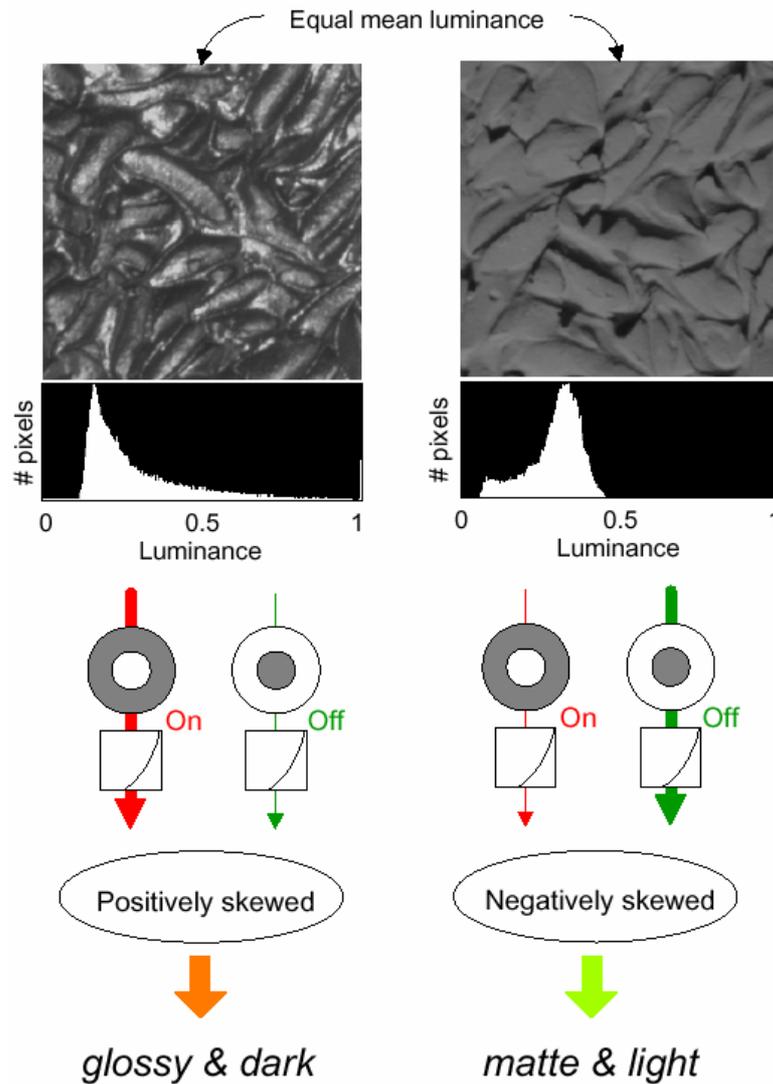


SUPPLEMENTARY INFORMATION



Supplementary Figure 1: A schematic drawing of the main finding of this study. Here we show that the perceived lightness and glossiness of natural surfaces can be predicted, and controlled, by the skewness of the luminance or subband histograms of the surface image. Skewness, or similar statistics, are easily computed by simple neural circuits that compare the pooled outputs of on-center and off-center visual sensors in early visual processing. This notion is supported by the observation that adaptation to visual stimuli with skewed histograms alters the perceived lightness and glossiness of surfaces that are subsequently viewed. We suggest that the human visual system utilizes simple image statistics such as skewness as a part of estimating the material qualities of natural surfaces.

SUPPLEMENTARY METHODS

We handmade 24 samples of stucco-like surfaces (6 gray levels x 4 glossiness levels). The gray level was manipulated by the amount of black pigment in the marble-based white paste, and the glossiness was manipulated by the amount of clear acrylic media coating. The diffuse and specular reflectance of each stucco sample (Fig. 1b in the main text) was estimated from the image taken by a 16-bit linear camera (BITRAN BS-42N) with a polarizer whose direction was parallel or orthogonal to the direction of another polarizer attached to the illumination source. By using the orthogonal images ($I_o(x,y)$, which do not include specular reflection), the diffuse intensity (D) was estimated as the mean intensity of the images ($D = \text{mean}[I_o(x,y)]$). The diffuse reflection was defined as the ratio of the diffuse intensity of each stucco sample to that of the reference (BaSO₄; ~98% albedo). The specular intensity S was estimated as the 99 percentile intensity of the difference between the parallel image $I_p(x,y)$ and the orthogonal image $I_o(x,y)$ normalized by the diffuse intensity D ($S = 99\%(\max[I_p(x,y)-I_o(x,y),0])/D$).

The gray-scale stucco images (256 x 256 pixels) used for the experiments were taken by the same 16-bit linear camera without the polarizer, three times for each stucco under slightly different natural illumination conditions. In the experiments, the image was presented on a CRT monitor (SONY GDM-F500R, refresh rate 100Hz, luminance range of 0.1-82 cd/m²) driven by a graphics card (Cambridge Research System, VSG2/3), with 8-bit luminance resolution for the luminance range of each image.

The standard deviation (SD) and skewness of the luminance histogram of an image were defined as follows.

$$SD = \sqrt{\frac{\sum (I(x,y) - m)^2}{N}}, \quad skew = \frac{\sum (I(x,y) - m)^3}{N \cdot SD^3} \quad (1),$$

where $I(x,y)$ is the luminance of a pixel, m the mean luminance, and N the number of pixels (256 x 256).

Lightness and glossiness ratings: In the first experiment (Fig. 2a-c), one of 63 images (24 stuccos x 3 illuminations, excluding 9 images whose dynamic range significantly exceeded that of the display monitor) was presented at the center of the monitor with the size of 7.3 x 7.3 deg. The mean luminance of all images was normalized to 16.3 cd/m², and the background luminance was kept low (< 0.1 cd/m²). Six subjects (two of the authors and 4 naïve subjects) rated the lightness or the glossiness of the surface using a five-level physical scale (0-4), at least four times for each image. The stimulus was presented until the subject made a response. The measurements of lightness and glossiness were made in separate sessions. Physical samples were shown to the subject as reference — five OSA patches (-6, -4, -2, 0, 2, 4/0/0) for the lightness measurement. For glossiness measurements the most matte (0) and the most glossy (4) stucco patches served as endpoints of the scale. Fig. 1b and 1c show image statistics and rating data for 23 images under one of three illumination conditions. Fig. 1d shows the rating data for all 63 images.

