

Image statistics do not explain the perception of gloss and lightness

Barton L. Anderson

School of Psychology, The University of Sydney,
NSW, Australia



Juno Kim

School of Psychology, The University of Sydney,
NSW, Australia



A fundamental problem in image analysis is to understand the nature of the computations and mechanisms that provide information about the material properties of surfaces. Information about a surface's 3D shape, optics, illumination field, and atmospheric conditions are conflated in the image, which must somehow be disentangled to derive the properties of surfaces. It was recently suggested that the visual system exploits some simple image statistics—histogram or sub-band skew—to infer the lightness and gloss of surfaces (I. Motoyoshi, S. Nishida, L. Sharan, & E. H. Adelson, 2007). Here, we show that the correlations Motoyoshi et al. (2007) observed between skew, lightness, and gloss only arose because of the limited space of surface geometries, reflectance properties, and illumination fields they evaluated. We argue that the lightness effects they reported were a statistical artifact of equating the means of images with skewed histograms, and that the perception of gloss requires an analysis of the consistency between the estimate of a surface's 3D shape and the positions and orientations of highlights on a surface. We argue that the derivation of surface and material properties requires a photo-geometric analysis, and that purely photometric statistics such as skew fail to capture any diagnostic information about surfaces because they are devoid of the structural information needed to distinguish different types of surface attributes.

Keywords: surface properties, material perception, image statistics, decomposition, histogram skew

Citation: 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://journalofvision.org/9/11/10/>, doi:10.1167/9.11.10.

Introduction

The pattern of light that reaches our eyes arises from the interaction of light with surfaces. The structure in images arises from a combination of surface optics, the illumination environment, surface geometry, and intervening media. Although these distinct physical sources are conflated in the image, the visual system appears to decompose the image into a representation of its constituent causes. We experience the world as a distribution of surfaces and objects that possess specific 3D shapes, material and reflectance properties, embedded in a particular illumination environment. Understanding how the visual system derives such information from images is a central goal of mid-level vision.

One longstanding theoretical approach asserts that the visual system engages in a form of “inverse optics” to decompose images into their underlying causes. This view took modern form in Barrow and Tenenbaum's (1978) intrinsic image analysis, which suggested that the visual system represents each physical cause in a separate “map.” One appeal of this approach is that it acknowledges our ability to extract information about illumination, shape, and surface optics when viewing natural scenes. However, the

mapping from the world onto an image is a destructive process, and hence no unique solution to the inversion problem exists.

The apparent intractability of this problem has inspired the search for computational “short-cuts” that might preclude the need to explicitly compute intrinsic images to extract surface and material properties. To this end, it was recently suggested that the visual system utilizes simple image statistics to help derive information about the lightness and gloss of surfaces (Motoyoshi, Nishida, Sharan, & Adelson, 2007). Motoyoshi et al. (2007) motivate their proposal as follows:

“The image of a surface arises from the combination of the surface geometry, the surrounding illumination, and the surface optics. Each of these components can be complex (for example, the reflectance at each point is characterized by a four-dimensional function known as the bidirectional reflectance distribution function). Each is typically unknown, and estimating any one using “inverse optics” requires knowing the others. To bypass this problem, we have looked for simple statistical image measurements that can provide information that is useful even if not complete (p. 206).”

Motoyoshi et al. (2007) further suggested that the visual system uses a specific image statistic—the skewness of the luminance histogram, or skewness of the sub-band filter outputs—to estimate the lightness and gloss of surfaces. The motivation for this hypothesis emerged from their investigation of uniform albedo stucco surfaces that were physically manipulated to have different albedos and gloss. They observed that as the gloss of a surface is increased, or the albedo of a surface is decreased, the relative contribution of the specular reflectance component to the overall reflectance of an object increases. This increases the relative number of pixels in the upper tail of the luminance histogram, causing histogram skew to become increasingly positive. Motoyoshi et al. found their observers' ratings of perceived gloss and lightness were correlated with skew (positively and negatively, respectively). They further showed that if the histogram of a uniform albedo surface is coerced to have a particular value of skew, observers' ratings of lightness and gloss were affected in much the same manner as if they viewed real surfaces that possessed the specific value of skew examined. Finally, they conducted adaptation experiments and found evidence for opponent aftereffects of stucco surfaces when observers adapt to images with a positive or negative skew.

These results provide intriguing correlational evidence for the role of skew computations in the perception of gloss and lightness and suggest that it may indeed be possible to find simple image statistics that are diagnostic of surface properties. Alternatively, we will argue that the data presented by Motoyoshi et al. (2007) likely reflect spurious correlations between skew, gloss, and lightness that arose from the particular way in which they constructed and evaluated their images. We will argue that histogram and sub-band skew contain all of the same ambiguities that such computations are supposed to help resolve and therefore do not provide any insight into how the visual system derives surface properties from images. We further argue that Motoyoshi et al.'s results, and those presented herein, are consistent with the kind of intrinsic image analysis that they purport to bypass with image statistics. In particular, we will present evidence that the perception of surface gloss depends critically on the consistency in the locations and orientations of specular highlights relative to the 3D surface geometry, information that cannot be deduced from the skew computations that they proposed to account for these percepts.

Skew, lightness, and gloss

In order to evaluate the conjecture that histogram or sub-band skew plays a role in the computation of surface properties, we begin by considering the meaning of skew and how it might provide leverage into the computation of surface properties. Skew is a measure of histogram asymmetry; it measures the balance between the positive and

negative tails of a distribution. Motoyoshi et al. (2007) considered two possible distributions from which skew might be computed: pixel histograms and sub-band filter responses. A pixel histogram is a distribution of the “counts” of discrete image intensities. All information about the relative spatial positions of the different intensities—i.e., all geometric information—is discarded. This means that the information contained in a pixel histogram is *purely photometric*. Thus, the claim that pixel histograms provide information diagnostic of surface properties is tantamount to asserting that surface properties can be deduced from purely photometric information, i.e., that the geometric *distribution* of pixel intensities is irrelevant. Although some geometric information is measured by sub-band filter outputs, Motoyoshi et al.'s experiments showed that such measures do not provide information about gloss and lightness much beyond that contained in luminance histograms (see [Skew and adaptation](#) section). Hence, the structural properties measured by sub-band skew do not appear to capture any relevant geometric constraints beyond those that might exist implicitly in histogram skew of the images that they studied. In what follows, we will therefore focus our discussion on the skew of luminance histograms and return to the role of sub-band skew when we consider non-uniform albedo surfaces.

The claim that purely photometric information can provide information about surface properties seems remarkable given the multitude of different physical variables that contribute to the luminance variation in an image. As noted above, image structure arises from a combination of 3D shape, illumination, surface optics, and atmospheric media. There would appear to be two possible reasons that correlations between skew, lightness and gloss might be observed: (1) the skew of luminance and/or sub-band histograms is strongly constrained by lightness and gloss but are only weakly affected by other possible sources of image variation; or (2) the other sources of image variation were artificially restricted in the images that they evaluated such that gloss and lightness were the only variables that had a significant impact on histogram shape. In order to determine which possibility is responsible for the results that they report, we must consider how the images that they studied were constructed.

The primary experiments reported by Motoyoshi et al. (2007) utilized a set of uniform albedo stucco surfaces that they constructed by hand. Surface albedo was manipulated by varying the amount of black pigment in a white paste, and gloss was manipulated by varying the amount of clear acrylic coating applied to the stucco surfaces. The surfaces contained pseudo-random 3D textures (see [Figure 1](#)) that were illuminated by a combination of ambient and directional light sources that were held approximately constant across their experiments. For the experiments involving the hand-constructed stucco surfaces, histogram skew accounted for a large proportion of the variance in observers' judgments of both lightness and gloss ($r^2 = .76$




Figure 1. Schematic illustrating the method used by Motoyoshi et al. (2007) to acquire natural images of hand-made stucco surfaces under directional and ambient illumination (A). The pixel histograms of the images were transformed to have either a positive (B) or negative (C) skew.

and .79, respectively). Importantly, however, when a range of different uniform albedo materials was evaluated, the proportion of variance accounted for dropped significantly ($r^2 = .64$ and $.37$ for lightness and gloss, respectively). What accounts for these correlations, and why did the consideration of more varied surfaces lead to such a sharp reduction in the correlation of skew with gloss?

