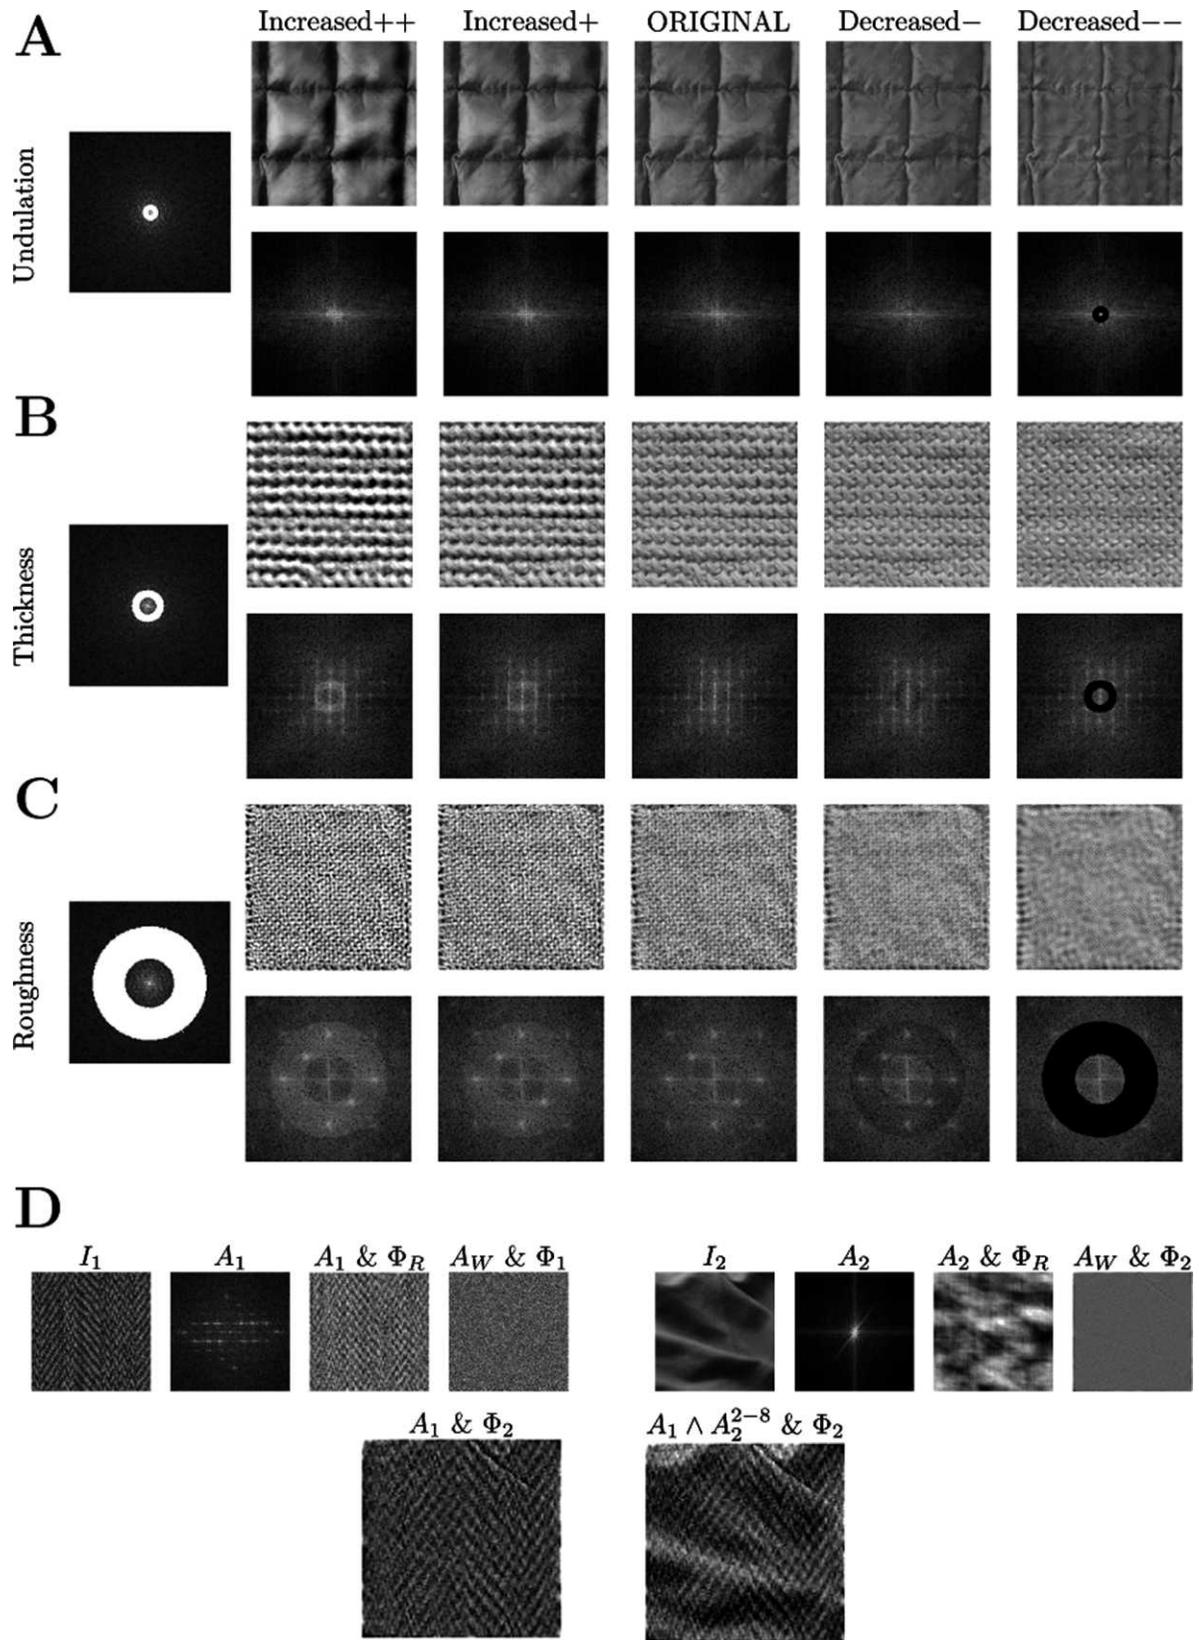
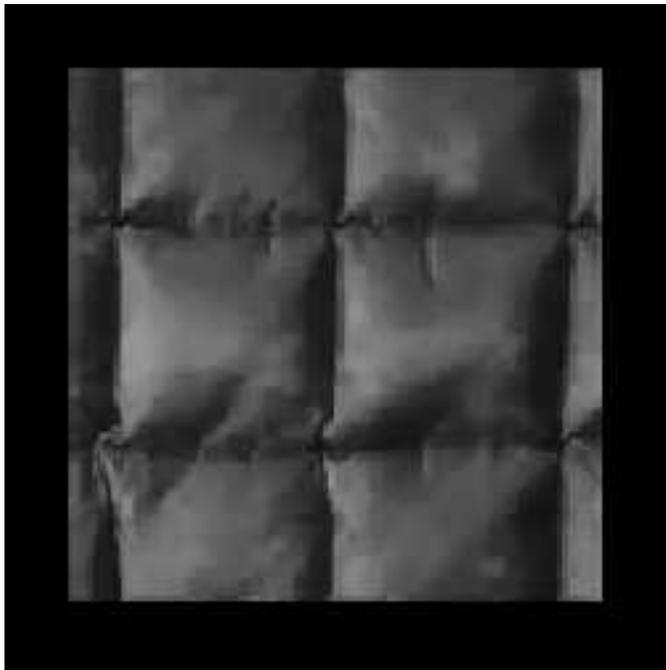


Figure 4. Original and manipulated images and their amplitude spectra. The middle column shows the original images, the first and second column show images with increased energy in the frequency bands, and the fourth and fifth column show images with decreased energy. (A) undulation band, (B) thickness band, (C) roughness band (see also Movies 1–3). (D) Transfer of properties between materials by using structures contained in the frequency band from 2–8 cpi.