In the second experiment (Fig. 2d), the skewness of the luminance histogram of images of stucco (medium gray, slightly glossy), fabric, and crumpled paper, was varied by the procedure of histogram matching to a Beta distribution, given by:

$$f(l) = \frac{1}{B(p,q)} l^{p-1} (1-l)^{q-1}, \quad B(p,q) = \int_0^1 l^{p-1} (1-l)^{q-1} dl, \quad q = 10 - p, \quad (2).$$

where l is the luminance, p is the parameter ranging from 1 to 9 (which changes the amount of skew approximately from -1.5 to 1.5). The mean (4.1 - 65 cd/m²) and SD (SD/mean = 0.05 - 0.8) were set by rescaling the Beta distribution. We chose the Beta distribution mainly because of its mathematical simplicity. We confirmed that other methods of skewness control (adjustment of intensity gamma, independent modulation

of the increment/decrement contrast) gave rise to similar effects. Seven subjects (two of the authors and 5 naïve subjects) rated the lightness and glossiness of the stimuli.

Aftereffects: The aftereffects in the perceived lightness and glossiness were tested by a cancellation method. The adapting stimulus was a texture consisting of a number of mean-balanced difference-of-Gaussian (DoG) patches given as follows:

$$DoG(x, y) = \exp\left(\frac{\sqrt{x^2 + y^2}}{2\sigma_c^2}\right) - \frac{\sigma_s^2}{\sigma_c^2} \cdot \exp\left(\frac{\sqrt{x^2 + y^2}}{2\sigma_s^2}\right) \quad (3).$$

where σ_c and σ_s are standard deviation of the centre (0.09 deg) and the surround (0.23 deg). The RMS contrast was normalized at 0.4. The resulting images had histograms with skewness of 2.0 or -2.0, and the mean luminance was 16.3 cd/m². The DoG patches were randomly located with a minimum centre-to-centre separation of 0.46 deg. In the experiment, two textures with opposite polarities were presented in the left and right sides (or right and left) with separation of 0.5 deg (Fig. 4a). The subjects viewed them for 100 sec while fixating at a small white dot in the centre of the display. To minimize the local light adaptation, the arrangement of DoG patches was randomly updated every 0.5 sec during the adaptation period. In each trial, after a 4.0 sec re-adaptation (top-up) period and a 0.8 sec blank of the uniform gray, a pair of the test stucco images (Fig. 4b) whose histograms were skewed in the opposite polarities with a particular amount (e.g., 0.5 in the left and -0.5 in the right), randomly chosen from a range of -1.5 to 1.5 (Fig. 4d; x-axis), was presented for 0.5 sec; the subjects then indicated by a button press which of the two surfaces appeared darker (or glossier). After the subject's response, the next trial started immediately.

The aftereffects were also examined by using stucco images as adapters. The stucco images had histogram skew between 1.5 and -1.5 (Fig. 4c). Each adapting image shifted its relative position every 0.5 sec within the image frame with wrap around. For both types of adapting stimuli, the perceived lightness and glossiness were tested in separate

sessions for each of six subjects, and at least 50 data points were collected for each point.

The inter-ocular transfer of the aftereffects was tested using a haploscopic display in which a pair of surface images presented to the either left or right eye and a pair of uniform gray to the other eye. The inter-ocular transfer was defined a ratio of the differential PSE between two opposite adaptations, which was calculated by a logit analysis, when the adapting and test images were presented to the same eye to that when they were presented to the different eyes.

SUPPLEMENTARY MOVIE LEGENDS

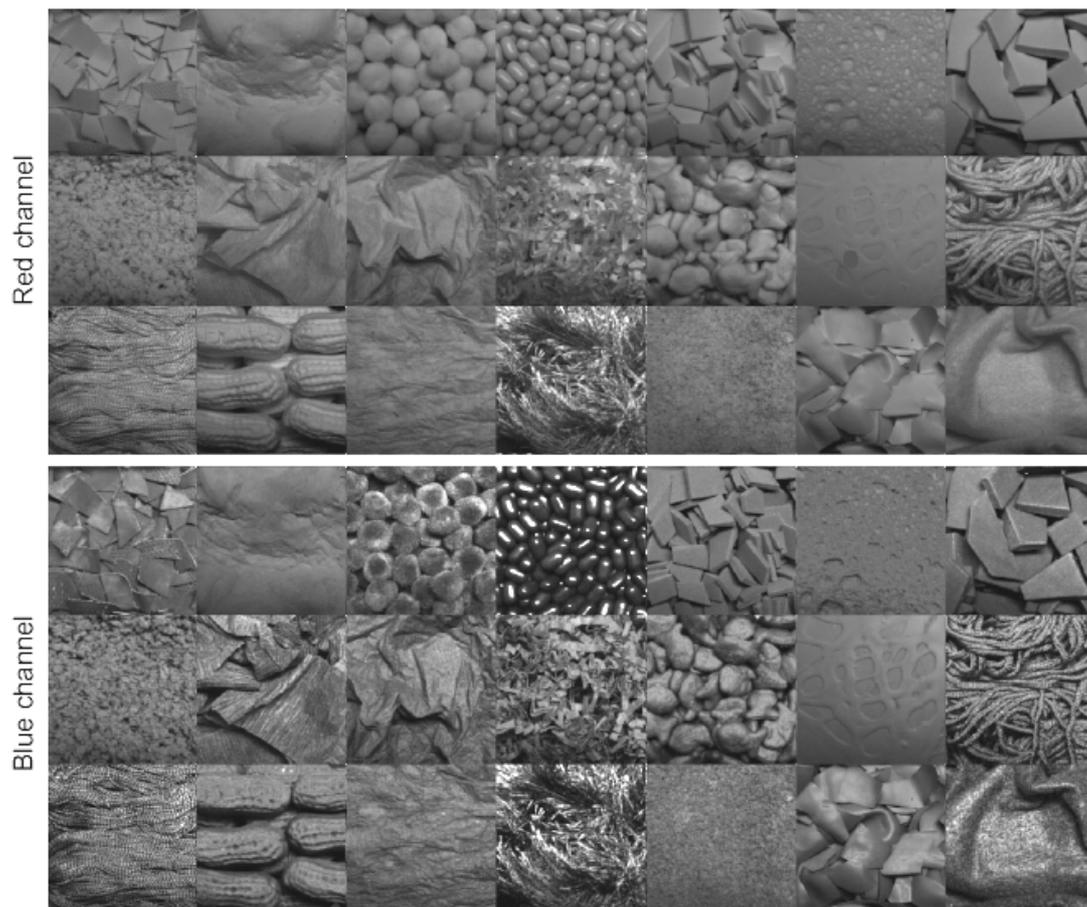
Aftereffects of lightness and glossiness. There are two pairs of movies. These movies demonstrate the aftereffects that follow adaptation to artificial DoG textures (“texture1.mov” and “texture2.mov”) and histogram-manipulated stucco images (“stucco1.mov” and “stucco2.mov”). Within each pair, the positions of adaptation textures are flipped (left-right or right-left). The movies are saved in QuickTime MOV format. The movies are playable at least with Quick Time 7 on Mac OS X and on Windows XP sp2.