First, consider the negative correlation that Motoyoshi et al. (2007) report between skew and perceived lightness. They compared lightness judgments of images with identical means but different skews and found that perceived lightness decreased as skew increased. Negatively skewed images (those with sparse dark regions) appeared lighter, and positively skewed images (with sparse light regions) appeared darker. This correlation is their most robust finding over the different materials they studied; the gloss

results were much less consistent (note r^2 values reported above). Although no theoretical justification is offered for the decision to equate the means of the images, it appears that the motivation for this decision lies in the attempt to equate the average image intensity across the images that observers compared, so that lightness judgments could not be based on this first order statistic. The question is whether equating means is a reasonable method of equating “average” image intensity for the images that they studied. When considered from a purely statistical perspective, it is well known that means are a poor estimate of central tendency for skewed distributions. Means are strongly affected by outliers, and their use as a measure of central tendency becomes increasingly inappropriate as the asymmetry of the underlying distribution—the skew—increases. Means are “pulled” in the direction of the outliers and hence provide a poor measure of where

the central mass of a distribution is located. This simple statistical fact alone could fully explain the observed negative correlation between skew and perceived lightness reported by Motoyoshi et al. As the skew of the histogram increases, the majority of the surfaces that are rated as “darker” contain a larger proportion of dark pixels and a few sparse bright pixels; the converse is true for negative skew and light surfaces. In other words, the majority of pixels in images that are rated as darker *are* darker than the images that are rated as lighter. The fact that observers’ judgments may be determined by the mere preponderance of pixels in their images is supported by their finding that the lightness effects they reported were also observed in phase scrambled versions of their stimuli, which did not even appear as uniform albedo surfaces, or even very clear surfaces of any kind. Thus, the correlation between skew and lightness they report is likely to be a statistical artifact of equating an inappropriate measure of central tendency of image luminance for skewed distributions and hence does not provide any evidence that skew computations play a role in perceived lightness.

Second, Motoyoshi et al. (2007) observed a positive correlation between skew and surface gloss for uniform albedo surfaces. In these images, the luminance variation in the images arises primarily from three variables: 3D surface geometry; surface sheen (gloss); and the illumination field. For the stucco images used in their main experiment, the illumination environment and surface geometry were held statistically constant. This implies that the illumination field and surface geometry could not contribute a significant source of variation to the shapes of the luminance histograms across the class of images studied. Virtually all of the variation in the luminance histograms across the images they tested arose from the one variable that was free to vary: surface gloss. Thus, if observers can visually distinguish the different degrees of surface gloss in their stimuli, their gloss judgments would have to correlate highly with histogram skew for the simple reason that the physical gloss of these surfaces is highly correlated with skew in this particular set of images. But this correlation would occur no matter what computations are actually responsible for observers’ perception of gloss; the result would inevitably follow from the fact that skew was strongly correlated with the physical gloss of surfaces judged by observers. The computations underlying the perception of surface gloss may entail sophisticated analyses of surface geometry that evaluate the consistency in the locations and orientations of highlights relative to this geometry and may not explicitly compute skew at all. Indeed, some suggestive evidence that such analyses are in fact essential was presented in Motoyoshi et al.’s supplementary material that accompanied their target article. When a range of different uniform albedo materials that varied more extensively in shape was evaluated, the correlation between skew and perceived gloss dropped off precipitously, accounting for only 37% of the variance in

observers’ gloss judgments (see also Sharan, Li, Motoyoshi, Nishida, & Adelson, 2008).

In what follows, we show that the skew of luminance histograms can be strongly affected by all of the variables that Motoyoshi et al. (2007) did not manipulate or vary, without any concomitant change in perceived gloss. We will argue that there are no data to support the claim that skewness computations play a causal role in the perception of lightness and gloss. Indeed, we will argue that there are currently no data to sustain a claim that the visual system computes either histogram or sub-band skew. Rather, we will argue that the gloss effects they reported arose from interactions between the perceptual analyses of 3D shape from shading and the locations and orientations of potential “highlights”—luminance extrema—in the types of images they studied.

Gloss, skew, and 3D shape

One of Motoyoshi et al.’s (2007) strongest sources of evidence for skew computations was their demonstration that the perception of gloss and lightness could be modulated by directly manipulating the skew of the luminance histogram. The input image could have any initial histogram shape, yet once transformed to have a particular skew, would appear to have the gloss of a real surface that generates the simulated value of skew (i.e., it would appear glossy if it had a positive skew, or matte if it had a negative skew). Note that the process of transforming the shape of the luminance histogram is itself a purely photometric manipulation; it simply re-scales the luminance values of pixels in particular locations of an image, leaving their spatial distribution unchanged. It would therefore appear that purely photometric transformations can affect the appearance of surface properties like gloss, potentially supporting the claim that skew provides diagnostic information about gloss and lightness. There is, however, an alternative explanation of Motoyoshi et al.’s findings that they did not consider. Rather than relying on 2D image statistics, the perception of surface properties may require a complex analysis of the compatibility of the information about a surface’s 3D shape, reflectance properties, and its illumination field (as suggested by intrinsic image models). Indeed, specular highlights of real objects are intrinsically tied to the 3D geometry of surfaces (Koenderink & van Doorn, 1980); they “cling” to regions of high surface curvature and are unstable for low regions of curvature (i.e., move or stretch along regions of low curvature with a change in the position of the light source or viewing position). Beck and Prazdny (1981) showed that the perception of gloss was most compelling when the orientations of the highlights were elongated along the direction of minimal curvature of a smooth 3D object (a vase) and were less effective at inducing the percept of gloss when oriented along the direction of high curvature (cf. Todd, Norman, & Mingolla, 2004). This

finding seems at odds with a claim that the computation of gloss could be derived from purely photometric information. But if gloss perception depends critically on the relationship of highlights to 3D shape, why were Motoyoshi et al.'s histogram transformations effective in altering the perceived gloss of their surfaces?

Although histogram transformations appear to be purely photometric, a more careful analysis reveals that there are implicit geometric constraints embodied in such transformations. Any image of a surface contains a specific geometric distribution of intensities. This geometric structure introduces correlations between the intensities of neighboring pixels. More precisely, the particular type of correlational structure between the intensities in an image is what defines the geometric structure in an image. Transforming the shape of luminance histograms largely preserves the correlations in the intensities of neighboring pixels; only the relative *amplitudes* of the intensities are affected, not their geometric distribution. In other words, manipulating the shape of the luminance histogram alters the strength of local contrasts (luminance gradients) in the image, but it does not alter their spatial distribution. This means that, although manipulating the luminance histogram is itself a purely photometric transformation, its efficacy in shaping perception of surface attributes may have arisen because such photometric transformations work in concert with the geometric structure present in the image to modulate perceived gloss, not because skew per se is computed to infer gloss.

Thus, there are at least two possible explanations for the efficacy of histogram manipulations in modulating the perceived gloss of surfaces. One possibility is that the visual system computes skew to infer gloss. Motoyoshi et al. (2007) state that their “putative skewness cells would be selective for contrast sign, but not for position and not necessarily for orientation” (p. 208). Thus, if the visual system relies on skew computations to infer surface gloss, the perceived gloss of a surface should be relatively unaffected by the positions and orientations of highlights on a surface. On the other hand, if the visual system computes surface gloss by evaluating the *consistency* between the 3D shape of a surface, and the positions and orientations of highlights, then the perception of gloss should be strongly dependent on the locations and orientations of the highlights relative to the geometry of the surface.