Movie 1. Effect of increasing and decreasing the energy in the frequency band 0.57–2.29 cpd on volume perception.

comparison for the flat textured image reveals that the amplitude spectrum determines its dominant texture. When the phase spectrum of the folded material replaces the phase spectrum of the textured material, lines are seen at the locations of the folds. Now if the undulation band (2–8 cpi) of the folded amplitude spectrum replaces the corresponding band in the textured amplitude spectrum, the textured material appears softly folded.

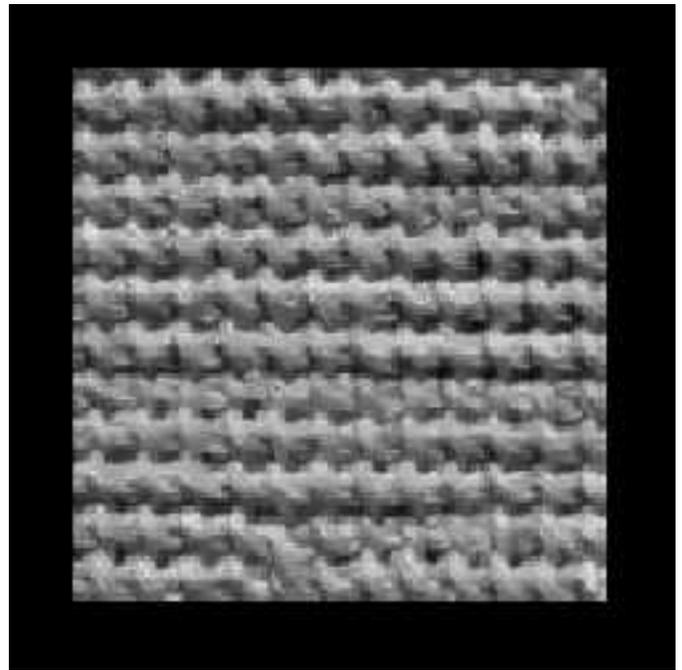
Adaptation to band-limited noise

To determine whether the variations in the magnitude of energy at certain spatial scales are sufficient to influence the perception of material properties, we tested whether adaptation to each of the three frequency bands can alter the perception of their correlated property. In particular, we tested whether adapting to a specific band of spatial frequencies decreases the perceived magnitude of the associated material property.

Methods

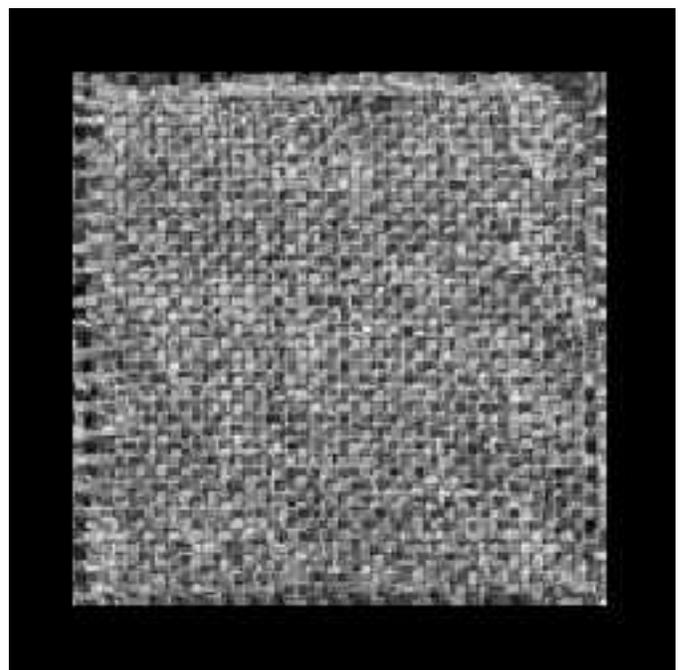
Stimuli

We used gray-scale versions of images that had been classified consistently in the material property classification experiment as stimuli. The stimuli are shown in



Movie 2. Effect of increasing and decreasing the energy in the frequency band 2.29–4.28 cpd on thickness perception.

Figure 5. For each of the three frequency bands, we used images of two fabrics (middle column of Figure 5), plus four versions of each of these images with increased (first and second column in Figure 5) or decreased (fourth and fifth column in Figure 5) relative energy in the specific frequency band. The images in the



Movie 3. Effect of increasing and decreasing the energy in the frequency band 6.57–15.14 cpd on roughness perception.

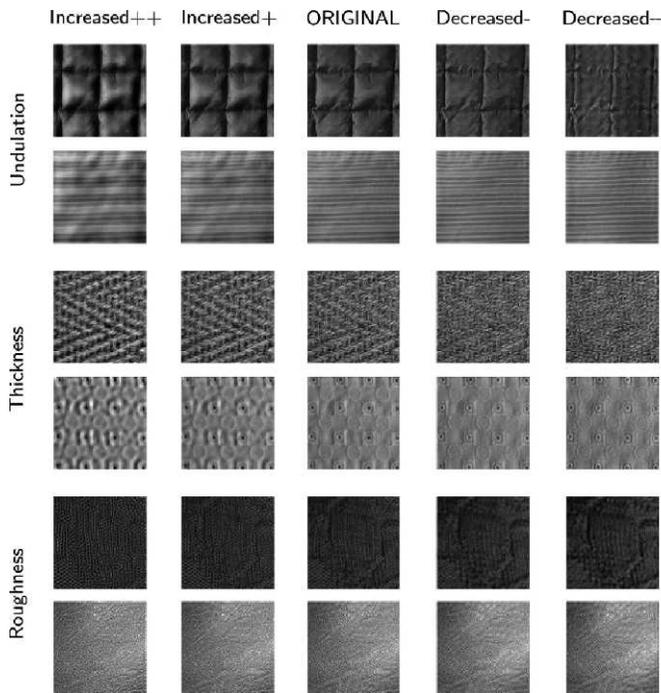


Figure 5. Images used in the adaptation and distance experiment. The middle column shows the original images. The first and second column depict manipulations of the images with increased energy in the frequency bands, and the fourth and fifth column show manipulations of the images with decreased energy in the frequency bands. ++ and -- indicate the maximal possible increase or decrease without having to correct for out-of-range pixel value in the resulting images; + and - indicate versions of the images intermediate to the original and the maximally changed images.

first (++) and the last column (--) show manipulations resulting from the largest increase and decrease, respectively, that was possible without having to correct for out-of-range pixel values. The two other frequency band manipulations (+ and -) represent manipulations intermediate to the original and the maximal possible manipulations. The total energy for each image was kept constant as was the mean of the images. The viewing distance was 70 cm. At this distance the images (150×150 pixels) subtended 3.5° of visual angle.

