

Effects of illumination on the categorization of shiny materials

J. Farley Norman

Department of Psychological Sciences, Western
Kentucky University, Bowling Green, KY, USA



James T. Todd

Department of Psychology, Ohio State University,
Columbus, OH, USA



Flip Phillips

Department of Motion Picture Science, Rochester
Institute of Technology, Rochester,
NY, USA



The present research was designed to examine how patterns of illumination influence the perceptual categorization of metal, shiny black, and shiny white materials. The stimuli depicted three possible objects that were illuminated by five possible high-dynamic-range imaging light maps, which varied in their overall distributions of illuminant directions and intensities. The surfaces included a low roughness chrome material, a shiny black material, and a shiny white material with both diffuse and specular components. Observers rated each stimulus by adjusting four sliders to indicate their confidence that the depicted material was metal, shiny black, shiny white, or something else, and these adjustments were constrained so that the sum of all four settings was always 100%. The results revealed that the metal and shiny black categories are easily confused. For example, metal materials with low intensity light maps or a narrow range of illuminant directions are often judged as shiny black, whereas shiny black materials with high intensity light maps or a wide range of illuminant directions are often judged as metal. To discover the visual information on which these judgements are based, we measured several possible image statistics, and we found two that were highly correlated with the observers' confidence ratings in appropriate contexts. We also performed a spherical harmonic analysis on the different light maps to quantitatively predict how they would bias observers' judgments of metal and shiny black surfaces.

Effects of illumination on the categorization of shiny materials

During the past decade, there has been a growing amount of interest in the ability of observers to perceptually identify different types of surface

materials. For example, [Sharan, Rosenholtz, and Adelson \(2009, 2014\)](#) have shown that observers can rapidly identify material categories, such as metal, glass, or fabric, from briefly presented photographs of objects in natural settings (see also [Wiebel, Valsecchi, & Gegenfurtner, 2013](#)). Their results reveal that observers can achieve 80% accuracy with presentation times as low as 40 msec.

Although these findings provide clear evidence that observers can identify surface materials, they do not reveal the specific sources of information on which these judgments are based. There is some evidence to suggest that the perception of material properties may involve heuristic processes that can sometimes produce systematic errors. Much of this research has focused on the perception of gloss ([Adams et al., 2018](#); [Doerschner et al., 2010](#); [Marlow & Anderson, 2013](#), [Marlow & Anderson, 2015](#); [Mooney & Anderson, 2014](#); [Nishida & Shinya, 1998](#); [Olkkonen & Brainard, 2010](#), [Olkkonen & Brainard, 2011](#); [Pont & te Pas, 2006](#); [Zhang et al., 2015](#)) and the perception of translucency ([Fleming & Bühlhoff, 2005](#); [Marlow, Kim, & Anderson, 2017](#); [Xiao et al., 2014](#)). The results show clearly that observers' judgments of these properties can be influenced by factors that are physically independent of an object's material composition, such as its three-dimensional (3D) shape or its pattern of illumination.

[Todd and Norman \(2018\)](#) have recently reported that the pattern of illumination can also influence the perceptual distinction between shiny metal and shiny black dielectric materials, such as obsidian. It is important to recognize that the reflection of light from both of these material classes is entirely specular, so that there is only a tiny range of incident angles for each local surface region that will reflect any light toward the point of observation. To better understand the differences between these materials, it is useful to

Citation: Norman, J. F., Todd, J. T., & Phillips, F. (2020). Effects of illumination on the categorization of shiny materials. *Journal of Vision*, 20(5):2, 1–16, <https://doi.org/10.1167/jov.20.5.2>.






Figure 1. Specular reflectance as a function of the incident angle for silver, chromium, and obsidian in unpolarized light. The numbers in parentheses show the real and complex coefficients of the IOR.

consider how reflectance varies as a function of the incident angle of illumination based on the Fresnel equations. Figure 1 shows the reflectance functions for silver, chrome, and obsidian (i.e., volcanic glass). Note that silver reflects almost 100% of the incident illumination at all incident angles, whereas obsidian (or shiny black plastic) reflects close to 0% except at relatively high incident angles. Chrome is roughly midway between those two extremes.