In order to distinguish between these two possibilities, we created variants of the St. Matthew images presented in Motoyoshi et al.'s (2007) paper. The gloss and matte variants of these images were created from a common image, which was forced to have either a positive or negative skew. The largest difference in the matte and gloss variants of this image arises from the strength of the luminance maxima in the two images, which occur in the same positions in the two images. This allows us to “peel off” the highlights by taking a difference image of the matte and glossy images. This difference image

containing the specular highlights can be added to the matte image to produce the glossy variant of the St. Matthew image when the two images are aligned. However, by combining the difference image with the matte image, we can parametrically vary the degree of alignment between highlights and the 3D surface geometry as specified by shape from shading computations, which can be used to determine whether (or the extent to which) the visual system evaluates the congruence between the locations and orientations of highlights, and the 3D surface geometry, when computing the gloss of a surface.

Methods

Participants

Twenty naïve first year psychology students participated in this experiment for 1 hour course credit (25% of their target commitment per semester). All participants indicated that they conformed to the strict pre-condition selection criteria of normal color vision and the absence of neurological pathology.

Stimuli

Visual stimuli were constructed using the original matte and glossy versions of the sculpted profile of St. Matthew developed in the Motoyoshi et al. (2007) study. We were able to obtain higher resolution versions of these images courtesy of the first author of that paper, sampled into 512×512 regions which we used here. As exemplified in Figure 2, in order to isolate the specular layer from the glossy surface, we simply subtracted the intensity map of the matte version of the St. Matthew image (B) from the glossy version (A). The residual image (C) formed a “gloss map” containing all the local luminance increment information corresponding to the specular highlights in the original glossy image. Psychophysical stimuli consisted of images produced by a linear combination of the original matte image with different angular or linear offsets of the computed gloss map (Figures 2D and 2E, respectively). Due to the misalignment of image edges after the rotational and translational operations were performed, a suitable cropping region was imposed (red circle and square in Figures 2D and 2E, respectively) to retain the portion of the image produced from superimposition of both the gloss map and the original matte image. Angular offsets were -30 , -15 , -10 , -5 , 0 , $+5$, $+10$, $+15$, and $+30$ degrees, where 0 degrees conformed to the original glossy image. Sample images produced with 0 and -30 degree rotational offsets are shown in Figures 3A and 3B, respectively. Horizontal translations of highlights were in the leftward direction and ranged between 0 and




Figure 2. Schematic of the method used to extract a specular highlight map (C) from 2D glossy (A) and matte (B) images of the same surface geometry (top panel). Spatial manipulations of these extracted highlights by angular rotation (D) and linear translation (E) could be performed before additively re-combining them with the original matte image. The red circle and square show the cropping limits used to discard regions of the two images that could not be superimposed due to edge misalignments.

60 pixels of offset (i.e., 0% to 12% of the image width, approximately). After the test images were constructed, their overall luminance histogram skewness was modified to +0.6 precisely via the application of an appropriate Gamma transformation. Image contrast was finally adjusted multiplicatively to RMS values of approximately 0.15 for rotational stimuli and 0.17 for translational stimuli.

Procedure

A two-alternative forced choice (2AFC) paired comparisons experiment was conducted using custom software written in Visual C++ 6.0, which ran under Windows XP on a MacPRO computer. Stimuli were presented on a 21-in. Sony Trinitron CRT monitor (1280 × 1024 resolution), which was linearized via calibration with the Spyder 2 Pro (DataColor). The calibration was further verified using the Light Scan version 2 (Cambridge Research Systems), which

indicated a black point of approximately 0.4 cd/m² and a white point of 82.4 cd/m².

At any one time during the 2AFC task, only one stimulus image (13.5 cm width/diameter) was visible against a uniform gray background (luminance = 41.0 cd/m²). Subjects viewed images at a distance of approximately 50 cm from the display (producing stimulus images subtending an approximate visual angle of 15 degrees). The participant could toggle between the two images in a 2AFC pairing by depressing the spacebar on a standard 101 PC keyboard at which time the screen was cleared to neutral gray for 1.0 s before the alternative image was displayed. When both images had been viewed at least once, the subject could select the glossier of the two images by making it visible and pressing the up arrow key to indicate their preference, where the screen was again cleared for 2.0 s before advancing to the next trial. For both rotational and translational stimuli, the order of presentation was randomly permuted and counter-balanced with up to 72 images ($n^2 - n$, where the




Figure 3. The result of rotating highlights attributed to specular reflection on the glossy version of St. Matthew (A). After matching histogram skewness, the same highlights (here, rotated by -30°) appear as light surface pigmentation on a completely matte surface (B). Image skew is normalized to $+0.73$ and RMS contrast adjusted to 0.12 in the two images matched in mean luminance.

number of images $n = 9$) being presented in one block of stimulus trials. The experimental session was completed after participants completed four blocks of trials in each condition or when they exceeded the duration of their time slot.

Data analysis

Paired comparisons data from each participant were pre-processed by taking as a proportion the number of times a stimulus was selected as glossier over the total number of times it appeared in each 2AFC experimental session. This proportion directly reflects the probability that a given image will be chosen as glossier for any two images randomly selected from our image pool. Given that gloss was expected to be contingent on the relationship between the closeness in orientation of the specular highlights and surface geometry specified by the diffuse shading profile, we also simulated the observers' predicted responses by correlating the peak responses of locally oriented filters to the matte (diffuse) and specular image layers used to construct the test images. This analysis is motivated by the suggestion by a number of authors that the perception of shape from shading is based on the pattern of isophotes (Koenderink & van Doorn, 1980) or "shading flow" (Ben-Shahar & Zucker, 2001; Breton & Zucker, 1996; Fleming, Dror, & Adelson, 2003; Fleming, Torralba, & Adelson, 2004; Huggins, Chen, Belhumeur, & Zucker, 2001). More specifically, we constructed "orientation fields" of the matte image and gloss map by convolving each location in the image with a Gabor filter (four pixels wavelength, zero phase offset, 1:1 aspect ratio, and one octave bandwidth) for angles ranging between 0 and 170 degrees (at steps of 10 degrees). The orientation field value at a pixel was taken as the angular filter value that yielded the greatest response over this range. Point-wise correlation between the diffuse and gloss orientation

fields was performed for all points within the gloss map that were greater than zero, yielding an index of consistency in the orientation field responses between the gloss and matte images. These correlations were multiplicatively scaled between 0 and 1, giving rise to ideal performance in our paired-comparisons experiment. Due to inevitable response noise, these ideal responses were re-scaled by estimating observers' false positives and false negatives from the data set we obtained, yielding the prediction curves in Figures 4 and 5. Range scaling of the prediction curves was performed according to:

$$P_i = O_i \cdot (1 - E_0 - E_1) + E_1, \quad (1)$$

where P_i is the probability of selecting image i with an orientation field correlation of O_i and error parameter E_0 (estimate of false positives) was taken as the inverse of the probability the image with the smallest (here, zero) angular/linear transformation of the gloss map was selected as glossier. Error parameter E_1 (estimate of false negatives) was taken as the average probability the image(s) with the largest angular/linear transformation of their respective gloss maps was selected as glossier.