When you start a movie, you will see a pair of adapting images for 16 sec. The image positions are updated every 0.5 sec. Please gaze at the small white fixation dot between the two images. When the adaptation phase completes, a uniform blank field follows. A pair of identical stucco images is then presented for 0.5 sec — we found that a brief presentation allows stronger aftereffects. Subjects report that the stucco on the right looks darker and glossier than the one on the left when the left adaptation image has a positively skewed luminance histogram (“texture1.mov” and “stucco1.mov”). For “texture2.mov” and “stucco2.mov” this report is reversed. Repeating the same movie might make the aftereffects stronger. Note however that successive observations of movies with opposite effects will reduce the aftereffects significantly because of cancellation.

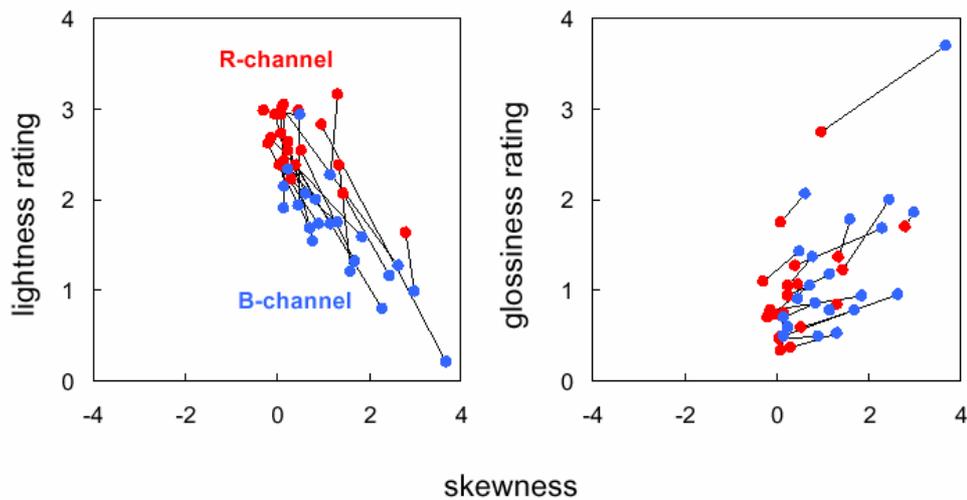
Note also that we generated movies with the assumption that they will be played on a linear monitor. To see the original effects, please set the gamma of the monitor so as to make the display as linear as possible. Even with non-calibrated displays, one can see the aftereffects we report, but artifacts of mean-luminance adaptation might creep in.

SUPPLEMENTARY DATA

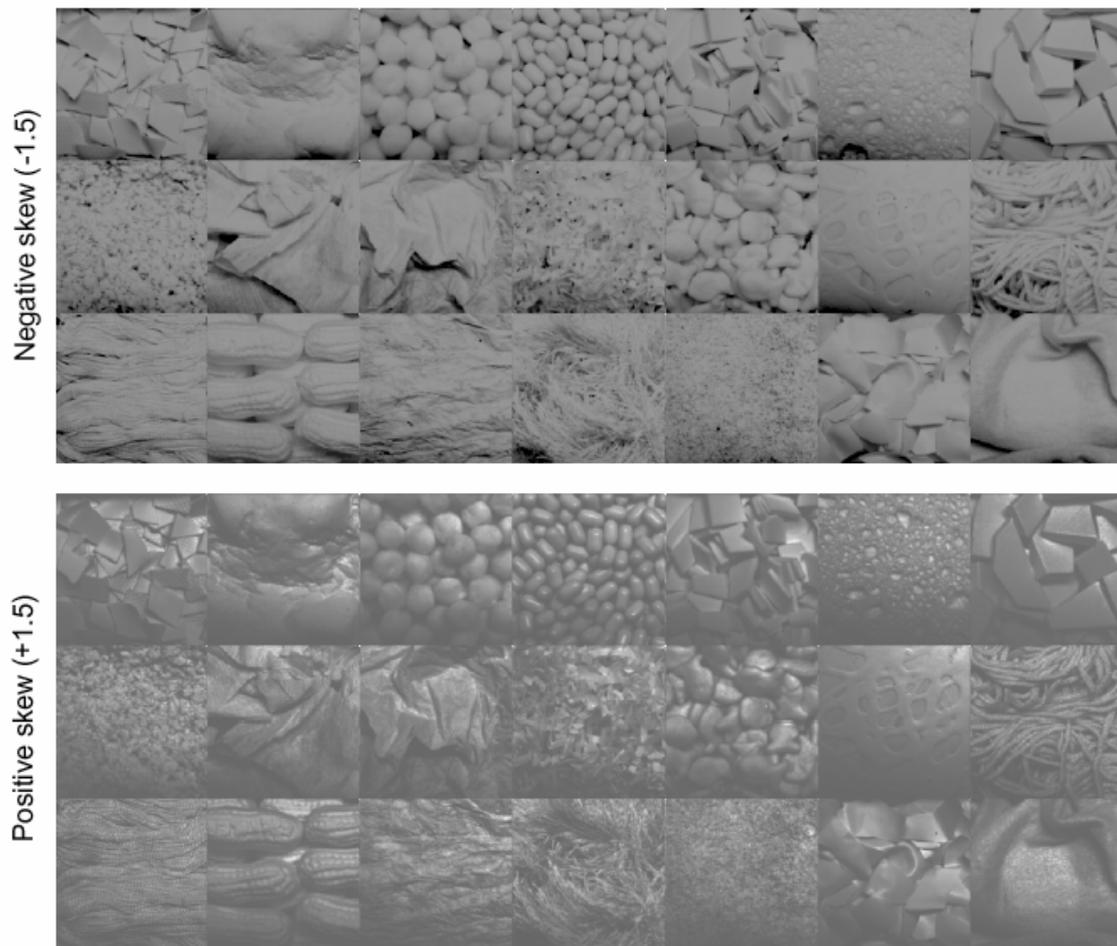
A. The perceived lightness and glossiness for 42 natural surface images. We measured the perceived lightness and glossiness for 42 different images of natural surfaces, and found similar effects of skewness on lightness and glossiness perception, as reported in the main text.