The perceptual distinction between metal and obsidian is similar in some respects to the classic problem of lightness constancy, in that the luminance of any given surface patch is determined by the product of its reflectance and illumination. By analogy, this suggests that metal and shiny black materials could potentially be confused by selectively manipulating the intensity of illumination. Consider, for example, the two images of a boy's bust in the top row of Figure 2. Both objects are illuminated by an high-dynamic-range imaging (HDRI) light probe of an esplanade, and they both have the same magnitude of illumination and the same camera exposure. The one on the left depicts a low roughness chrome material that is perceived as metal. The image on the right depicts a low roughness obsidian material that is perceived as shiny black. It is important to keep in mind that variations in reflectance can be offset by variations in the magnitude of illumination or camera exposure. The bottom left panel of Figure 2 shows a low roughness chrome material with a five-fold reduction in illumination that is perceived as shiny black, and the bottom right panel shows a low roughness obsidian material with a five-fold increase in illumination that is perceived as metal.

It is interesting to note that there is a possible source of information that could potentially be used to perceptually distinguish the metal surface from

the shiny black one with high illumination. Because of the Fresnel effect, the peak specular reflections on metal surfaces tend to be in surface regions that face toward the point of observation, whereas the peak specular reflections of shiny black surfaces tend to be in peripheral regions near smooth occlusion boundaries that face away from the point of observation (see Todd & Norman, 2018). With careful inspection, this can be observed in Figure 2. Note that the shiny black object in the lower right panel has a distinct brightening along much of its occlusion contour, which is not present in the metal object in the upper left panel. This effect is rather subtle, however, and does not seem to have much impact on how these objects are perceived.

Another important factor that can influence the perceptual distinction between metal and shiny black materials is the distribution of illuminant directions. Whenever that distribution is relatively broad, the specular reflections on shiny black materials will be sparser than those that occur on metal. This is because shiny black materials have a wide range of incident angles from 0° to 60° for which only a tiny portion of the incident illumination is reflected (Figure 1). Although this is a potentially useful source of information, a similar sparse pattern of specular reflections can also occur on metal surfaces that are illuminated from a sparse set of directions. Consider the three images of a metal object that are shown in Figure 3. The object in the left panel is illuminated by an HDRI light map of an exhibit hall with a sparse pattern of lights in an otherwise dark environment. It is perceived as a shiny black material. The object depicted in the middle panel is illuminated by a light map of an esplanade, which has an intermediate range of directions. The depicted material in that case is perceived as metal. Finally, the image in the right panel is illuminated by a light map of a snowfield on a cloudy day, which is similar to a Ganzfeld because there is incident light from almost all directions. When that image is shown to human observers, they express uncertainty about whether it is metal, shiny white, or something else. Note that these perceptual distinctions are based entirely on the pattern of illumination because all of the depicted objects are composed of the same metal material, and they all have exactly the same 3D shape.

The research described in the present article was designed to provide a more rigorous exploration of the effects demonstrated in Figures 2 and 3. Observers made material category confidence ratings for objects composed of metal, shiny black, or shiny white materials over a wide range of illuminations. The results reveal that the perceptual categorization of these materials is only loosely coupled to the ground truth because observers' judgments are also heavily influenced by the pattern and intensity of illumination.