Results and discussion

Probabilities of individual images being selected as glossier after rotation of specular highlights are shown in Figure 4 as a function of angular offset between the specular highlight map and the diffuse (matte) surface images. This shows that perceived surface gloss was highest when there was no angular misalignment between specular highlights and the matte version of the St. Matthew image (0 degrees offset). As the angle of offset between the specular highlight map and the original diffuse image

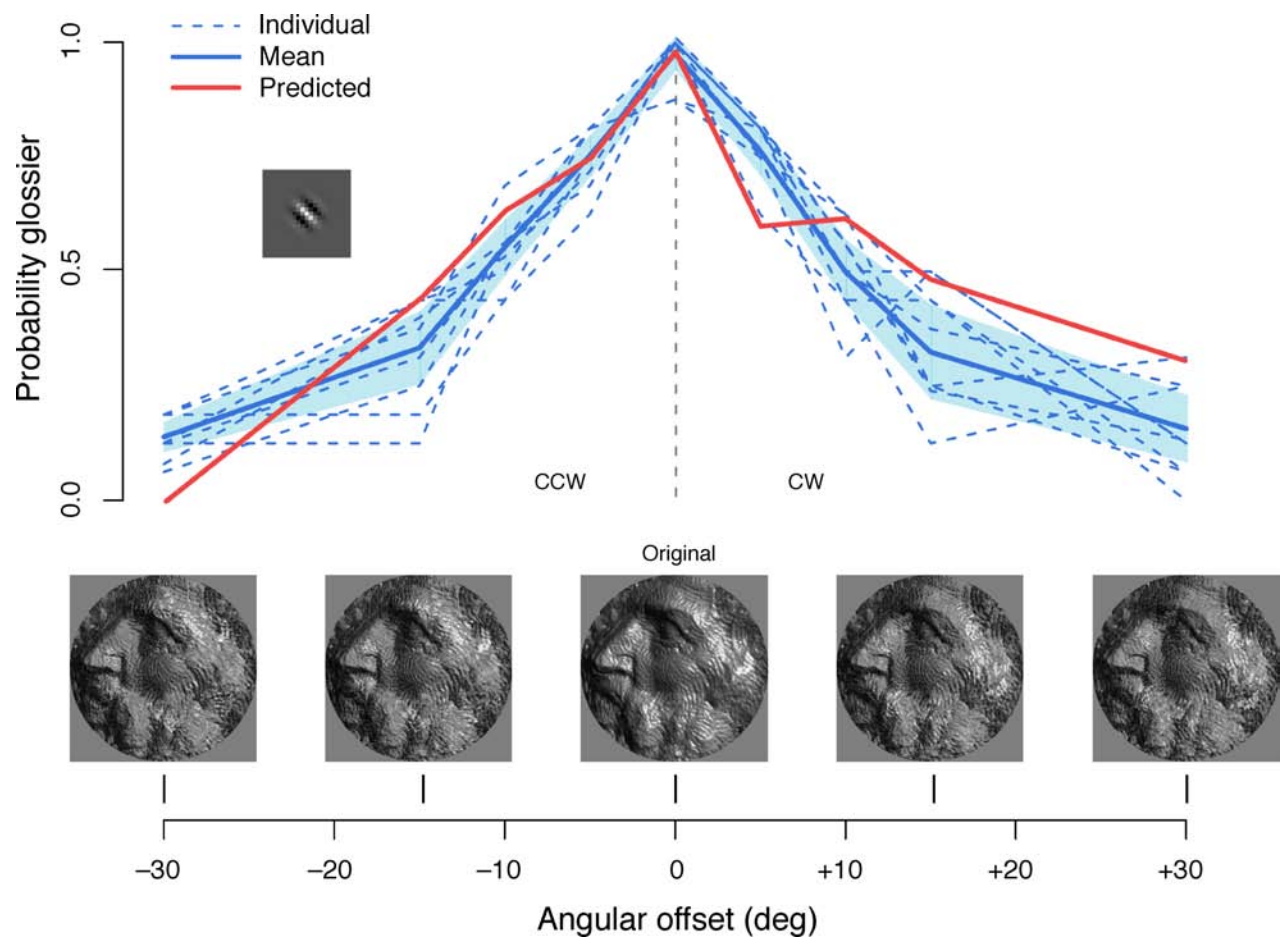


Figure 4. Probabilities of images being selected as glossier was found to decrease with increasing angular rotation of the specular highlights relative to the diffuse (matte) image of St. Matthew used in the Motoyoshi et al. (2007) study. Results shown here are from 2AFC trials consisting of randomly presented pairs of images differing in specular highlight offset in the counter-clockwise (CCW) and clockwise (CW) directions. Individual response probabilities from each participant are shown in broken blue lines, while the mean and 95% confidence band for these response probabilities is indicated by the solid blue line and faint blue band, respectively. Zero angular offset (i.e., 0 degrees) indicates presentation of the original glossy image where no rotational transformation of highlights was performed. This value is observed to be consistent with the optimum correlation between polar responses of oriented filters as shown in red.

increased either in the clockwise (CW) or counter-clockwise (CCW) direction, the probability of the resulting image appearing glossy decreased rapidly. The probability of an image appearing glossy fell below chance at angular rotations in excess of ± 10 degrees over the 60 degree range we examined. This pattern of perceived gloss in surfaces follows what would be expected from the correlation between the orientation fields of specular and diffuse layers after successive CW and CCW angular rotation of specular highlights of increasing amplitude away from their natural locations and orientations.

Similar effects were observed for the translation of specular highlights. Figure 5 shows the probability of surfaces in 2D images being selected as glossier in the 2AFC scenario where the highlight map was translated horizontally in the leftward direction at increments of 2% total image width. As the magnitude of horizontal (leftward) translation of highlights was increased, the overall

probability that the surface visible in the image appeared glossy was found to decrease monotonically over the 12% horizontal width of the image. The superimposed red trace shows the predicted outcome from the correlation of oriented filter responses to the matte and diffuse image layers over successive increments in horizontal translation.

The results of this experiment demonstrate that the perceived gloss of surfaces is highly sensitive to changes in the orientation of luminance increments with respect to the underlying diffuse shading profile of a surface that specifies the 3D geometric surface structure. Both rotational and translational offsets of the luminance increments in the difference map led to a consistent decline in the perceived gloss in the depicted surfaces. Observers' judgments are in good qualitative agreement with the prediction line based on the correlation between the orientation field of the gloss map and the orientation field of the diffuse shading profile. This suggests that the visual system may use such