Supplementary Figure 2. Gray-scale images were taken from red and blue channels of linear digital photographs of 21 orange materials including food, plastic plates, strings, etc. By using the two color channels, we could generate a pair of images of the same object with different shades of gray. The mean luminance was normalized at 8.2 cd/m^2 , and the observers were asked to rate the lightness and glossiness. Eleven subjects participated in the experiment.

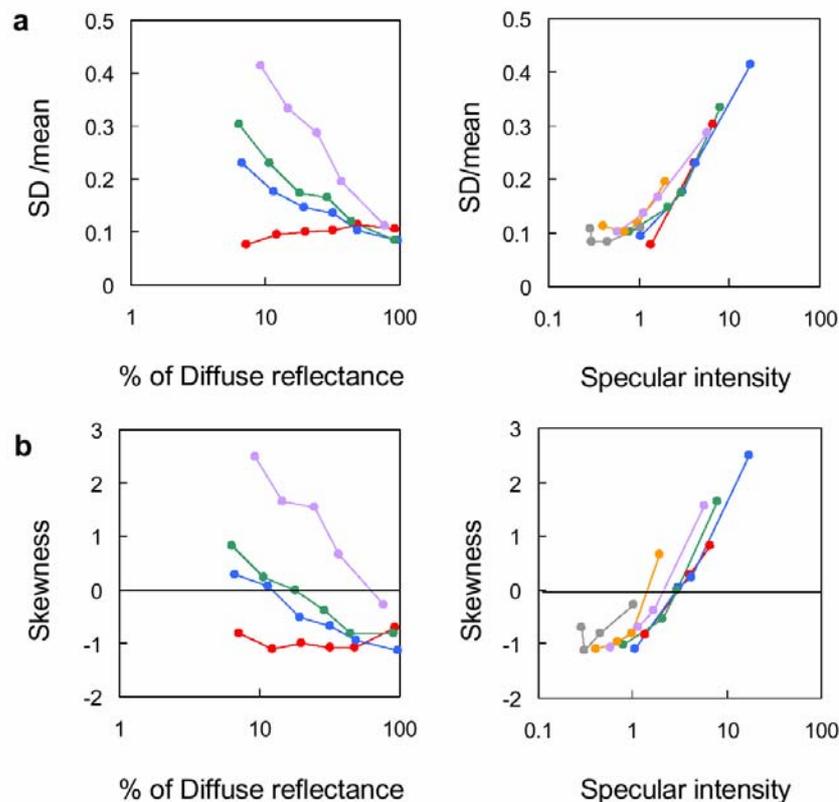


Supplementary Figure 3. The lightness and glossiness ratings for 42 surface images plotted against the skewness of the luminance histogram. The data points for the same object (Red: R-channel, Blue: B-channel) are connected by a solid line. Skewness was negatively correlated with lightness judgments for the same object (indicated by the negative slope of the connecting line), as well as across objects ($r = -0.80$). Skewness was positively correlated with glossiness judgments for each object as well as across all objects ($r = 0.61$). The lower correlation with glossiness ratings, relative to that with the lightness ratings, may be ascribed to the higher sensitivity of gloss perception to the spatial structure of the stimulus.



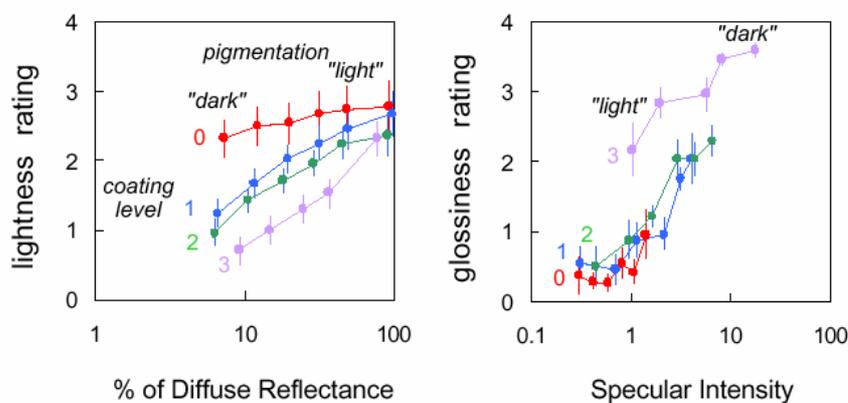
Supplementary Figure 4. The red channel of each surface image was skewed positively or negatively by the process of histogram matching. The two panels clearly demonstrate the effect of skewness on perceived lightness (positively skewed looks darker) and glossiness (positively skewed looks glossier) for all objects, although the magnitude of the perceptual effect may vary across objects especially for glossiness.