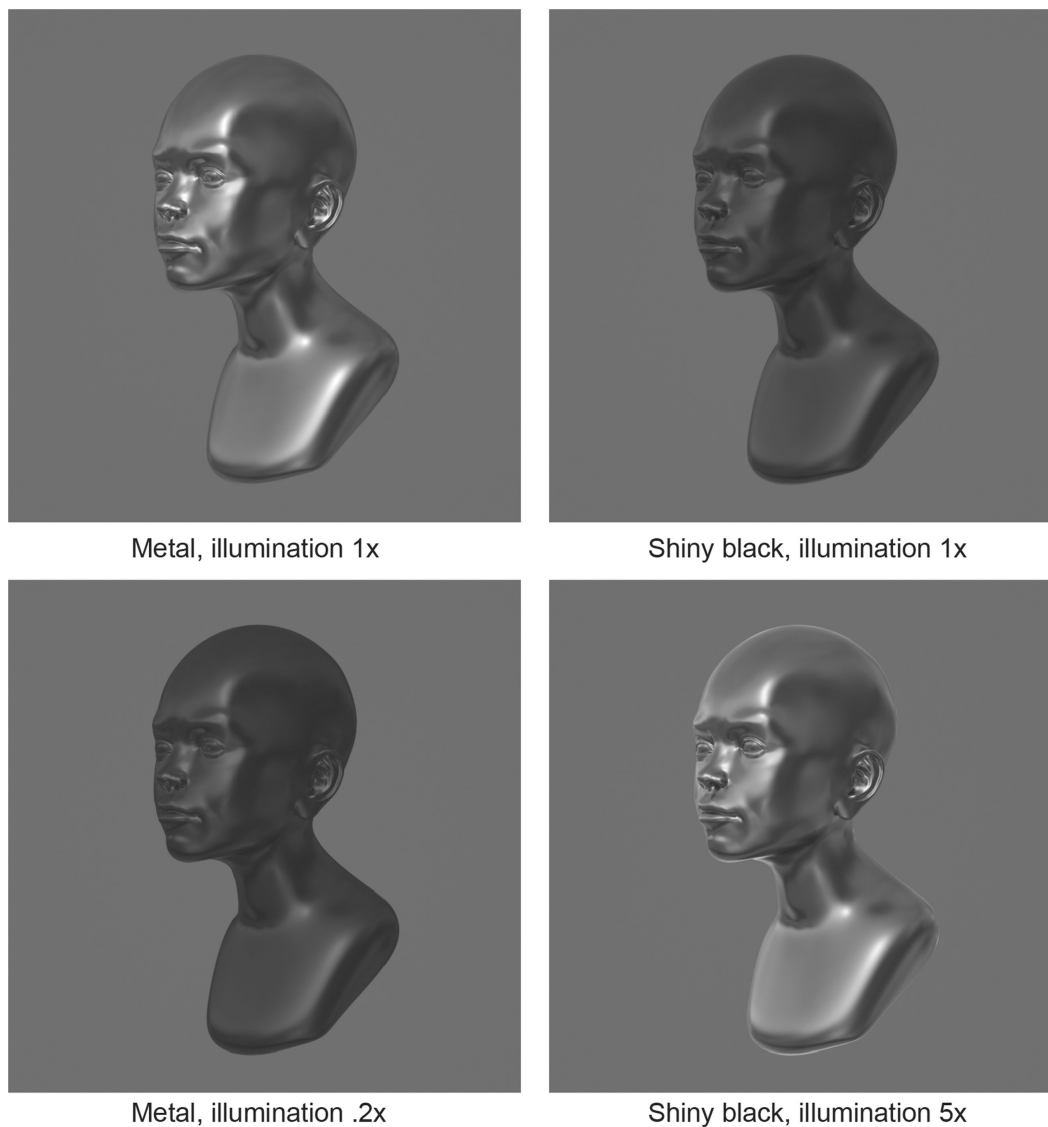


Figure 2. Four images of a boy's bust made with chromium and shiny black materials. All of the objects are illuminated by the same HDRI light map, although the overall intensity was varied as shown in the caption below each image.




Figure 3. Three images of a chromium object illuminated by three different HDRI light maps. The one on the left has a narrow range of illuminant directions and is perceived as shiny black; the one in the middle has an intermediate range and is perceived as metal; and one on the right has a very broad range of directions, and is perceptually ambiguous between metal, shiny white, and something else.