Noise

For the adapting stimuli, we used band pass-filtered dynamic white noise, and the corresponding notch-filtered dynamic white noise, containing all frequencies except for those in the band. If adaptation to the band of spatial frequencies attenuates perceived magnitude of a material property, then adapting to all frequencies except for those in the band should have the opposite effect. For each frequency band, we generated 120

images of isotropic white noise, and from each of these images we created both a band pass-filtered version and a notch-filtered version. The same ideal band-pass filters were used that had been used for the image manipulations in Figure 5. Inverting them produced the corresponding ideal notch filters. Each image in the sequence was shown for 250 ms. The order of the noise images in the sequences was randomized for each observer. The luminance of the band pass-filtered and notch-filtered noise patches was equated for mean and standard deviation, and they were the same size as the material images.

Procedure

In each trial of this experiment, we presented a test stimulus paired with one of five comparison stimuli. The test stimulus was always one of the original images (middle column of Figure 5). The set of comparison stimuli consisted of the four manipulated versions of the test stimulus (++, +, -, --) and the test stimulus itself. The test and comparison stimuli were always presented for 800 ms simultaneously aligned horizontally or vertically.

In the baseline condition, the test stimulus was presented simultaneously with each of the five comparison stimuli, and on each trial observers picked which of the two materials had more undulation, was thicker, or was rougher, respectively.

In the adaptation condition, observers repeated the measurements of the baseline condition but before the presentation of each pair of stimuli they adapted to band pass-filtered dynamic white noise presented at the location of the test stimulus, and to the complementary notch-filtered dynamic white noise presented at the location of the comparison stimulus (Figure 6A). Thus, the test stimulus was always spatially aligned with the band pass-filtered noise, and the comparison stimuli were always spatially aligned with the notch-filtered noise.

Each test stimulus was presented 10 times together with one of the manipulated comparison stimuli and 20 times together with itself as comparison stimulus. The different frequency bands/material properties were tested in different blocks. Before each trial, a message on the screen instructed the observers which material property they had to judge. The two different images belonging to each of the three frequency bands were tested blocked within the same frequency/material property block. The first adaptation phase after the beginning of a block lasted for 60 s, while subsequent adaptation phases between trials lasted for 10 s. After the end of the adaptation phase there was a gap of 250 ms before the presentation of the image pairs. Between blocks there was a pause of 2 min. To further reduce carry-over effects between successive

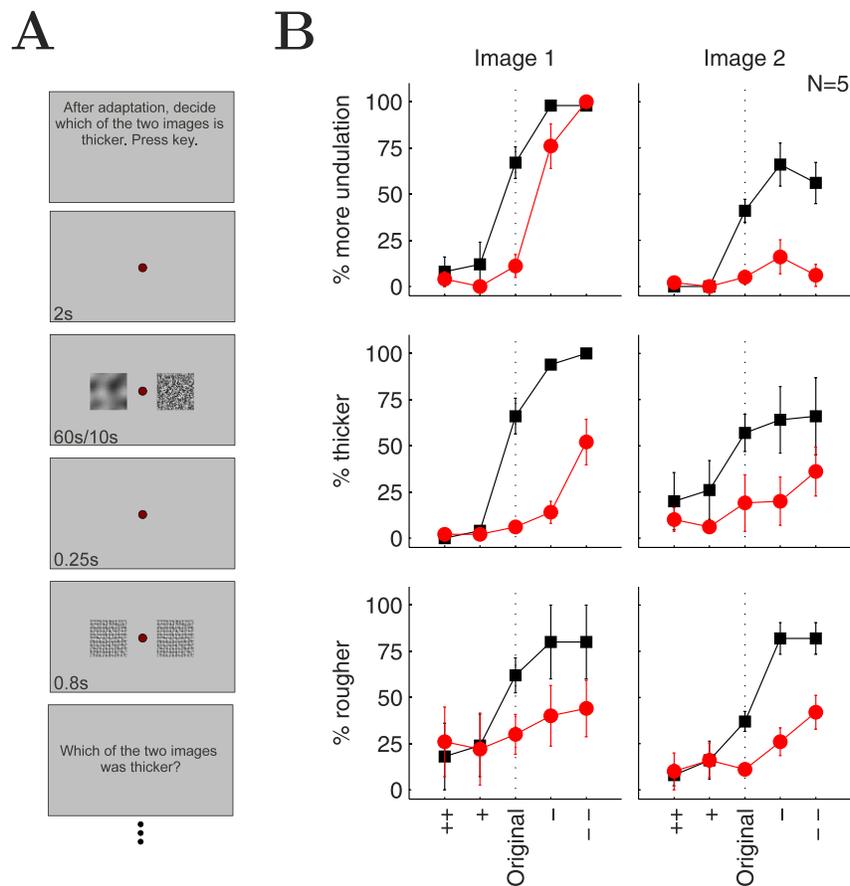


Figure 6. Experimental sequence and results of the adaptation experiment. (A) Different frequency bands were tested in different blocks. Across blocks, the location of the noise patches and stimuli was alternated between left and right of the fixation point, and above and below. Initial adaptation was 60 s, with 10 s top ups. The stimuli were presented for 0.8 s. (B) Baseline (black) and postadaptation (red) psychometric curves for material property comparisons of two fabrics per frequency band, averaged across five observers. The y axis shows the percentage of trials in which the original image (test stimulus) was seen as being rougher, thicker, and more undulated, respectively, than the images (comparison stimuli) indicated on the x axis. Error bars show \pm one SEM.

blocks, the positions of the noise patches and stimuli were alternated between a presentation to the left and right, and a presentation above and below of the fixation point. All presentation sequences (frequency bands, images, positions of the adaptors) were randomized between observers. Some of the observers started with the adaptation condition, others started with the baseline condition. Except for the adaptation phase, the experimental sequence was the same in the baseline and the adaptation condition. The baseline and the adaptation condition were measured on different days.

Observers

Five paid observers (two female, three male) participated in the experiment. Three of them had already participated in the classification experiment.

Data analysis

For the data analysis, we used the percentage of trials in which the test stimulus was seen as being more undulated, thicker, and rougher, respectively, for each combination of test and comparison stimuli. The data were averaged across observers. Our main interest was the effect of adaptation on the responses to the test stimuli when the comparison stimuli were identical to the test stimuli, i.e., when the original images were compared to themselves.

Results

The psychometric curves for the baseline condition (Figure 6B, black lines) show that the image manipulations had the expected effect on the material property judgments. The original images were seen as

less undulated, thinner, or softer than the ++ and + comparison stimuli, and they were seen as more undulated, thicker, or rougher than the -- and - comparison stimuli. The psychometric curves for the adaptation condition (Figure 6B, red line) show that after adaptation, observers' judgments of undulations, thickness, and roughness were depressed. The effect is particularly clear in comparing the pre-versus postadaptation results for the original images (dotted line in Figure 6B). A repeated-measures analysis of variance (ANOVA) with factors adaptation and material property for the data averaged across the two images for each property showed a significant effect of adaptation, overall, $F(1, 5) = 167.81$, $p < 0.001$, and separately for undulation, $F(1, 5) = 203.95$, $p < 0.001$, thickness, $F(1, 5) = 94.62$, $p < 0.001$, and roughness, $F(1, 5) = 11.35$, $p = 0.03$. There was a significant effect of image, $F(1, 5) = 32.57$, $p < 0.005$, and a significant interaction between image and adaptation, $F(1, 5) = 32.00$, $p < 0.005$. For roughness, the effect of image was also significant, $F(1, 5) = 44.51$, $p = 0.003$.