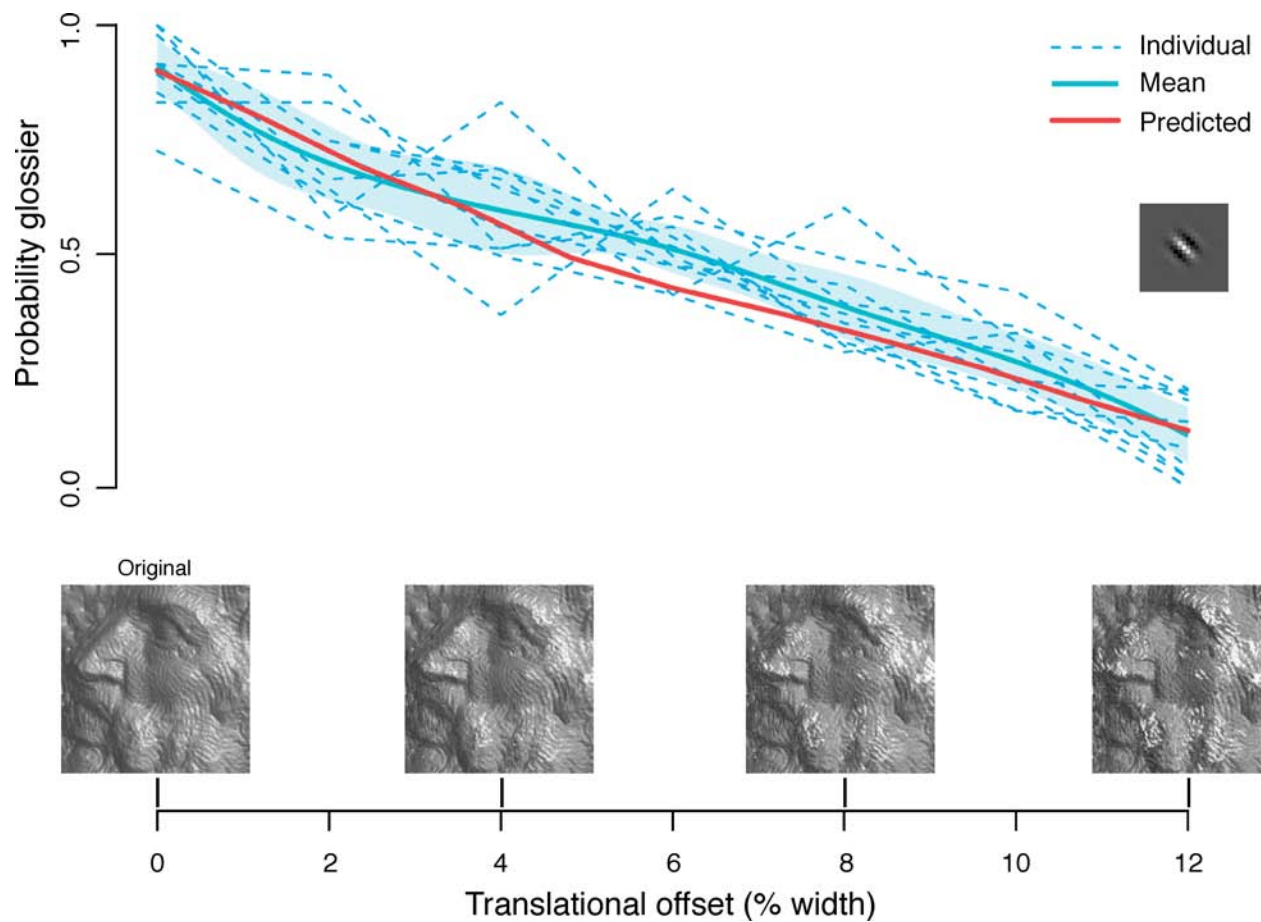


Figure 5. Probabilities of images being selected as glossier was found to decrease with increasing translational offset of the specular highlights (as a percentage of total image width) relative to the diffuse (matte) image of St. Matthew used in the Motoyoshi et al. (2007) study. Results shown here are from 2AFC trials consisting of randomly presented pairs of images differing in specular highlight offset to the left. Individual response probabilities from each participant are shown in broken blue lines, while the mean and 95% confidence band for these response probabilities is indicated by the solid blue line and faint blue band, respectively. Zero angular offset (i.e., 0%) indicates presentation of the original glossy image where no linear translation of highlights was performed. This value is observed to be consistent with the optimum correlation between polar responses of oriented filters as shown in red.

orientation fields, or similar mechanisms tuned for spatial luminance gradients, to evaluate the plausibility that a region of high luminance is consistent with a specular highlight or arises from another source (such as a pigment).

Gloss, skew, and highlight placement

The preceding results reveal that histogram skew does not correlate with the perception of surface gloss unless the luminance maxima are positioned and oriented in a manner consistent with the surface's diffuse shading profile. This suggests that histogram skew per se does not play a direct role in the perception of gloss in these images. However, it should be noted that the images containing displaced and rotated highlights no longer elicit percepts of uniform albedo surfaces. Although the manipulations we performed were all done using uniform

albedo surfaces, Motoyoshi et al. (2007) restricted their attention to surfaces that *appeared* as uniform albedo surfaces. We will return to this issue in the [General discussion](#), where we will argue that claims of this form are logically circular. At this juncture, we were interested in determining whether it is possible to elicit a percept of gloss in images that appear as uniform albedo surfaces with some highlights, even if the image has a negative skew.

To answer this question, we constructed a new variant of the St. Matthew image. If the visual system computes gloss by evaluating the consistency between the locations of the luminance maxima in an image relative to the luminance gradients that specify the 3D shape of surfaces, then the percept of gloss should be elicited by the presence of local luminance increments at the correct positions in the image even if the overall skew is consistent with a matte surface. This could be accomplished by placing some sparse luminance maxima at the locations where they would (approximately) appear on a glossy variant of the

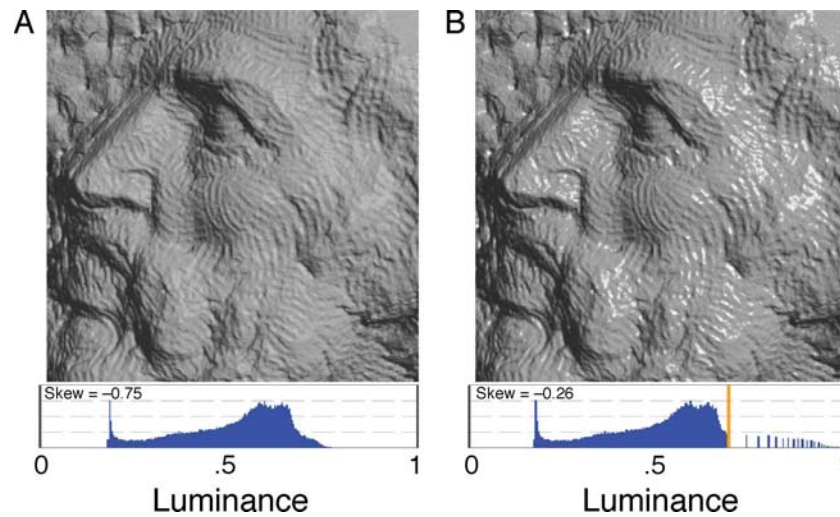


Figure 6. Log transformation applied to the upper tail of the luminance histogram for the matte image of St. Matthew changes the appearance of the luminance maxima in the image from diffuse reflection (A) to specular reflection (B). The specular appearance of the luminance maxima in B promotes the overall impression of the surface as being glossier than the surface in A.

image relative to the 3D geometry of the surface as specified by shading information. To accomplish this goal, we began with a negatively skewed matte version of the St. Matthew image. We then clipped the upper end of the histogram containing the brightest pixels and increased their intensity by applying a saturating non-linearity to the upper values in the histogram (in this example, a log transform; see Figure 6). A cutoff value for the application of this non-linearity was chosen such that the overall skew of the histogram would be negative following this transformation (here, the overall skew was -0.26). The resulting image (Figure 6B) appears as a uniform albedo, relatively glossy surface, despite having a negative skew.

This demonstration provides further evidence that skew does not play any direct causal role in the perception of surface gloss. Rather, the perception of gloss appears to involve a computational analysis that evaluates the consistency between the 3D shape of a surface and the locations and orientations of highlights.

Non-uniform albedo surfaces

One of the putative benefits of exploiting image statistics as sources of information about surface properties like gloss is that they could help circumvent the need to perform the kind of “inverse optics” analysis embodied in intrinsic image models. For image statistics to be useful, they must be able to distinguish between the different causes of luminance variation in an image. However, to this juncture, the only images that have been evaluated to support the claim that skew plays a role in the perception of gloss were those generated by uniform albedo surfaces with essentially identical surface relief, viewed in an approximately constant illumination field. It is therefore possible that the

correlations observed between skew and surface gloss may have only arisen because the other possible sources of image variation were never considered or evaluated.