Figure 4. Five HDRI light maps used in the present experiment with variable distributions of illuminant directions and intensities.

Methods

Material simulations

All of the rendered images presented in this article were created using the Maxwell Renderer developed by Next Limit Technologies (Madrid, Spain). Maxwell is an unbiased renderer in that it does not use heuristics to speed up rendering times at the cost of physical accuracy. The depicted scenes were illuminated by the five desaturated HDRI light maps shown in Figure 4. They depict an empty exhibit hall, an atrium, an esplanade, an empty white room, and a snowfield on a cloudy day. These light maps were chosen specifically to vary the distributions of illuminant directions. For example, the exhibit hall light map has a very sparse set of illuminant directions; the atrium and the esplanade light maps have an intermediate range of illuminant directions; and the white room and snowfield light maps have a broad range of illuminant directions.

To appear metallic or shiny, it is important for a material to have a low level of roughness (Todd & Norman, 2018), but there are very few objects in the natural environment that are perfectly smooth with a

roughness of zero. Thus the materials in the present study were modeled using a relatively low roughness of 15. These included a chromium material, whose real and imaginary components of the index of refraction (IOR) were 3.21 and 3.30, respectively. A shiny black material was also included with an IOR of (1.51, 0). Finally, a shiny white material was simulated with a specular reflectance that was identical to the shiny black material, combined with an equal proportion of diffuse reflectance with a roughness of 95 and an IOR of (3, 0). Figure 5 shows examples of these three materials for a single object illuminated by the atrium light map. A more detailed discussion of how the complex IOR influences reflections is provided in the Appendix.

Apparatus

The experimental stimulus images were displayed by an Apple Mac Pro computer (Dual Quad-Core processors, with ATI Radeon HD 5770 hardware-accelerated graphics; Apple Inc., Cupertino, California, United States) using an Apple 27-in. LED Cinema Display (2560 × 1440 pixel resolution). The monitor was located at a 100-cm viewing distance. The luminous




Figure 5. Images of shiny black (left), metal (middle), and shiny white (right) materials, with the same 3D shape and pattern of illumination.

intensity of the monitor, measured over an area of 25° , had a minimum intensity (for black) of 1 cd/m^2 and a maximum intensity (for white) of 136 cd/m^2 .

Procedure

The stimuli depicted three different 3D objects, including the bust of a young boy, a circular disk with seven small bumps, and a randomly deformed sphere (Figures 2, 3, and 5). The material composition of these objects could be metal, shiny black, or shiny white, and they were illuminated by the five possible HDRI light maps shown in Figure 4. The relative intensities of the light maps were adjusted so that the metal objects would all have approximately the same maximum luminance that was just below the maximum image intensity of 255. This was done to ensure that there was no saturation in the depicted specular highlights. These same intensities were employed for the shiny black and shiny white stimuli. However, we also included additional images of metal surfaces for which the base illumination intensity was decreased by a factor of five, and shiny black surfaces for which the base illumination intensity was increased by a factor of five (Figure 2). We made small downward adjustments to the high illuminations when necessary to avoid saturation of the specular highlights.

The rendered images were globally tone mapped for the Apple monitor into the sRGB 2.1 color space with a D65 white-point and a γ of 2.2. No other global histogram adjustments (e.g., tint or burn) or local sharpening or contrast enhancement operators were used. Because the intensity of the light maps were adjusted to prevent saturation of the specular highlights, and we did not compress the dynamic range of intensities, this likely caused some loss of information at lower intensities that might have been visible on a display device with a higher dynamic range. All of the depicted stimulus objects were presented against a uniform gray background (with an intensity

of 100), which was created using an environment map. Although the background intensity can have a large effect on lightness judgments, it is not currently known whether similar effects may occur for the categorization of shiny materials. Thus the backgrounds were held constant in the present experiment to prevent any possible interactions with the other manipulated display variables.