Retinal versus material spatial frequency

We have expressed frequency bands in retinal spatial frequencies (cpd), but because all measurements were done at one distance, they could equivalently have been expressed in material spatial frequency (cpi), a concept analogous to object spatial frequencies (Burbeck, 1987). A change in distance alters retinal spatial frequency but leaves the material spatial frequency unchanged. To determine whether the perceived material properties are determined by material spatial frequencies or retinal spatial frequencies, we conducted a control experiment in which we presented images of materials at three different distances.

Methods

Stimuli

We used the same set of stimuli as in the adaptation experiment (Figure 5). The original images again served as test stimuli while the comparison stimuli consisted of the four manipulated images (++ , + , - , --) and the original image.

Procedure

We varied the viewing distance by using two matched CRTs, one monitor displaying the test stimulus at 33, 66, or 132 cm from the observer, and a

second monitor simultaneously displaying the comparison stimulus at 66 cm from the observer (Figure 7A). The sizes of stimuli for the three distances subtended 7.5°, 3.8°, and 1.9° of visual angle, respectively. The stimuli were presented for 2 s in order to give the observers time to inspect the fabrics on both monitors. Each combination of the test stimuli and the manipulated comparison stimuli was presented 10 times, and the test stimulus was presented 20 times with itself as comparison stimulus. The different material properties were blocked. The sequence of viewing distances and the sequence of material properties were balanced between observers. The judgments for different distances were done in different sessions with at least a day between sessions. On each trial, observers reported whether the test stimulus was more undulated, thicker, or rougher than each of the five comparison stimuli.

Observer

Measurements were made on three uninformed observers (one female, two male) who had not participated in any of the previous experiments.

Results

For the data analyses, we used the percentage of trials in which the test stimulus was seen as being more undulated, thicker, and rougher, respectively, for each combination of test and comparison stimuli. The data were averaged across observers. Our main interest was the effect of distance on the responses to the test stimulus on the test monitor when the comparison stimulus on the comparison monitor was identical to the test stimulus, i.e., when the original images were compared to themselves.

Figure 7B shows that at best there was a weak overall effect of distance. As in the adaptation experiment only the results for the original images are used in the data analysis. A repeated-measures ANOVA with distance and property as factors, for data averaged across the two images, was just significant, $F(2, 6) = 7.63$, $p = 0.043$, due to a weakly significant effect of distance for undulation, $F(2, 6) = 7.60$, $p = 0.043$, but not for thickness, $F(2, 6) = 0.64$, $p = 0.572$, or roughness, $F(2, 6) = 1.49$, $p = 0.328$. The interaction between distance and image was significant, $F(2, 6) = 16.35$, $p = 0.001$. The distance effects for the properties thickness, $F(2, 6) = 0.64$, $p = 0.572$, and roughness, $F(2, 6) = 1.49$, $p = 0.328$, were not significant. The interaction between distance and image for roughness was significant, $F(2, 6) = 7.39$, $p = 0.045$.

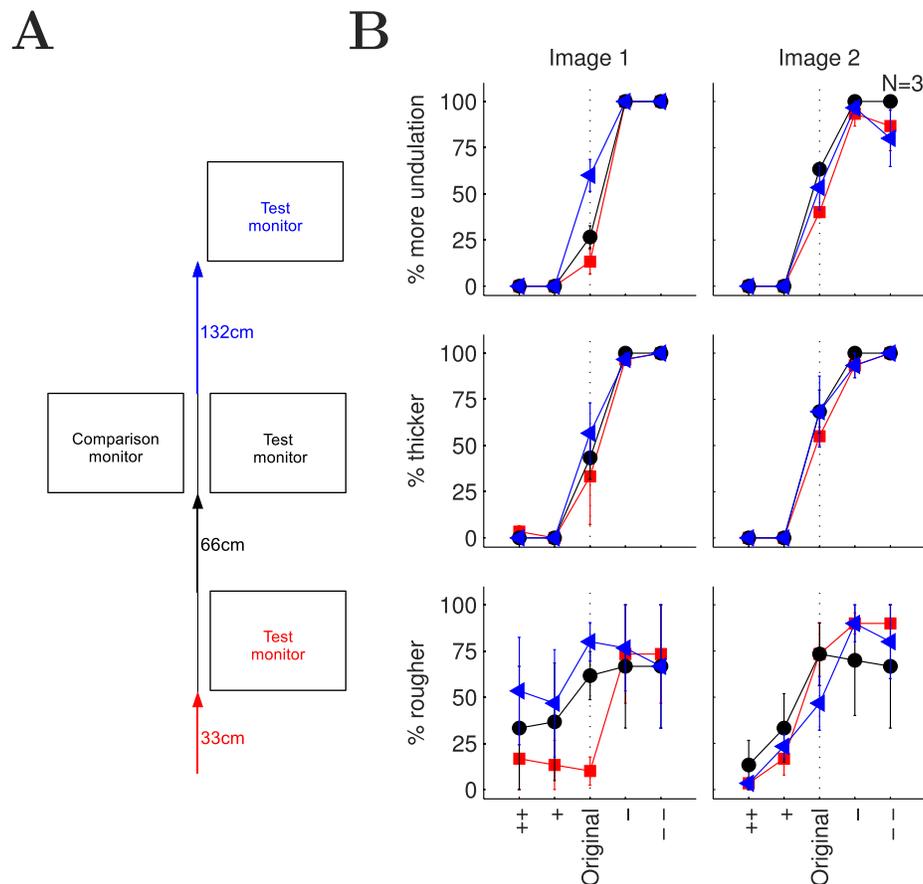


Figure 7. Experimental setup (A), and results of the distance experiment (B). The results for the different images are shown in separate columns. The x axis indicates the type of comparison stimulus shown on the reference monitor. The y axis shows the percentage of trials in which the test stimulus presented on the test monitor was chosen as being more undulated, thicker, or rougher, respectively, than the comparison stimuli. Colors and symbols indicate the different conditions: both monitors at the same distance (black, circles), test monitor closer to the observer (red, squares), test monitor farther from the observer (blue, triangles). Symbols indicate the mean across three observers. Error bars show \pm one SEM.