B. Comparison of the effects of image statistics on lightness and glossiness perception of stucco-like surfaces. Here is supplementary data for the experiments shown in Fig. 2. In addition to the effects of skewness, we found a minor effect of the standard deviation (SD) of the luminance histogram on both lightness and glossiness judgments. The mean luminance had a significant effect on lightness, but not on glossiness. We found little, if any, effect of kurtosis.

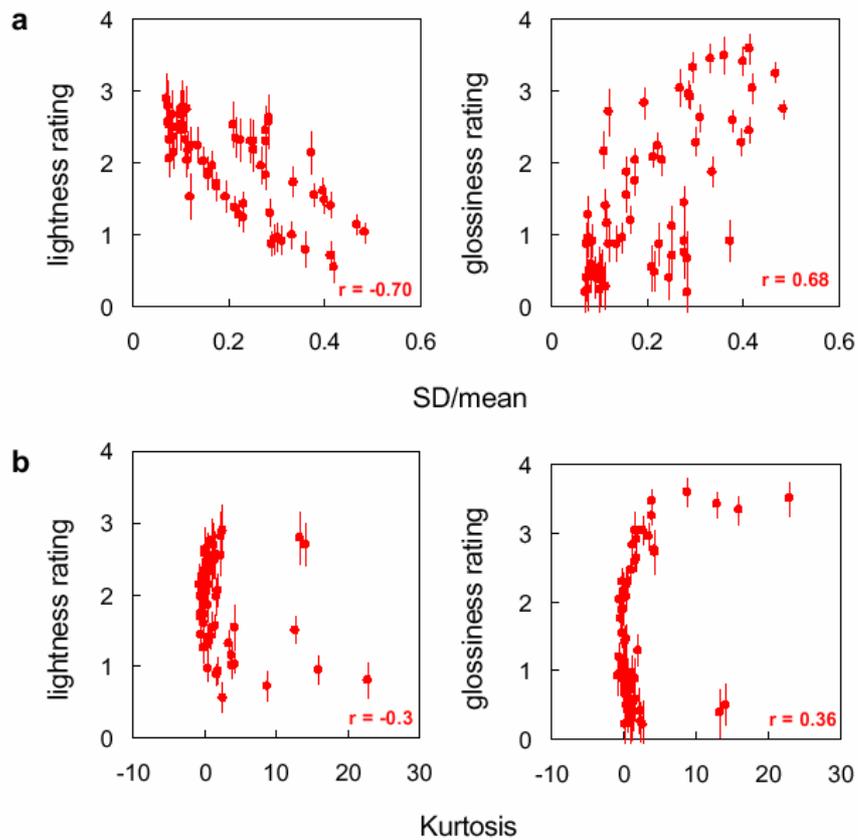


Supplementary Figure 5. Physical image statistics of all stucco-like surfaces (6 gray levels x 4 glossiness levels). The standard deviation (a) and the skewness (b) of the luminance histogram of mean luminance normalized images are plotted as a function of the magnitude of diffuse reflectance (left) and the magnitude of specular reflectance (right). Different colors in the left panel represent the thickness of acrylic coating that controls the magnitude of glossiness (red: none, blue: thin, green: middle, purple: thick). Different colors in the right panel represent the pigmentation that controls the gray level

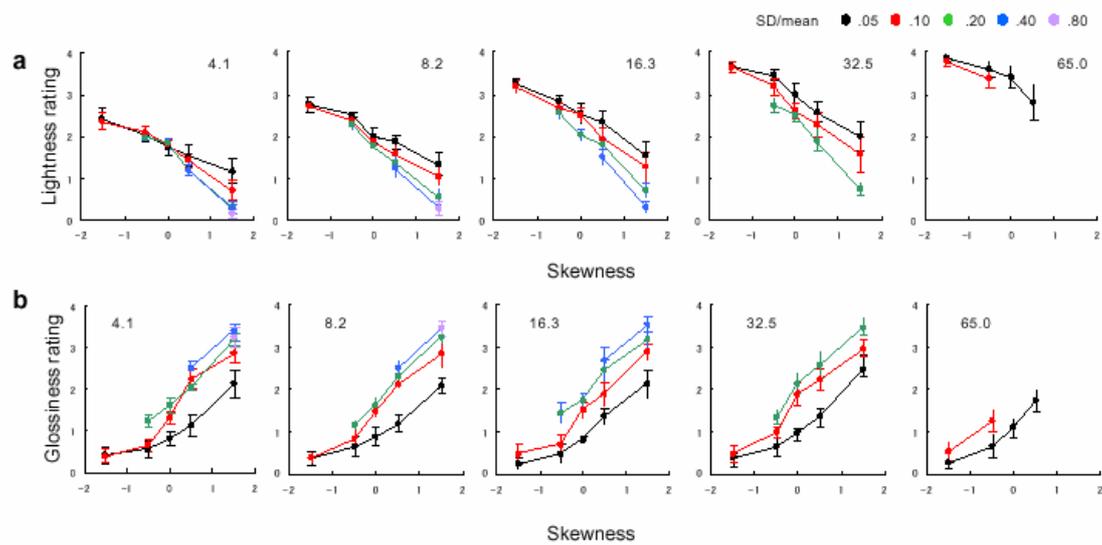
(red: black, blue: dark gray, green: middle gray, purple: light gray, orange: very light gray, gray: white). This figure illustrates that both standard deviation and skewness are negatively correlated with the magnitude of diffuse reflectance, and positively correlated with the magnitude of specular reflectance. The exceptions are matte surfaces (indicated by red symbols in the left panel and by data at the lowest specular intensity in the right panel), whose statistics do not seem to be affected by a change in physical reflectance parameters.