On each trial, observers were presented with a single image, and were required to categorize the depicted material by adjusting four sliders with a hand-held mouse. Each of the sliders represented a different category labeled metal, shiny black, shiny white, or something else, and a digital readout was also provided for each one. Observers were instructed to adjust the sliders to indicate their confidence rating for each of the four possible categories. These confidence ratings were constrained by the program so that the four different ratings would always sum to 100%.

Observers

The 75 stimulus images were judged by one of the authors (JFN), and seven other observers who were completely naive about the purpose of the experiment or how the displays were generated. All observers possessed normal or corrected-to-normal visual acuity. During each experimental session, observers made judgments for all of the 75 stimuli. At the beginning of each session, the details of the response task were explained, and observers were shown real physical examples of metal, shiny black, and shiny white materials. All observers participated in two sessions on separate days.

Results

Figure 6 shows the average confidence rating for each of the four possible response categories

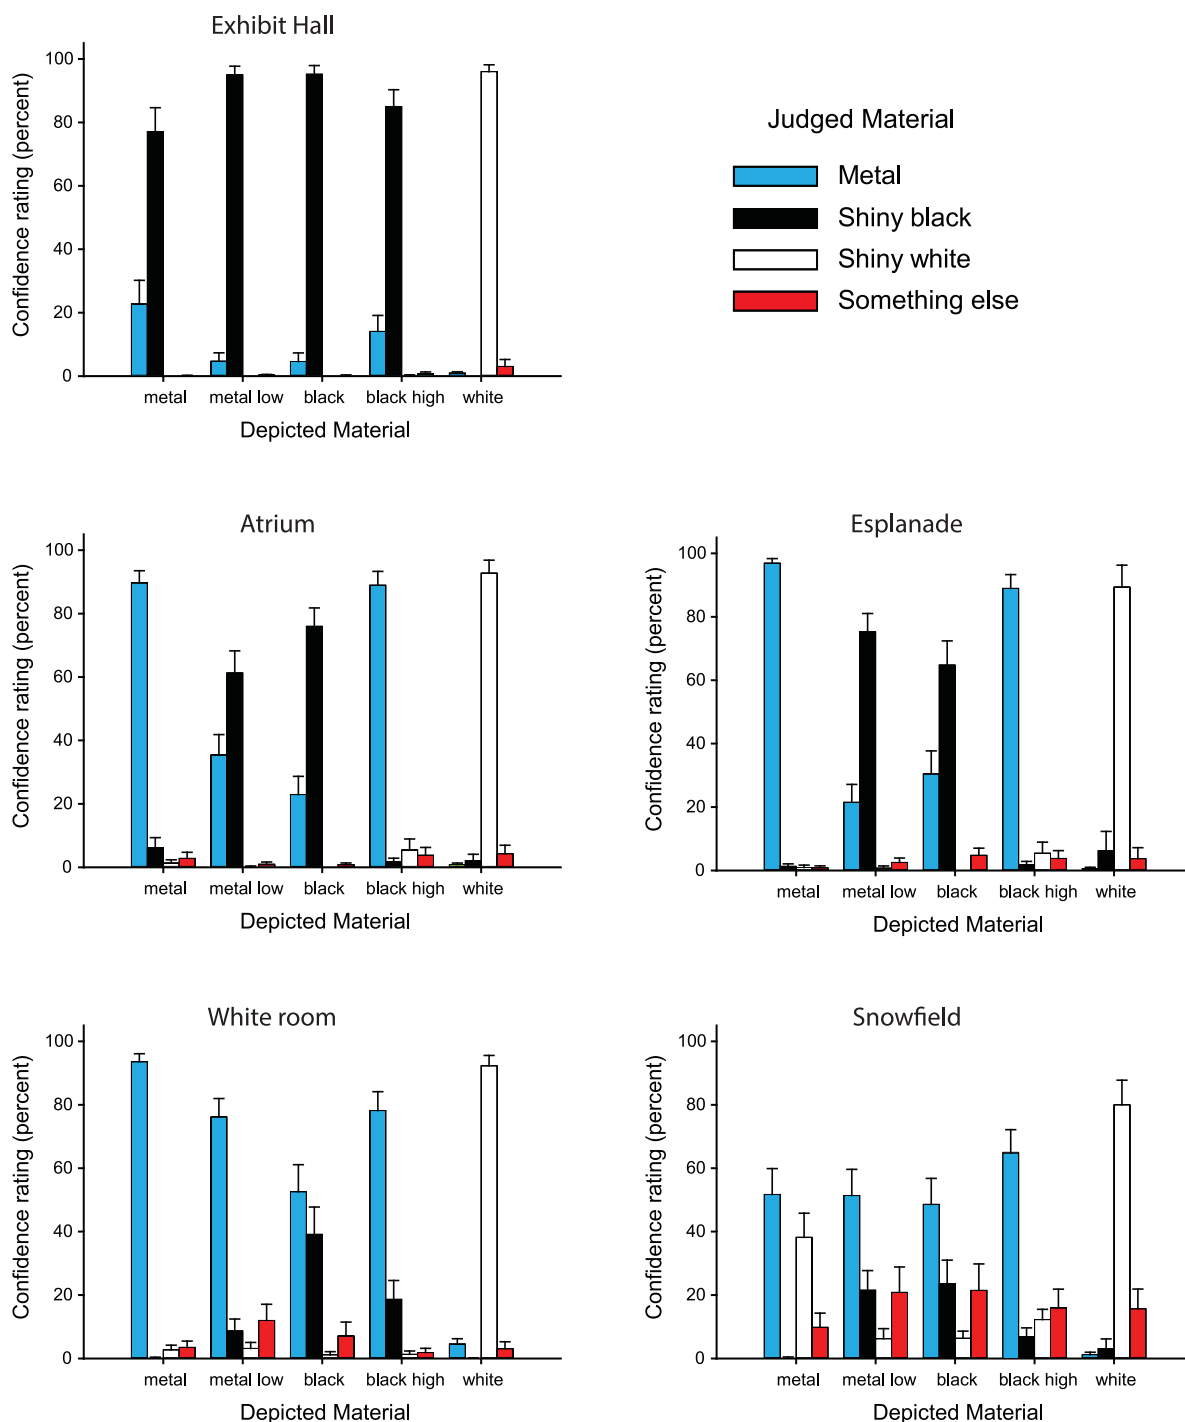


Figure 6. The average confidence rating for each of the four possible response categories for all of the different illumination and material conditions, collapsed over objects and observers. Error bars show the standard error of the mean for each condition.

collapsed over objects and observers. Each of the individual bar graphs shows the average data for a single light map for all possible combinations of materials and illumination intensities. Let us first consider the results for the shiny white material. Note in Figure 6 that these were categorized as shiny white with a high confidence rating for all five light maps. This

finding suggests that the categorization of shiny white materials is only minimally influenced by the pattern of illumination.

There was much more confusion between the metal and shiny black materials, and the pattern of illumination had a much larger influence on the perception of those materials. For example, when

images were rendered using the exhibit hall light map with a sparse distribution of illuminant directions, all of the metal objects were categorized as shiny black with a high confidence rating. A quite different pattern of results was obtained with the atrium and esplanade light maps, which had intermediate distributions of illuminant directions: the metal material with a base illumination intensity was rated primarily as metal, and the shiny black material with the base illumination intensity was rated primarily as shiny black. However, when the metal objects were presented with an illumination intensity that was five times lower, they were rated primarily as shiny black, and when the shiny black objects were presented with an illumination intensity that was five times higher, they were rated primarily as metal. These effects are all demonstrated in [Figure 2](#), which was rendered using the esplanade light map.