Discussion

The results for undulation and thickness in Figure 7B show no significant effect of distance on material perception in the tested range. There was predictable variability in the data for roughness. The two images used for roughness show a tendency to be affected in opposite ways by the increase in the viewing distance. For the first image, distance causes an increase in perceived roughness. This image, which was originally rated as soft, has the critical variations at low to middle frequencies. These are shifted to higher retinal spatial frequencies with increasing distance, and that might have resulted in a rougher appearance. For the second image, there is a nonsignificant tendency for perceived roughness to decrease systematically for the largest viewing distance. The dominant variations in the second image are already in a high frequency region, so increasing the distance may have moved them beyond the window of visibility, thus increasing the relative

effect of lower frequencies, and that could have resulted in a softer appearance. Overall, material judgments remain stable over a range of distances from which an observer would commonly examine materials. This suggests that visual inferences of material properties are more likely to be based on estimated material spatial frequencies than on retinal frequencies. This finding is in accordance with data from experiments investigating spatial-frequency discrimination (Burbeck, 1987), spatial frequency memory masking (Bennett & Cortese, 1996), size constancy (Blakemore, Garner, & Sweet, 1972), recognition memory for shapes (Milliken & Jolicoeur, 1992), and for the extraction of upper case letters from noise (Parish & Sperling, 1991). It also seems to be in line with our everyday experience where material properties do not change massively with viewing distance. There are of course sensory limitations to this that we intend to investigate more closely in future experiments

Agreement between material property rankings and frequency bands

To strengthen the generality of the frequency-based analysis of material properties, we conducted a ranking experiment to derive the frequency bands from experimental data. In this experiment we used printouts of some of the images shown in Figure 1. This experiment was less well controlled than the monitor based experiments, but had the advantage that observers made relative judgments of the material properties without being required to explicitly label a material as belonging to a certain property category. The central question was whether the relative judgments were correlated with the amplitude distributions in the three frequency bands.

Methods

Stimuli

Gray-scale images of 161 of the fabrics (Figure 1, leaving out only the printed fabrics) were printed on white letter-sized paper using a standard laser printer. The printed images had a size of 4.5×4.5 cm and thus were, when held at arm's length, approximately of the same retinal size as the stimuli in the monitor based experiments.

Procedure

The sequence of the images was randomized in the stack of papers given to the observers. The observers' task was to sort the images independently in nine point classes from "least" to "most" undulated, thick, and rough. The sorting was carried out on a large table so that images could be seen simultaneously in order to facilitate direct comparisons between them. No specific instructions regarding the viewing distance and the properties were given to the observers. However, to illustrate the properties we showed samples of real fabrics to the observers: A fabric lying flat on the desk, and the same fabric with folds to illustrate the undulation property, a thinner and a thicker fabric to illustrate thickness, and a softer and rougher fabric to illustrate roughness.

Observers

Material rankings were done by three observers (one female, two male). One of the observers had already participated in the classification and the adaptation experiment, one observer had participated in the

	W_t	χ^2	df	p
Undulation	0.619	297	160	< 0.01
Thickness	0.651	312	160	< 0.01
Roughness	0.692	332	160	< 0.01

Table 1. Interrater concordance for material property rankings. W_t is Kendall's coefficient of concordance (W) corrected for ties within raters.

distance experiment, and the third observer was uninformed as to the purpose of the experiment.

Results

Table 1 shows the inter-rater reliability as measured by Kendall's coefficient of concordance (W), which is a measure of agreement for ordinal data among several judges who are assessing a set of n objects (Kendall, 1948). It ranges from zero (no agreement) to one (complete agreement).

Table 2 shows the correlations (Kendall's rank correlation coefficient τ) between the rankings for the three properties separately for the three observers. The rank cross-correlations between the different properties were below 0.3 for all observers, with the undulation ratings being negatively correlated with the roughness ratings.

To examine the correlation between a material's rank for a property and its amplitude distribution, Figure 3 (right column) shows the amplitude distributions of materials assigned to the different scale levels separately for the different properties. The band-wise amplitudes were determined from thresholded amplitude spectra. Since we found that most of an image's dominant structure is contained in the top 30% of amplitudes, only those top 30% of the amplitude were used. For better visibility, the median energy distribution across observers' rankings was collapsed into three scale levels (Figures S1A1–S1C3 in Appendix A show the rankings separately for each observer and property). In Figure 3, for all three properties the "most" levels (Scale Levels 7–9) exhibited the highest relative amplitudes in the three bands identified in previous sections, and the "least" levels (Scale Levels 1–3) exhibited the lowest relative amplitudes in those bands, thus corroborating our choice of frequency bands.

	Undulated ~ Thick	Undulated ~ Rough	Thick ~ Rough
Obs. 1	0.162**	−0.185**	0.202**
Obs. 2	0.248**	−0.096	0.257**
Obs. 3	0.132*	−0.293**	0.160**

Table 2. Rank correlations (Kendall's rank correlation coefficient τ) between material property rankings.

Validation of frequency bands from ranking data

The perceptual and adaptation results justify linking the material properties of roughness, thickness, and undulation to specific frequency bands. However, the frequency bands were chosen by visual inspection, so the question remains whether they can also be justified on statistical principles. To capture the most general linear relationship between property rankings and amplitude spectra, we used canonical correlation analysis (CCA). For two sets of variables, CCA finds mutually orthogonal functions consisting of pairs of linear combinations (variates) of each of the variables that have maximum correlation with each other (see Appendix B.1). The number of functions is limited to the cardinality of the smaller set, so that three functions encapsulate all of the 3×37 correlations across the 161 images between the dependent variables consisting of the three median property rankings across observers and the independent variables consisting of the relative image amplitudes divided into 37 equal-width frequency bands. Canonical correlation is more general than multivariate regression and linear discriminant analysis in allowing linear combinations of dependent as well as independent variables, and is invariant to affine transformations of the variables, which is useful when dealing with ranking measurements.

Figure 8A shows the correlations (loadings) of the independent variables with the independent variate, and Figure 8B shows the squared loadings, i.e., the proportion of variance a variable shares with the variate. The first to third rows show the loadings for the first to third canonical functions. The loadings of the dependent variables on the dependent variate are shown in Table 3. The correlation between the first canonical variates is 0.805 ($p < 0.01$) (Appendix B.1, Table S2). The first dependent variate has a high positive correlation with the roughness rankings, and a strong negative correlation with the undulation rankings. In agreement with the bands (shaded areas) that we derived from inspection, the loadings of the first independent variate are highly negative in the undulation band and highly positive in the roughness band. This variate accounts most for variance in the frequencies belonging to the undulation band. The correlation between the second canonical variates is 0.564 ($p < 0.05$). The independent variate has a high positive correlation with the thickness ratings, and the dependent variate has the highest loadings and squared loadings in our thickness band. The correlation between the third canonical variates is 0.520 ($p > 0.05$). All rankings correlate positively with the third dependent variate, but the correlation with roughness is highest. Corresponding to that, the third independent