Supplementary Figure 6. The lightness and glossiness judgments for mean-normalized images of all stucco-like surfaces. The lightness rating is plotted as a function of the magnitude of diffuse reflectance (left), while the glossiness rating is plotted as a function of the magnitude of specular reflectance (right). Different colors represent the degree of acrylic coating. Selected data points are shown in Fig. 2b in the main text. Both lightness and gloss perception are well correlated with the corresponding physical properties for all surfaces except the matte ones.



Supplementary Figure 7. The lightness and glossiness ratings for all stucco-like surfaces (a) plotted against the mean-normalized SD (the 2nd-order moment), and (b) plotted against the kurtosis (the 4th-order moment; $\sum (L(x,y) - mean)^4 / (N \cdot SD^4) - 3$) of the luminance histogram. The perceived lightness and glossiness are correlated with the mean-normalized SD, but these correlations are weaker than that with skewness (Fig. 2c). The ratings have little correlation with kurtosis and higher-order moments.



Supplementary Figure 8. The lightness (a) and glossiness (b) ratings for a stucco image whose luminance histogram was matched to Beta distributions of varying values of skewness and standard deviation. Although Fig. 2d in the main text shows the effect of the skewness only, the full data shown here include the effects of the mean luminance and the mean-normalized SD. The data points are plotted against the skewness. Symbols in different color represent the results for different SDs. Each panel shows the results for a given mean luminance. For any combinations of the mean and the mean-normalized SD, the skewness has clear and consistent effect on the perceived lightness and glossiness. The standard deviation of the mean-normalized image has a minor effect on the both lightness and glossiness. The surface appears dark and glossy as the mean-normalized SD increases. However, this effect is evident only when the skewness is positive; no surface was seen dark and glossy when the skewness is negative. The results also show that as the mean luminance is increased, the perceived lightness increases, while the perceived glossiness remains the same.

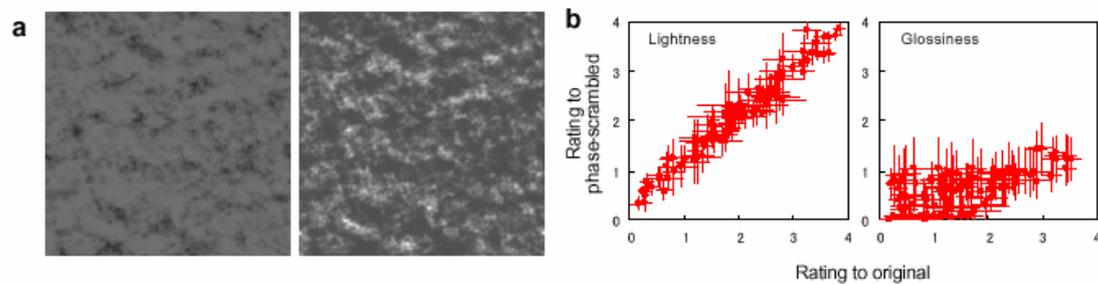
C. Randomizing spatial structures. While skewness is predictive of perceived surface qualities, it can be computed on arbitrary images, whether or not they look like surfaces. Our findings were made in the case where the image is perceived as a surface of uniform albedo with some highlights. To see how stimulus spatial structure interacts with the luminance histogram in lightness and glossiness perception, we altered the spatial structure of a stucco image by (a) randomizing the phase component, or by (b) randomizing the pixel positions.

(a) Phase-randomized stimuli: The spatial phase of the stucco image was randomized while the magnitude of the frequency spectrum was kept intact (Supplementary Fig. 9a). The luminance histogram was matched to skewed Beta distributions (skewness ranged from -1.5 to 1.5). The mean luminance of phase-randomized images was varied from 4.1 to 65.0 cd/m^2 , and the standard deviation from 0.05 to 0.8. Even after these image manipulations, the images still looked like real surfaces, and the observers could judge the both lightness and glossiness easily. Six observers participated in the experiment.

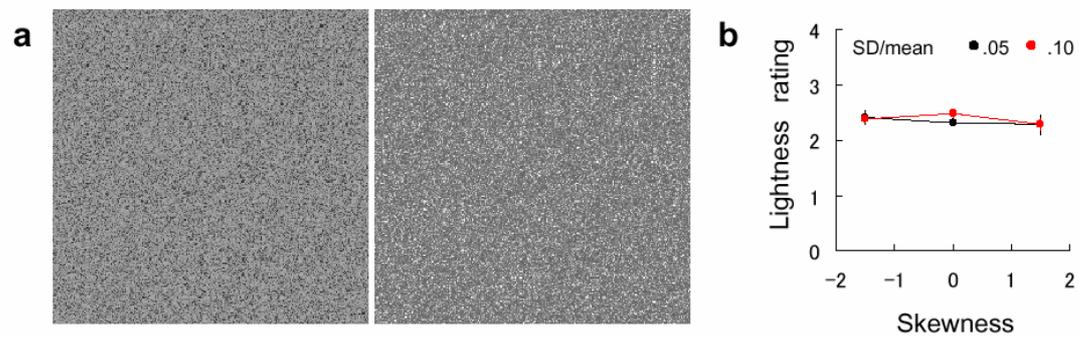
Supplementary Fig. 9b shows the relationships between the ratings of lightness (left) and of glossiness (right) for the original images (x-axis) and their phase-randomized versions (y-axis). While the lightness ratings for the phase-randomized images are almost identical to those for the original images, the glossiness ratings were very low for the phase-randomized images; observers reported that bright spots in the phase-randomized images looked like dust or snowflakes, not specular highlights.

(b) Pixel-randomized stimuli: The position of pixels was randomized (Supplementary Fig. 10a), and the luminance histogram was matched to skewed Beta distributions (skewness varied from -1.5 to 1.5), with the mean luminance set to 8.2 cd/m^2 . Since the pixel-randomized images did not look like real surfaces, the observers could not judge gloss at all. They could only make lightness judgments. Six observers participated in the experiment.

Supplementary Fig. 10b shows the lightness ratings plotted against the histogram skewness. The perceived lightness does not vary over a range of skewness. The results are consistent with a previous study that measured the effects of luminance histogram on overall luminance judgments of IID textures (Nam, J.H. & Chubb, C., 2000, *Vision Res*, 40, 1695-1709). Although we forced observers to make lightness judgments, since the images did not look like real surfaces, subjects may have judged the overall brightness of the image instead.