When the images were rendered with the snowfield or white room light maps, which had broad distributions of illuminant directions, the results were the opposite of those obtained with the exhibit hall. That is to say there was a general bias to judge all of the shiny black materials as metal. Observers' judgments for the snowfield light map stand out from the others in several respects. First, the primary confusion for the base illumination metal objects was shiny white rather than shiny black. Second, the combined metal and shiny black confidence ratings were lower than in the other conditions; and third, that was the only light map for which the depicted materials were categorized as "something else" with a rating that was significantly above zero. It is important to note that the distribution of illuminant directions for the snowfield light map is close to a Ganzfeld. Images of objects that are illuminated in that manner can look a bit weird (e.g., see right panel of [Figure 3](#)) because they have so little contrast. These findings highlight an interesting problem of how observers can distinguish between diffuse and specular reflections when they both have the same color. The problem arises because specular reflections from a Ganzfeld are quite similar to diffuse reflections, although they may sometimes be distinguishable due to specular inter-reflections in concave regions (e.g., see right panel of [Figure 3](#)).

To provide a simple quantitative measure of how each light map biased the observers' judgments, we calculated the average metal confidence rating for all of the metal and shiny black materials, as well as the average shiny black confidence rating. The ratio of these two averages provides a bias index for any particular light map. For the five light maps used in the present study, the bias index was 0.13 for the exhibit hall, 1.63 for the atrium, 1.66 for the esplanade, 4.16 for the white room, and 4.51 for the snowfield. These values indicate that the exhibit hall produces a strong bias to perceive purely specular surfaces as shiny black. The atrium and

the esplanade produce small biases to perceive purely specular surfaces as metal, whereas the white room and snowfield produce much stronger biases to perceive those surfaces as metal.

[Marlow and Anderson \(2013\)](#) have argued that visual information for the perception of gloss has three component dimensions. One of these, called specular contrast, refers to perceived differences between the diffuse and specular components of reflection. Note for example that images of shiny black materials typically have much higher specular contrast than those that depict shiny white materials. The second component, called specular sharpness, refers to the perceived steepness of the luminance gradients along the edges of highlights. The third component of their model is called specular coverage, and it refers to the proportion of an object's surface that is perceived to be covered by specular reflections. This is the component that is most affected by the pattern of illumination, and we suspect it may be an important source of information for distinguishing different types of shiny materials, such as metal or obsidian.

It is important to keep in mind that the three components of gloss proposed by [Marlow and Anderson \(2013\)](#) are all perceptual properties, and they cannot be measured in visual images without obtaining perceptual judgments. Although their analysis was an important inspiration for the one described here, we have focused instead on physical measures of image structure rather than perceptual ones. One of these that we refer to as the percentage of bright pixels (PBP) is designed as a physical analog to the Marlow and Anderson concept of coverage. The logic of this measure for the categorization of shiny materials is based on the relative reflectance curves of metals and dielectric materials shown in [Figure 1](#). Note that dielectric materials (e.g., obsidian) produce negligible amounts of specular reflection except at high incidence angles, whereas metals produce substantial specular reflections at all incidence angles. It is important to keep in mind that there is only a tiny range of incident angles for each local surface region that will produce specular reflections toward the point of observation. This occurs when the surface normal comes close to bisecting the angle between the direction of illumination and the viewing direction. If the light field has a reasonably broad range of illumination directions, then most local regions on metal surfaces will contain visible specular reflections. However, that is not the case for dielectric materials. Because of the Fresnel effect, specular reflections on those surfaces will be primarily located in peripheral regions near smooth occlusion contours, where there is a sufficiently high angle between the surface normal and the viewing direction.

Our specific method for measuring the PBP involves setting a threshold intensity value and counting the number of pixels with an intensity above that threshold, excluding the background. To test the perceptual




Figure 7. The metal and shiny black confidence ratings plotted as a function of the PBP for each of the five light maps used in the present experiment.

relevance of this measure, we calculated the PBP for each of the shiny black and metal stimulus images used in the experiment. The shiny white stimuli were excluded from this analysis because they contained both diffuse and specular reflections. After some trial and error, we found that a threshold of 50 produced the best fits to the empirical data. The left panel of Figure 7 shows the shiny black confidence ratings as a function of the PBP, and a similar plot for the metal ratings is shown in the right panel. For the shiny black judgments, there was a strong linear correlation with the PBP ($R^2 = 0.84$). This relation was more complex for the metal judgments, producing an R^2 of 0.69. The outliers in that case included all the stimulus objects illuminated by the snowfield light map. These all had PBP values in excess of 90%, yet the average metal confidence rating in those conditions was only 54%.