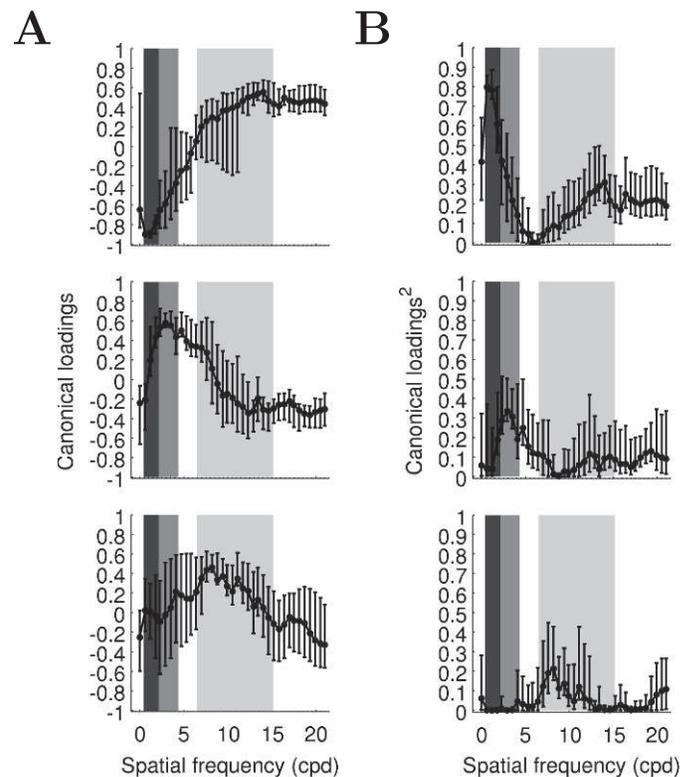


Figure 8. Results of canonical correlation analysis. (A) Loadings for the independent variable (amplitudes in 37 frequency bands for 161 images). (B) Squared loadings for the independent variable. Error bars denote 95% confidence intervals resulting from 1,000 replications of the canonical correlation analysis. In each bootstrap, a new sample was created by sampling with replacement from the data set.

variate has fairly low loadings, but the highest absolute values are in our roughness band, and the selectivity is more apparent in the squared loadings. Overall, the CCA analysis validates our choice of frequency bands. If the loading functions were used in image manipulations, the effects would be similar to using our frequency bands. The opponent relationship between roughness and undulation may be the result of a physical constraint, or a bias in the sample not having many rough appearing fabrics with folds. Note that in the image manipulations (Figure 4), when we increased the energy in the undulation band while keeping the total energy constant, it effectively reduced the energy in the other bands.

If a larger set of material properties could be shown to depend on image spatial frequency, CCA could be

	1. Variate	2. Variate	3. Variate
Undulation	−0.860	−0.074	0.505
Thickness	−0.055	0.816	0.575
Roughness	0.781	−0.120	0.612

Table 3. Loadings of the dependent variables.

used as the foundation of extracting optimal frequency bands by running ranking experiments on a much larger sample of materials. In Appendix B.2 we present one possible procedure to derive orthogonal frequency bands for the properties discussed in this paper by applying singular value decomposition to the independent loading functions weighted by the dependent loading functions.

General discussion

The main results of this paper are that material properties, such as roughness, thickness, and undulation, are characterized by specific scales of luminance variations. These results were derived from perceptual scaling of a large sample of fabric images, combined with image analyses and statistical dimensionality reduction. The 2-D luminance variations arise from the 3-D textures of the materials, and are used by the visual system in perceiving local 3-D structure, as confirmed by movies showing that judgments of 3-D roughness, thickness, and undulations vary continuously as a function of relative contrast in corresponding 2-D frequency bands, and are predictably altered by adaptation to dynamic noise at the corresponding scales. The appearance changes that result from the manipulations of the amplitude distributions in the three frequency bands are all caused by changes of the shading components at different spatial scales. This is most obvious for the undulation band, where variations give rise to 3-D structures that extend across large parts of the sample, but it also applies to variations of the mesostructure resulting in changes in apparent thickness and roughness. The perceived material properties are thus functions of the 3-D structures of the materials, which are mainly conveyed by shading cues.

The results of the canonical correlation analysis indicate that the correlation between perceptual ranks and frequency-band amplitudes is not perfect. This may partly be due to the paper sorting experiment being a noisy procedure. However, we are almost certain that there are additional low-level factors, including limitations imposed by the contrast sensitivity for spatial frequencies (Watson & Ahumada, 2005), and cross-band and frequency-dependent cross-orientation masking effects on the saliency of spatial frequencies (Li & Zaidi, 2009). In addition, it remains to be tested whether high-level mechanisms, such as recognition, potentiate specific low-level cues.

Since the earliest stage of cortical visual processing consists of neurons that filter the visual scene in terms of spatial frequencies and orientations (Hawken & Parker, 1987; Schiller, Finlay, & Volman, 1976), it is

not surprising that spatial frequencies play an important role in pre-attentive texture discrimination (Julesz, 1962), and texture matching (Richards & Polit, 1974). More recently, direct scene categorization schemes proposed correlations between specific configurations of power spectra and perceptual scene dimensions such as naturalness and openness (Oliva & Torralba, 2001).

While the visual perception of roughness, thickness, or undulation has not been investigated extensively, the estimation of the roughness of surfaces or terrains from images has long been an important topic of research in machine vision. A wide array of methods has been employed to this end, including spatial frequency analysis. In this context, it has been found that spatial frequency analysis was often inferior to other methods, e.g., statistics derived from gray-tone co-occurrence probabilities (for reviews see Haralick, 1979; Tuceryan & Jain, 1998). However, the focus of this line of research was on the reliable identification of physical structures as, for example, required in remote sensing. Here, we were primarily concerned with material appearance and not with the veridical recovery of surface properties. Surface properties and illumination geometry are conflated in the spatial frequency information. The amplitude distribution changes systematically with changes in pose, scale, and illumination, and that seems correlated with the resulting changes in material appearance. An interesting case is presented by slanted surfaces. When 2-D textures are slanted, the spatial frequencies are increased in the image, orientation flows are created (Li & Zaidi, 2004), and the brightness is reduced for Lambertian surfaces. However, the case is more complicated when 3-D structures are slanted as the structure determines the change in brightness (Nayar & Oren, 1995) and spatial frequency (Dana, Van Ginneken, Nayar, & Koenderink, 1999). To analyze the interaction of pose, scale, and illumination on material perception, we chose images from the KTH-TIPS database (Fritz, Hayman, Caputo, & Eklundh, 2004). In general, slanting materials increased the spatial frequencies in the image at short distances, and the effect is small or even absent for larger distances (Appendix C, Figures S3A–D). Illumination from the side emphasized the finer structure of the fabrics, thus causing a shift to higher image frequencies. Interestingly, the energy peaks for the fabrics were generally located in one of the previously identified frequency bands, and the perceived qualities followed the bands, e.g., when the spatial frequency peak moved to frequencies higher than the roughness band, the material appeared increasingly flat and smooth. Given the systematic interaction between the amplitude distribution and changes in pose, scale and illumination-geometry caused by the complexity of real materials (Anderson, 2011), the amplitude distributions might also provide

cues to the separation of surface-reflectance and illumination when observers can view multiple slants and tilts (Barron & Malik, 2011).

Our results seem to be the first recognition of the role played in shape from shading by relative contrast at different spatial frequencies. Orientation flows created by isophotes have also been shown to play a role in conveying 3-D shape (Dragnea & Angelopoulou, 2005; Kunsberg & Zucker, 2012; Pentland, 1984b; Šára, 1995), and we would probably get greater adaptation if the orientations in the noise were matched to the orientations in each material, but in this case the noise would be specific to each material rather than to general properties. Our frequency band manipulations predict specific properties across materials because they do not alter the orientation structure resident in each material, but it is unlikely that specific orientation flows are associated with specific material properties.