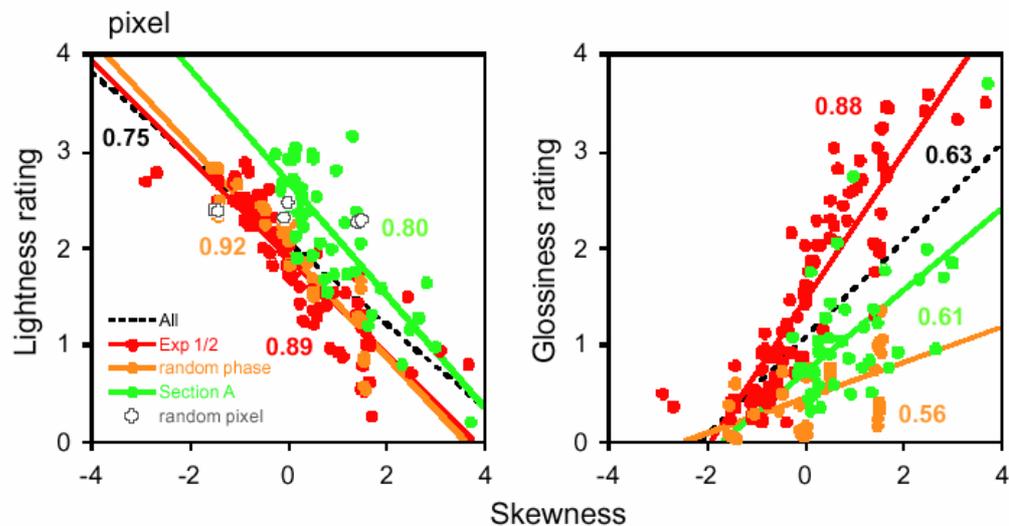


Supplementary Figure 9. Experiments with phase randomized stimuli (a) Examples of stimuli used in the experiment. The spatial phase of a stucco image is randomized while the frequency spectrum kept intact, and the luminance histogram is skewed negatively (left) and positively (right), respectively. (b) The ratings of lightness (left) and glossiness (right) for phase randomized stimuli are plotted against those for the original stucco images. The data was collected for a range of mean, standard deviation, and skewness values. The lightness ratings are not affected by phase randomization whereas the glossiness ratings decrease considerably when phase is randomized; e.g., most subjects judged the surface in the right panel of (a) as very dark, but not glossy.

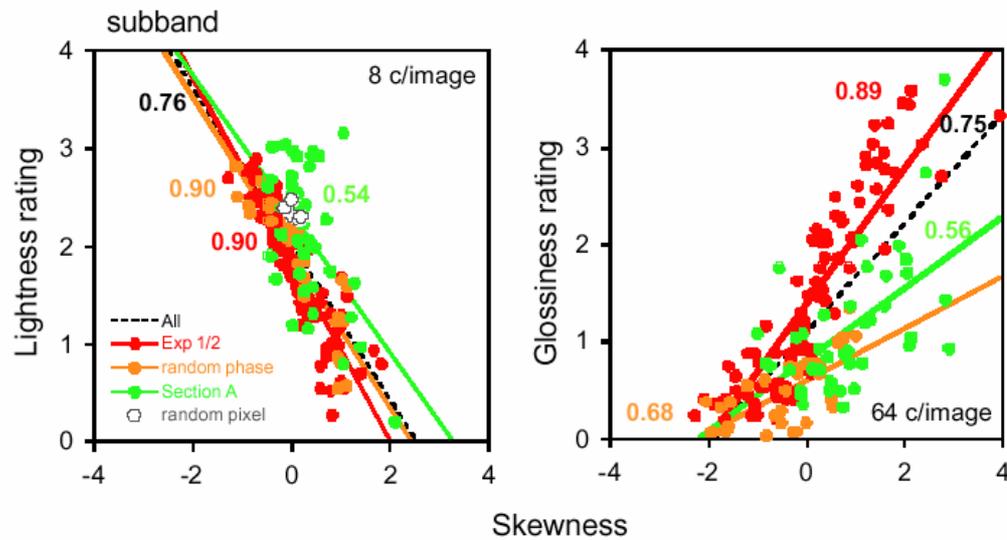


Supplementary Figure 10. Experiments with pixel randomized stimuli (a) Examples of stimuli used in the experiment. The luminance histogram is skewed negatively (left) and positively (right), respectively. (b) The lightness rating as a function of skewness in the luminance histogram. The ratings did not vary much over the range of skewness values we tested.

D. Summary of the data and comparison of the effects of luminance skewness and subband skewness. All the ratings data reported in this paper are plotted together against luminance skewness (Supplementary Fig. 11), and subband skewness (Supplementary Fig. 12). The lightness ratings show high correlation with luminance skewness for all image sets, and are well approximated by linear functions with similar slopes. The outliers (pixel-randomized images) can be explained by subband skewness. On the other hand, glossiness ratings are not quantitatively consistent across categories. They increase with luminance skewness and are always low for images with negative luminance skewness, but the rates of increase are different for different categories.. The perceived glossiness increases rapidly and consistently for the original and histogram controlled stuccos, but slowly and diversely for other materials and for phase-randomized stuccos. Subband skewness does not seem to account for glossiness data any better than luminance skewness. We believe that in addition to skewness computations, glossiness perception may involve an analysis of image structures beyond subband filtering to distinguish the effects of specular highlights from other factors.