Consider the image of unpolished granite depicted in Figure 7A. The luminance variations in the image arise primarily from the various colors of pigment within the granite but also contain shading information that arises from the 3D meso-structure within the rock. If histogram skew provides information about surface gloss, then varying the skew of the granite’s luminance histogram should impact on the perceived gloss of the granite texture. To test this hypothesis, we manipulated the histogram (preserving RMS contrast) in a manner that varied its skew between strongly positive to strongly negative. As can be seen in the images depicted in Figure 7B, this manipulation has no impact on the perceived gloss of the granite. Rather, granite images that are coerced into a positively skewed histogram appear to contain predominantly dark pigment interspersed with sparse light pigment, whereas negatively skewed images appear to be composed of predominantly light pigment interspersed with sparse dark pigment. This lightness difference still holds true (though less pronounced) when the means of these transformed histograms are subsequently matched (Figure 7C), in which case the matte appearance of the image also remains constant across very large changes in skew.

These demonstrations provide evidence that skew does not provide any unambiguous evidence for the gloss of a surface. Rather, these demonstrations reveal that image statistics like skew are incapable of distinguishing between variations in surface albedo and variations in surface gloss. In the demonstrations presented in Figure 7, the manipulations of the skew of the luminance histogram altered the perceived distribution of pigments in the granite, not perceived gloss. Similar effects can be experienced with the uniform albedo stucco images that were the focus of




Figure 7. Natural surfaces as in this image of granite are seen to be characteristically matte even though their luminance histograms are positively skewed (A). Gamma transformations of histogram skew over a range from positive to negative produce no change in the lack of gloss in these images (B). Instead, the primary perceptual change is an apparent increase in the overall lightness of the surface with decreasing skewness. However, this inverse relationship between lightness and skew is reversed for the lighter surface pigmentation when these histograms are subsequently matched in terms of mean luminance (C).




Figure 8. The negatively skewed image of a uniform-albedo stucco surface used in the Motoyoshi et al. (2007) study is matte in appearance (A) even after all the luminance information is numerically inverted to produce the appearance of an inhomogeneous matte surface with the opposite geometric profile (B). The inevitable inversion of luminance gradients from A to B produces a change in the physical appearance of local geometry from concave to convex, and vice versa. Regions of A that are masked by shadow appear in B as lighter pigmented material embedded in a darker surface.


Motoyoshi et al.'s (2007) studies (Figure 8A). The luminance histogram for a matte stucco surface is given below the image, which is negatively skewed. The histogram of this image can be transformed into a positively skewed histogram with the same shape and opposite skew by simply inverting the intensities of all of the pixels, creating the luminance negative of the seed image. As can be seen in Figure 8B, these positively skewed images do not generate a percept of gloss. Rather, this positively skewed surface appears just as matte as the original surface. Note that the regions that were shadows in the original image now appear as lightly colored pigment in the inverted image. Any explanation of this effect must embody mechanisms that can differentiate between luminance variations that arise from changes in reflectance and surface gloss, distinctions that are beyond the scope of skew computations.

The illumination field and histogram skew

The general problem confronted in deriving surface properties from images is that the 3D shape, reflectance properties, and the illumination field are conflated in the image. In Motoyoshi et al.'s (2007) experiments, the only variable that was free to vary across the class of images studied was surface gloss. The 3D shape, albedo, and illumination field were held constant, so surface gloss was the only source of variation that could contribute to the shape of the luminance histogram. The preceding demonstrates that any arbitrary skew can be created with surfaces that are non-uniform albedo. Here, we show that luminance histograms of arbitrary skew can be constructed by

manipulating the surface geometry and position of the predominant light source with surfaces of constant albedo.

Examples of positively skewed images that appear matte are presented in Figure 9. Consider first Figure 9A, which depicts an image of a roll of toilet tissue, which appears (and is) very matte (the histograms in this image are of the objects in the pictures, not of their backgrounds). The positive skew of this image arises from the interaction between the location of the light source and the cylindrical geometry of the roll. The bright tail of the luminance histogram arises from the highly illuminated region along the top of the roll, whereas the mass of the luminance distribution is darker because the majority of the surface only receives illumination from secondary (diffuse) illumination. Similarly, the stucco image in Figure 9B is similar to the matte surfaces employed in Motoyoshi et al.'s (2007) experiments but illuminated from a more oblique angle than they explored in their studies. From this illumination angle, only some of the largest regions of relief (bumps) along the surface are directly illuminated, giving the histogram a positive skew. This effect is especially pronounced with a simulated surface rendered with an obliquely oriented collimated light source (Figure 9C). Nonetheless, these surfaces appear completely matte. A similar interaction between the position of the light source and histogram skew is observed for glossy surfaces. This can be seen most readily for relatively flat glossy surfaces. If the surface normal of a glossy surface is oriented in a manner such that it bisects the angle between the viewing position and the predominant light source direction, the surface will strongly reflect the light source, causing the majority of pixels in an image to appear in the upper end of the luminance histogram, generating a negatively



skewed histogram. Such surfaces nonetheless appear glossy.

In summary, all of these demonstrations lead to the same conclusion: histogram skew does not provide diagnostic information about whether a surface is matte or glossy; this information can only be derived by computations that evaluate the consistency between its estimates of surface reflectance, 3D shape, and the illumination field.

Skew and adaptation

One of the core pieces of evidence used to support the claim that skew is explicitly computed was Motoyoshi et al.'s (2007) adaptation experiments. They reported that observers' judgments of surface gloss were strongly affected by being exposed to skewed adapting stimuli that were matched for mean luminance and contrast (in their studies, the standard deviation of luminance). They tested two types of adapting stimuli: those depicting surfaces of the type tested and an array of difference of Gaussian "blobs" (DoGs). Motoyoshi et al. presented positively and negatively skewed adaptors simultaneously and reported negative aftereffects to the adapting stimuli. Motoyoshi et al. argued that this result provides evidence that skew is explicitly computed by the visual system and used to infer surface properties such as lightness and gloss. We have recently completed a large series of adaptation experiments in an attempt to understand the effects reported by Motoyoshi et al. Due to the volume of experiments conducted and the variety of complex and subtle issues that arise in attempting to interpret the outcome of experiments of this kind, these experiments will be described in a subsequent paper. In the present paper, we focus on whether the logic of these adaptation experiments and the model they proposed to compute skew can provide any insight into the computation of surface properties such as gloss and lightness.

To derive insight into the kind of information that can be revealed by the computation of skew, consider the model proposed by Motoyoshi et al. (2007). The model consists of three stages: the detection of local increments and decrements by circularly symmetric receptive fields (on/off cells); the transformation of these responses by an accelerating non-linearity (a half-squaring operation); a pooling of the responses from each stream (separately) over space; and a differencing between these pooled responses. The first stage of this model detects the local sign of

Figure 9. Luminance histograms of images of uniform-albedo surfaces that appear matte, such as toilet paper (A) or uniform-albedo stucco cement (B) can be made to have positive skew when these surfaces are illuminated with a single light source positioned obliquely with respect to the surface normal oriented toward the image capture device. Similar effects are observed with a simulated surface illuminated from a collimated light source angled 70° from the average surface normal (C).

contrast; the non-linearity exaggerates the relative contribution of the extreme responses in the on and off streams; the pooling over space effectively “blurs” the responses over larger image regions; and the differencing stage detects an imbalance between the on and off streams. The intuition of this model is straightforward: it simply takes differences in the pooled responses of the on- and off-center-surround receptive fields. The existence of the distinct on- and off-pathways is well known; the main theoretical claim is that the imbalance of on and off filter responses provides information about surface gloss (and lightness). The problem with this claim is that the very same imbalance in responses can be generated by all of the other contributions to the image that were not considered in their original analysis, including surface pigmentation or the position of the illuminant. The only way the imbalance in filter responses can reliably provide information about surface gloss is if it is known *a priori* that the surface viewed has a uniform albedo, and the illumination direction is within a fixed range of possible angles relative to the surface. But this is precisely the kind of analysis that image statistics like skew were supposed to circumvent. Thus, all of the same ambiguities in the 3D shape, surface optics, and illumination field still exist for any given value of skew and cannot be resolved by the kind of model proposed.