The neural substrate for the proposed frequency analysis remains to be investigated. fMRI studies have implicated fusiform and parahippocampal regions as part of the network that processes surface properties (Cant & Goodale, 2007; Cavina-Pratesi, Kentridge, Heywood, & Milner, 2010), but there have been no physiological studies of how the neural network extracts material structure from the outputs of V1 neurons which are tuned to fairly narrow spatial frequency bands and process retinal images in parallel. Two-dimensional textures have been analyzed and synthesized by filter responses at multiple scales and orientations (De Bonet, 1997; Heeger & Bergen, 1995; Portilla & Simoncelli, 2000) with some success. These methods may not capture all the local detail of structured textures that human observers perceive, but by substituting receptive fields of V1 neurons for the filters, explicit neural models can be built and tested for the scale based estimation of material properties.

The somatosensory system supports fine discrimination of surface textures, comprising at least four major dimensions: roughness versus smoothness, hardness versus softness, stickiness versus slipperiness, and warm versus cool (Hollins, Bensmaïa, Karlof, & Young, 2000). There is a close correspondence between the haptic and visual estimates of the roughness of real surfaces (Bergmann-Tiest & Kappers, 2007). For both senses, the medium material spatial frequency band had the highest correspondence with the observers' roughness orderings, and there was a drop in estimated roughness at high spatial frequencies corresponding to the decrease in sensitivities of both the visual and tactile systems (Bergmann-Tiest & Kappers, 2007; Connor & Johnson, 1992). These results suggest that either both visual and tactile systems directly estimate roughness from the scale of the surface microstructure, or that roughness is exclusively a haptic percept and its visual estimation is based on material recognition and

retrieval of roughness information from memory (Bergmann-Tiest & Kappers, 2007). The effects of haptic feedback can be large on visual percepts, but they seem to be short-lived (Meng & Zaidi, 2011), so whether visual roughness perception is calibrated by long-term haptic experience, still needs to be investigated.

Analyzing and synthesizing natural patterns is essentially an unsolved problem (Mumford & Desolneux, 2010). In computer graphics, fabric rendering, especially in animated sequences is an especially important and difficult problem (Selle, Su, Irving, & Fedkiw, 2009). Some approaches use 3-D texels to achieve realism, but at the cost of computational speed (Durupinar & Gdükbay, 2007). Our material manipulations illustrate a novel and rapid method for altering properties in images of real materials, and for transferring properties across materials. Our methods could also enhance the realism of synthetically generated materials by endowing them with recognizable properties. As the image manipulations (Figure 4) imply, this will be more possible when there is at least some structure at the relevant scales (Appendix D).

Keywords: material perception, spatial frequency, adaptation, image statistics

Acknowledgments

We thank Hal Sedgwick and Stephen Engel for discussions about this work. This work was supported by NEI grants EY07556 and EY13312 to QZ, and DFG Research Fellowship GI 806/1-1 to MG. The main sections of this paper were presented at VSS 2011, 2012, 2013, and ECVF 2012.

Commercial relationships: none

Corresponding author: Martin Giesel.

Email: mgiesel@sunyopt.edu.

Address: Graduate Center for Vision Research, SUNY College of Optometry, New York, NY, USA.

References

- Adelson, E. H. (2001). On seeing stuff: The perception of materials by humans and machines. *Proceedings of the SPIE: Human Vision and Electronic Imaging VI*, 4299, 1–12.
- Anderson, B. L. (2011). Visual perception of materials and surfaces. *Current Biology*, 21, R978–R983.
- Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness.

- Journal of Vision*, 9(11):10, 1–17, <http://www.journalofvision.org/content/9/11/10>, doi:10.1167/9.11.10. [PubMed] [Article]
- Barron, J. T., & Malik, J. (2011). High-frequency shape and albedo from shading using natural image statistics. *2011 IEEE Conference on CVPR*, 2521–2528.
- Beck, J., Prazdny, K., & Ivry, R. (1984). The perception of transparency with achromatic colors. *Perception and Psychophysics*, 35, 407–422.
- Bennett, P. J., & Cortese, F. (1996). Masking of spatial frequency in visual memory depends on distal, not retinal, frequency. *Vision Research*, 36(2), 233–238.
- Bergmann-Tiest, W. M., & Kappers, A. M. L. (2007). Haptic and visual perception of roughness. *Acta psychologica (Amsterdam)*, 124, 177–189.
- Blakemore, C., Garner, E. T., & Sweet, J. A. (1972). The site of size constancy. *Perception*, 1(1), 111–119.
- Binns, H. (1937). Visual and tactual ‘judgement’ as illustrated in a practical experiment. *British Journal of Psychology*, 27, 404–410.
- Breton, P., & Zucker, S. W. (1996). Shadows and shading flow fields. *Proceedings CVPR ‘96 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 782–789.
- Burbeck, C. A. (1987). Locus of spatial-frequency discrimination. *Journal of the Optical Society of America A*, 4(9), 1807–1813.
- Cant, J. S., & Goodale, M. A. (2007). Attention to form or surface properties modulates different regions of human occipitotemporal cortex. *Cerebral Cortex*, 17, 713–731.
- Cavina-Pratesi, C., Kentridge, R. W., Heywood, C. A., & Milner, A. D. (2010). Separate processing of texture and form in the ventral stream: Evidence from fMRI and visual agnosia. *Cerebral Cortex*, 20, 433–446.
- Connor, C. E., & Johnson, K. O. (1992). Neural coding of tactile texture: Comparison of spatial and temporal mechanisms for roughness perception. *Journal of Neuroscience*, 12, 3414–3426.
- Dana, K. J., Van Ginneken, B., Nayar, S. K., & Koenderink, J. J. (1999). Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18, 1–34.
- De Bonet, J. S. (1997). Multiresolution sampling procedure for analysis and synthesis of texture images. *Proceedings of SIGGRAPH*, 97, 361–368.
- De Haan, E., Erens, R. G. F., & Noest, A. J. (1995). Shape from shaded random surfaces. *Vision Research*, 35(21), 2985–3001.
- Doerschner, K., Fleming, R. W., Yilmaz, O., Schrater, P. R., Hartung, B., & Kersten, D. (2011). Visual motion and the perception of surface material. *Current Biology*, 21, 2010–2016.
- Dragnea, V., & Angelopoulou, E. (2005). Direct shape from isophotes. *Proceedings of the ISPRS Workshop Ben-COS05*, 45–50.
- Durupinar, F., & Güdükbay, U. (2007). Procedural visualization of knitwear and woven cloth. *Computers & Graphics*, 31, 778–783.
- Erens, R. G. F., Kappers, A. M. L., & Koenderink, J. J. (1993). Perception of local shape from shading. *Perception & Psychophysics*, 54(2), 145–156.
- Fleming, R. W., & Bühlhoff, H. H. (2005). Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception*, 2, 346–382.
- Fleming, R. W., & Paulun, V. (2012). Goop! On the visual perception of fluid viscosity. *Journal of Vision*, 12(9): 949, <http://www.journalofvision.org/content/12/9/949>, doi:10.1167/12.9.949. [Abstract]
- Fritz, M., Hayman, E., Caputo, B., & Eklundh, J.-O. (2004). The KTH-TIPS database [Database]. Retrieved from <http://www.nada.kth.se/cvap/databases/kth-tips>.
- Gamer, M., Lemon, J., Fellows, I., & Singh, P. (2012). irr: Various coefficients of interrater reliability and agreement (R package version 0.84) [Computer software]. Retrieved from <http://CRAN.R-project.org/package=irr>.
- Gibson, J. J. (1986). *The ecological approach to visual perception*. New York, NY: Psychology Press.
- Haralick, R. M. (1979). Statistical and structural approaches to texture. *Proceedings of the IEEE*, 67(5), 786–804.
- Hawken, M. J., & Parker, A. J. (1987). Spatial properties of neurons in the monkey striate cortex. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 231(1263), 251–288.
- Heeger, D. J., & Bergen, J. R. (1995). Pyramid-based texture analysis/synthesis. *Proceedings of SIGGRAPH*, 1995, 229–238.
- Hirschfeld, H. O. (1935). A connection between correlation and contingency. *Proceedings of the Cambridge Philosophical Society*, 31, 520–524.
- Ho, Y. X., Landy, M. S., & Maloney, L. T. (2008). Conjoint measurement of gloss and surface texture. *Psychological Science*, 19, 196–204.
- Hollins, M., Bensmaïa, S., Karlof, K., & Young, F. (2000). Individual differences in perceptual space for tactile textures: Evidence from multidimen-