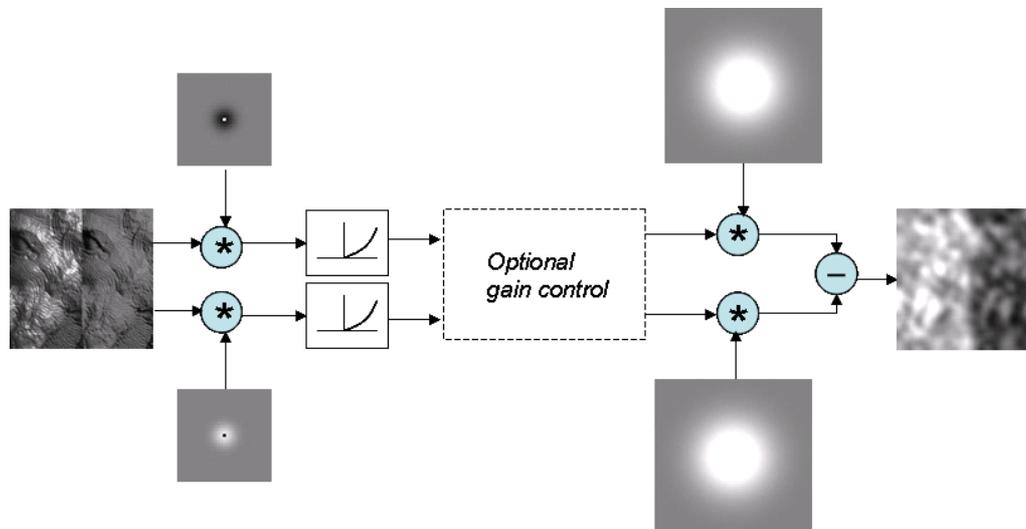
Supplementary Figure 11. The lightness (left) and glossiness (right) ratings are plotted against skewness of the pixel luminance histogram. The red circles show the results of the experiments that used stucco images (original and histogram controlled, 93 images), green circles are the results for images of orange objects (Supplementary Fig. 3, 42 images), orange circles the results for phase-randomized stucco images (Supplementary Fig. 9, 32 images), and open circles for pixel-randomized images (Supplementary Fig. 10, 6 images). The mean luminance of each image was 8.2 cd/m^2 , except for the original stucco images (16.3 cd/m^2). The lines depict the fitted linear functions (no lines were fitted to pixel-randomized data), and the colored numbers depict correlation coefficients between the ratings and the luminance skewness for each category. The number in black is the correlation coefficient for all categories of data (173 images for lightness, 167 images for glossiness).



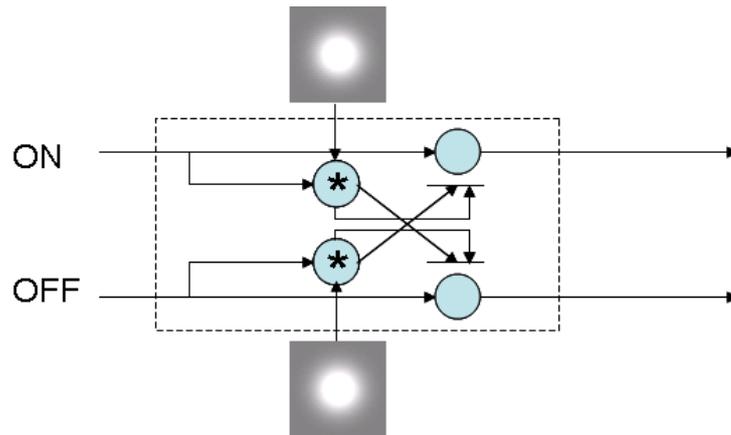
Supplementary Figure 12. The lightness (left) and glossiness (right) ratings plotted against skewness of the subband histogram. The subband images were obtained with a Gaussian band-pass filter of the bandwidth of 2.0 octaves. The subbands are 8 c/image for lightness and 64 c/image for glossiness. Although these subbands provide the best account of our results, the differences from other subbands were small.

SUPPLEMENTARY DISCUSSION

A proposed skewness detection mechanism. Here we elaborate on the qualitative model outlined in the main text (Supplementary Fig. 13). We know that skewness, as defined by Karl Pearson, is the expected value of X^3 if X is a zero mean, unit variance random variable. Based on that intuition, we developed our model as follows. The input image is convolved with a set of on-center and off-center DoG (difference of Gaussian) filters. (A similar computation can be implemented with oriented filters.) If the ratio of the size of the Gaussian in the center to the size of that in the surround is small i.e. the filter is more punctate, the filtering operation is akin to removing the local mean from the signal. From a physiological point of view, signals are non-negative; therefore we add the half-squaring (half-wave rectification followed by a squaring non-linearity) operation and use the on-off filter pair to preserve the sign information. We chose the squaring non-linearity but in principle any accelerating non-linearity suffices for computing histogram asymmetry. At present, it is impractical to exactly specify the shape of the non-linearity, since it depends on a number of unknown parameters, such as coefficients of spatial filters and the index in the skewness formula. Next, an optional gain control stage that allows divisive normalization (Heeger, 1993, *J Neurophysiol*, 70, 1885-1898) follows where the on and off stream outputs are divided by the spatially pooled signal from both streams (see Supplementary Figure 14). Local energy is computed by squaring and blurring the output of the either the on or the off filters. Alternatively, the on and off filter outputs can undergo half-squaring followed by blurring and addition to yield the local energy. In the gain control stage in Figure 14, the squared on and off-streams are normalized by the local energy. At the end of this stage, the signal in both streams is the squared version of a zero (local) mean, unit (local) variance signal. The signal is then spatially pooled and subtracted. The final signal is an estimate of local skewness, similar to Pearson's definition with a squaring in the numerator instead of cubing.



Supplementary Figure 13. Proposed model for skewness detection. The input image is convolved with on-center and off-center linear filters followed by half-squaring operation. Next, the images in both on and off-channels undergo an optional gain control stage (Heeger 1993) where a signal that equals local energy is divided out. Local energy is computed by squaring and blurring the image right after filtering and before half-wave rectification in either channel or alternatively by inserting the gain control stage shown in Supplementary Fig. 14. The result in both channels is then spatially pooled and the streams are subtracted. The difference signal is an estimate of local asymmetry, similar in spirit to Pearson's skewness.



Supplementary Figure 14. The gain control stage shown here allows a division by local energy by pooling over spatial neighbors both within the on and off streams as well as across streams. This stage could follow the half-squaring operation in Supplementary Fig. 13.