With regard to adaptation experiments, it should be noted that any adaptation of the on-center receptive fields sensitive to local increments could generate the after-effects reported, irrespective of whether skew was explicitly computed or not. By definition, a glossy surface will contain local highlights. Thus, if mechanisms that detect local increments were desensitized, one would expect any subsequently viewed stimulus that requires observers to make judgments that depend on the operation of these mechanisms to exhibit a diminished response.

General discussion

In the preceding, we have argued that there is currently no evidence to support the view that histogram or sub-band skew is computed by the visual system to infer surface gloss or lightness. The theoretical basis of our arguments arises from a belief that the computation of surface properties requires a *photo-geometric* analysis. We contend that surfaces with different material and reflectance properties possess characteristic patterns of *correlations* between the intensities in the image that allow them to be distinguished. Histogram skew is, by definition, a purely photometric quantity that does not capture any of the spatial correlations that convey the structure of a surface. The correlation between skew and surface properties like gloss and/or lightness only occurs when the other variables that contribute to image structure are held constant.

In the Motoyoshi et al. (2007) studies using images of hand-made surfaces, gloss was the only parameter that was systematically varied; the illumination field, albedo, and surface geometry were all held constant. In this restricted context, there is a strong correlation between the physical gloss of surfaces and histogram (or sub-band) skew. Given that observers were able to discriminate the different degrees of surface gloss of these surfaces, their gloss ratings would have to correlate strongly with skew simply because skew is correlated with gloss in these particular images, even if skew is not explicitly computed by the visual system. Such data are a statistical inevitability that follows from the correlation between skew and physical gloss used in their studies and hence cannot provide any evidence for the role of skew in the perception of surface properties. Critically, the correlation between skew and gloss drops off precipitously when a broader range of surfaces and illumination environments is evaluated. We have shown that any histogram skew can be generated by non-uniform albedo surfaces containing a distribution of different surface pigments, and that the kinds of skew computations advocated by Motoyoshi et al. are incapable of distinguishing non-uniform albedo surfaces from surfaces that vary in gloss.

We have also shown that essentially any value of skew can be obtained with the kinds of matte surfaces investigated by Motoyoshi et al. (2007) by simply varying the position of the light source, without any concomitant change in perceived surface gloss. If skew were a cue to surface gloss, one would expect that such variations should lead to illusory shifts in the perceived gloss of a surface, but no such effects are observed. Finally, the correlation between skew and gloss of uniform albedo surfaces is much weaker when a broader range of surface geometries (i.e., surfaces with different meso-structure) is evaluated (r^2 values falling from .79 to .37), even when other variables (like the illumination field) are constrained.

In addition to the perception of gloss, we have argued that the correlation between skew and perceived lightness can be fully explained by the decision to equate the means of the images. It is well known that means are poor estimates of central tendency of skewed distributions and are “pulled” in the direction of extreme values (the tails of a skewed distribution). The negative correlation observed with skew and lightness could simply reflect a bias for observers to judge lightness on the basis of the luminance associated with the majority of pixels in the image. This interpretation is supported by the fact that the negative correlation between skew and perceived lightness is unaffected by phase scrambling their stimuli, which destroys the percept of a uniform albedo surface.

The theoretical intuition that motivates our argument against the claim that skew provides information about surface properties can perhaps be better appreciated by considering another complex perceptual problem—the perception of transparency. The perception of transparency requires the satisfaction of both geometric and photometric

constraints (Anderson, 1997, 2003a, 2003b; Anderson & Winawer, 2005, 2008; Gerbino, Stultiens, Troost, & de Weert, 1990; Metelli, 1974; Metelli, Da Pos, & Cavedon, 1985). The assertion that the percept of surface gloss is predicted by purely photometric (structure-blind) statistics like histogram skew is analogous to a claim that the perception of transparency can be predicted based solely on the image intensities in an image, irrespective of their distribution. But it is a well-established fact that the presence or absence of transparent surfaces cannot be determined by evaluating purely photometric information. Metelli's (1974) seminal work demonstrated that the percept of transparency can be abolished by simply redistributing the intensities in the image, or by breaking the geometric continuity of contours defining the boundaries of the transparent surface or the underlying surface (see also Anderson, 1999, 2003a, 2003b; Anderson & Winawer, 2005, 2008; Kanizsa, 1979; Singh & Hoffman, 1998). Our rotation and/or displacement of the luminance maxima (highlights) of glossy surfaces is in the same spirit as these manipulations and essentially modern variants of Beck and Prazdny's (1981) seminal demonstrations of the importance of highlight location and orientation on the perception of surface gloss.

Motoyoshi et al. (2007) appear to acknowledge that some additional information is needed to understand the role of skew in surface perception:

“While skewness is predictive of perceived surface qualities, it can of course be computed on arbitrary images, whether or not they look like surfaces. A picture of fireworks against the night sky will be positively skewed, but one cannot meaningfully judge its albedo or gloss; the same is true of the adapting stimulus of Figure 4a. Our findings were made in the case where the image is perceived as a surface of uniform albedo with some highlights. We do not know what aspects of image structure determine “surfacedness” or “highlights”. When our images are phase-scrambled so as to retain sub-band power, but not phase structure, they are typically seen as plausible but not convincing surfaces. The lightness effects are retained, but glossiness is lost. When the images are pixel-scrambled they are seen as two-dimensional noise patterns without a unitary albedo or gloss (p. 209).”

There are a number of ways that the qualifications in this passage can be interpreted. Motoyoshi et al.'s (2007) statement that “our findings were made in the case where the image is perceived as a surface of uniform albedo with some highlights” suggests that their claims about skew, lightness, and gloss are restricted to images that are *perceived* as uniform albedo surfaces with some highlights. The problem with a statement of this kind is that statistics like skew are supposed to provide the information

needed to *determine* the nature of the surface being viewed. Their statement suggests that their theory only applies to certain kinds of perceptual *outcomes*, the very thing that skew computations are supposed to help explain. In the above passage, Motoyoshi et al. acknowledge that statistics like skew cannot determine whether the image contains a surface of any kind, so it is difficult to understand how skew could provide a source of information about a specific kind of surface (i.e., glossy or matte, light or dark). Clearly, whatever additional geometric structure—i.e., spatial correlations in photometric information—is required to define a surface, uniform albedo or otherwise, cannot be derived from statistics like image skew. Their claim reduces to one of the form: histogram skew predicts the perceived albedo and gloss of a uniform albedo surface with some highlights if and only if a uniform albedo surface with some highlights is perceived. Even if the circular logic of this claim was accepted, the claim itself is demonstrably false. Skew does not predict the perception of gloss for surfaces that are perceived as uniform albedo surfaces when the illuminant is free to vary. As can be seen in Figure 9, the skew of a matte surface can be strongly positive by merely varying the position of the light source, without producing any appearance of surface gloss.

The preceding analysis reveals that skew computations are insufficient to derive surface and material properties. It should be noted, however, that Motoyoshi et al. (2007) acknowledged that some geometric analysis may also contribute to the perception of gloss:

“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 (p. 17, Supplementary material).”