- sional scaling. *Perception & Psychophysics*, 62(8), 1534–1544.
- Horn, B. K. P. (1970). *Shape from shading: A method for obtaining the shape of a smooth opaque object from one view* (Technical Report MAC-TR-79). Massachusetts Institute of Technology.
- Ikeuchi, K., & Horn, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17, 141–184.
- Julesz, B. (1962). Visual pattern discrimination. *IRE Transactions on Information*, 8(2), 84–92.
- Kendall, M. G. (1948). *Rank Correlation Methods*. London, UK: Charles Griffin.
- Koenderink, J. J., & Doorn, A. J. (1996). Illuminance texture due to surface mesostructure. *Journal of the Optical Society of America A*, 13, 452–463.
- Koenderink, J. J., Doorn, A. J., Dana, K. J., & Nayar, S. (1999). Bidirectional reflection distribution function of thoroughly pitted surfaces. *International Journal of Computer Vision*, 31, 129–144.
- Krippendorff, K. (1980). *Content analysis: An introduction to its methodology*. Beverly Hills, CA: Sage.
- Kunsberg, B., & Zucker, S. W. (2012). The differential geometry of shape from shading: Biology reveals curvature structure. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 39–46.
- Lesch, M. F., Chang, W.-R., & Chang, C.-C. (2008). Visually based perceptions of slipperiness: Underlying cues, consistency and relationship to coefficient of friction. *Ergonomics*, 51, 1973–1983.
- Li, A., & Zaidi, Q. (2004). Three-dimensional shape from non-homogeneous textures: Carved and stretched surfaces. *Journal of Vision*, 4(10):3, 860–878, <http://www.journalofvision.org/content/4/10/3>, doi:10.1167/4.10.3. [PubMed] [Article]
- Li, A., & Zaidi, Q. (2009). Release from cross-orientation suppression facilitates 3D shape perception. *PLoS One*, 4(12), e8333.
- Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. *Current Biology*, 22, 1909–1913.
- Meng, X., & Zaidi, Q. (2011). Visual effects of haptic feedback are large but local. *PLoS One*, 6(5), e19877.
- Metelli, F. (1985). Stimulation and perception of transparency. *Psychological Research*, 7, 185–202.
- Milliken, B. M., & Jolicoeur, P. (1992). Size effects in visual recognition memory are determined by perceived size. *Memory & Cognition*, 20(1), 83–95.
- Motoyoshi, I., Nishida, S., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature*, 447, 206–209.
- Mumford, D., & Desolneux, A. (2010). *Pattern theory: The stochastic analysis of real world signals*. Natick, MA: A K Peters.
- Nayar, S. K., & Oren, M. (1995). Visual appearance of matte surfaces. *Science*, 267, 1153–1156.
- Nenadic, O., & Greenacre, M. (2007). Correspondence analysis in R, with two- and three-dimensional graphics: The ca package. *Journal of Statistical Software*, 20, 1–13.
- Nishida, S., & Shinya, M. (1998). Use of image-based information in judgments of surface-reflectance properties. *Journal of the Optical Society of America A*, 15, 2951–2965.
- Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision*, 42, 145–175.
- Olkkonen, M., & Brainard, D. H. (2011). Joint effects of illumination geometry and object shape in the perception of surface reflectance. *I-Perception*, 2, 1014–1034.
- Parish, D. H., & Sperling, G. (1991). Object spatial frequencies, retinal spatial frequencies, noise, and the efficiency of letter discrimination. *Vision Research*, 31(7), 1399–1415.
- Pentland, A. P. (1982). *The visual inference of shape: computation from local features* (Unpublished Ph.D. thesis). Massachusetts Institute of Technology, Cambridge, MA.
- Pentland, A. P. (1984a). Fractal-based description of natural scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6), 661–674.
- Pentland, A. P. (1984b). Local shading analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(2), 170–187.
- Portilla, J., & Simoncelli, E. P. A. (2000). A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision*, 40, 49–71.
- R Foundation for Statistical Computing. (2013). *R: A language and environment for statistical computing* [Computer software]. Retrieved from <http://www.R-project.org/>.
- Richards, W., & Polit, A. (1974). Texture matching. *Kybernetik*, 16, 155–162.
- Robilotto, R., Khang, B.-G., & Zaidi, Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *Journal of Vision*, 2(5):3, 388–403, <http://www.journalofvision.org/content/2/5/3>, doi:10.1167/2.5.3. [PubMed] [Article]

- Šára, R. (1995). Isophotes: The key to tractable local shading analysis. *Computer Analysis of Images and Patterns LNCS*, 970, 416–423.
- Schiller, P. H., Finlay, B. L., & Volman, S. F. (1976). Quantitative studies of single-cell properties in monkey striate cortex. III. Spatial frequency. *Journal of Neurophysiology*, 39, 1334–1351.
- Selle, A., Su, J., Irving, G., & Fedkiw, R. (2009). Robust high-resolution cloth using parallelism, history-based collisions, and accurate friction. *IEEE Transactions on Visualization and Computer Graphics*, 15, 339–350.
- Sharan, L., Li, Y., Motoyoshi, I., Nishida, S., & Adelson, E. H. (2008). Image statistics for surface reflectance perception. *Journal of the Optical Society of America A*, 25, 846–865.
- Tuceryan, M., & Jain, A. K. (1998). Texture analysis. In C. H. Chen, L. F. Pau, & P. S. P. Wang (Eds.), *The handbook of pattern recognition and computer vision* (pp. 207–248). Singapore: World Scientific.
- Vangorp, P., Laurijssen, J., & Dutré, P. (2007). The influence of shape on the perception of material reflectance. *ACM Transactions on Graphics*, 26, 1–9.
- Watson, A. B., & Ahumada, A. J. (2005). Standard model for foveal detection of spatial contrast. *Journal of Vision*, 5(9):6, 717–740, <http://www.journalofvision.org/content/5/9/6>, doi:10.1167/5.9.6. [PubMed] [Article]
- Wijntjes, M. W., & Pont, S. C. (2010). Illusory gloss on Lambertian surfaces. *Journal of Vision*, 10(9):13, 1–12, <http://www.journalofvision.org/content/10/9/13>, doi:10.1167/10.9.13. [PubMed] [Article]
- Zaidi, Q. (2011). Visual inferences of material changes: Color as clue and distraction. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2, 686–700.