Motoyoshi et al.'s (2007) construction “in addition to skewness computations” implies that skewness is but one kind of computation—a “cue”—that is involved in the perception of surface gloss. This construction suggests that skew is explicitly part of the computations involved in determining surface gloss but suggests that other analyses “beyond” those captured by sub-band filtering also play a role. This raises the question as to whether there is any compelling evidence that histogram or sub-band skew contributes a source of information about surface gloss independently from the “analysis of image structures beyond sub-band filtering” that we have argued are critical for the perception of gloss. Evidence of this type is needed to sustain any claim that skew plays a causal role in such percepts rather than simply being correlated with the variables used to compute skew. We contend that such data do not currently exist. If the perception of gloss *requires* the analysis of image structures beyond that captured by histogram or sub-band skew, then there is no

basis by which it can be asserted that sub-band skew plays any direct role in the perception of gloss. We have shown that skew fails to predict the perceived gloss of surfaces when the luminance maxima are not appropriately aligned with the underlying diffuse shading profile; when the illumination direction is systematically varied; or when non-uniform albedo surfaces are viewed. Any value of skew contains all of the same ambiguities in image interpretation that it was supposed to help resolve, namely, determining whether the luminance variations in an image arise from the illuminant, surface optics, or surface geometry.

Probably the “weakest” interpretation of the skew hypothesis is that the visual system tracks the relative difference in strength between mechanisms that detect increments and decrements in local contrast to infer skew. This interpretation is in keeping with the model Motoyoshi et al. (2007) offered as a method for computing skew, which essentially performs a weighted difference between on- and off-center cell responses that have been weighted by an accelerating non-linearity. It seems difficult to imagine any model of gloss perception that would assert that information about contrast sign is discarded. However, such information is incapable of specifying its environmental cause(s). Although the visual system clearly must retain information about decrements and increments to distinguish between features such as highlights and shadows, or dark and light pigments, such computations do not provide any leverage into resolving the ambiguities concerning the possible causes of a local increment or decrement in image contrast. All of the same ambiguities in the illumination field, surface optics, and reflectance (pigmentation) are preserved, which is precisely the problem of “inverse optics” that Motoyoshi et al. were attempting to “bypass.”

In the preceding, we have argued that the image statistics of histogram and/or sub-band skew do not provide any computational leverage into understanding the perception of gloss and albedo of surfaces. We have further argued that the computation of these properties relies on assessing the consistency between information specifying the 3D geometry of a surface, with the locations and orientations of luminance extrema that are potential highlights. This viewpoint is in keeping with a larger body of research that has demonstrated the critical role 3D geometry plays in the perception of surface gloss (Beck & Prazdny, 1981; Blake & Bühlhoff, 1990; Norman, Todd, & Orban, 2004; Todd et al., 2004; Wendt, Faul, & Mausfeld, 2008). From this perspective, there are a number of empirical and theoretical challenges that remain. One question involves understanding the precise nature of the computations used to perform such “consistency checks.” Using a simple and relatively crude means of estimating the orientation fields (namely, by using a single type of Gabor), we found that the orientation field of the gloss map and the diffuse shading profile of the matte surface must be correlated for gloss to be perceived. In these images, 3D shape was

specified by the 2D shading pattern. It is known, however, that the perception of gloss is strongly modulated by the 3D depth of the highlights relative to a surface body (Blake & Bühlhoff, 1990; Wendt et al., 2008), which suggests that more complex consistency checks must be performed in full cue scenes. It seems unlikely that image statistics of any kind could capture these types of dependencies. Any proposed statistic would have to be capable of reliably discriminating between the different contributions to image structure that are present in our experience of surfaces and materials, including the illumination field, 3D shape, and the reflectance properties of surfaces.

Acknowledgments

We kindly thank Isamu Motoyoshi for his open exchange over the many ideas presented in his paper. We would also like to thank R. Fleming for his views on many of the ideas presented in this paper and for generously providing the real world images presented in [Figure 9](#) demonstrating the role of illumination direction on histogram skew. We also thank J. Todd, G. Wendt, D. Wollschläger, and the reviewers for their comments on an earlier version of this manuscript. The research reported in this paper was supported by a grant from the ARC to B. Anderson.

Commercial relationships: none.

Corresponding author: Barton L. Anderson.

Email: barta@psych.usyd.edu.au.

Address: School of Psychology, 526 Griffith Taylor Building (A19), The University of Sydney, NSW 2006, Australia.

References

- Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: The role of contour junctions. *Perception*, *26*, 419–454. [[PubMed](#)]
- Anderson, B. L. (1999). Stereoscopic surface perception. *Neuron*, *24*, 919–928. [[PubMed](#)] [[Article](#)]
- Anderson, B. L. (2003a). The role of occlusion in the perception of depth, lightness, and opacity. *Psychological Review*, *110*, 762–784. [[PubMed](#)]
- Anderson, B. L. (2003b). The role of perceptual organization in White’s illusion. *Perception*, *32*, 269–284. [[PubMed](#)]
- Anderson, B. L., & Winawer, J. (2005). Image segmentation and lightness perception. *Nature*, *434*, 79–83. [[PubMed](#)]

- Anderson, B. L., & Winawer, J. (2008). Layered image representations and the computation of surface lightness. *Journal of Vision*, 8(7):18, 1–22, <http://journalofvision.org/8/7/18/>, doi:10.1167/8.7.18. [[PubMed](#)] [[Article](#)]
- Barrow, H. G., & Tenenbaum, J. M. (1978). Recovering intrinsic scene characteristics from images. In A. Hanson & R. Riseman (Eds.), *Computer vision systems* (pp. 3–26). New York: Academic Press.
- Beck, J., & Prazdny, K. (1981). Highlights and the perception of glossiness. *Perception & Psychophysics*, 30, 407–410. [[PubMed](#)]
- Ben-Shahar, O., & Zucker, S. (2001). On the perceptual organization of texture and shading flows: From a geometrical model to coherence computation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1048–1055). Kawaii, HI.
- Blake, A., & Bülthoff, H. (1990). Does the brain know the physics of specular reflection? *Nature*, 343, 165–168. [[PubMed](#)]
- Breton, P., & Zucker, S. W. (1996). Shadows and shading flow fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 782–789). San Francisco, CA.
- Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, 3(5):3, 347–368, <http://journalofvision.org/3/5/3/>, doi:10.1167/3.5.3. [[PubMed](#)] [[Article](#)]
- Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision*, 4(9):10, 798–820, <http://journalofvision.org/4/9/10/>, doi:10.1167/4.9.10. [[PubMed](#)] [[Article](#)]
- Gerbino, W., Stultiens, C. I., Troost, J. M., & de Weert, C. M. (1990). Transparent layer constancy. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 3–20. [[PubMed](#)]
- Huggins, P. S., Chen, H. F., Belhumeur, P. N., & Zucker, S. W. (2001). Finding folds: On the appearance and identification of occlusion, in CVPR'01. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (vol. 2, pp. 718–725).
- Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Praeger.
- Koenderink, J. J., & van Doorn, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta*, 27, 981–996.
- Metelli, F. (1974). The perception of transparency. *Scientific American*, 230, 90–98. [[PubMed](#)]
- Metelli, F., Da Pos, O., & Cavedon, A. (1985). Balanced and unbalanced, complete and partial transparency. *Perception & Psychophysics*, 38, 354–366. [[PubMed](#)]
- Motoyoshi, I., Nishida, S., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature*, 447, 206–209. [[PubMed](#)]
- Norman, J. F., Todd, J. T., & Orban, G. A. (2004). Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science*, 15, 565–570. [[PubMed](#)]
- 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*, 25, 846–865. [[PubMed](#)]
- Singh, M., & Hoffman, D. (1998). Part boundaries alter the perception of transparency. *Psychological Science*, 9, 370–378.
- Todd, J. T., Norman, J. F., & Mingolla, E. (2004). Lightness constancy in the presence of specular highlights. *Psychological Science*, 15, 33–39. [[PubMed](#)]
- Wendt, G., Faul, F., & Mausfeld, R. (2008). Highlight disparity contributes to the authenticity and strength of perceived glossiness. *Journal of Vision*, 8(1):14, 1–10, <http://journalofvision.org/8/1/14/>, doi:10.1167/8.1.14. [[PubMed](#)] [[Article](#)]