We also performed a similar analysis using the mean intensity of the images rather than PBP. Although these measures covary to some extent, the PBP measure counts all pixels with an intensity above 50 as equal, whereas the mean intensity weights the brighter pixels more heavily. The linear correlations of mean intensity with the shiny black and metal confidence ratings produced R^2 values of 0.61 and 0.50, respectively. Thus the mean intensity accounts for substantially less variance than the PBP measure.

Although the PBP measure does a reasonable job of distinguishing shiny black and metal materials in four of the five lighting environments we employed, it cannot predict observers' responses to metal surfaces illuminated by the snowfield light map, and it cannot distinguish metal and shiny white surfaces. To better understand those conditions, we identified all the displays with a PBP above 75%, and carefully observed them to see how they differ from one another. The upper row of Figure 8 shows three examples that depict a metal and shiny white material illuminated by the esplanade light map, and a metal material illuminated by the snowfield light map. Note that these images have

very different contrasts. We first tried to measure that by calculating the standard deviations of the image intensity distributions, but this did not provide a good account of the observers' judgments.

We then considered whether local contrast (as opposed to global) may be more perceptually relevant. To do that we applied an edge filter to all the displays with a PBP of 75% or higher to identify the regions with high local contrast. The bottom row of Figure 8 shows the results of that filtering for three images in the top row. Note that the metal (esplanade) one has the highest percentage of high contrast regions; the shiny white one has the lowest; and the metal (snowfield) one is somewhere in the middle. To formalize that we measured the percentage of high contrast regions for all of the images with a PBP greater than 75%, and correlated those measures with the observers' judgments. This was achieved by setting an intensity threshold of 220 on the edge filtered images, and counting the percentage of pixels below that value (excluding the background). It was immediately clear from this analysis that the three objects employed in the experiment produced noticeably different results, so we analyzed each object separately. The results are shown in Figure 9. The solid curves in that figure show the best fits to the data using logistic regression, which produced R^2 values of 0.88 for the cobblestone object, 0.82 for the boy's bust, and 0.95 for the distorted sphere. Note that these objects differ from one another in terms of local surface curvature, which can also affect local image contrast. It appears that observers may have considered that in making their judgments so that objects with many high curvature regions (like the cobblestones) require a greater percentage of high contrast regions to be perceived as metal.

It is best to be cautious about drawing any strong conclusions with respect to the specific analyses shown in Figures 7 and 9. There are many possible measures that are conceptually similar to the ones we adopted, and it is likely that one of these may eventually provide




Figure 8. A cobblestone surface with different materials and patterns of illumination. These include a metal material with the esplanade light map (left), a metal material with the snowfield light map (middle), and a shiny white material with the esplanade light map (right). The bottom row shows edge filtered versions of the images in the top row.

a better account of the perceptual categorization of shiny materials. One interesting distinction between our approach and others that have been proposed in the literature (e.g., [Marlow & Anderson, 2013](#)) concerns the need to separate the diffuse and specular components of reflection to make judgments about glossy materials. Our strategy has been to avoid that issue by analyzing diffuse and specular components as a single underlying pattern of luminance. Evaluating the success of that strategy will remain as an interesting issue for future research.

Discussion

It is important to keep in mind that for any local neighborhood of a purely specular surface, there is only a tiny range of incident angles that will reflect any light toward the point of observation. Thus the PBP for these materials is significantly influenced by the distribution of illuminant directions. The present experiment was designed to measure the perceptual significance of this effect by manipulating the patterns of illumination to produce systematic variations in the PBP. The results

reveal that the pattern of illumination can influence the categorization of shiny materials almost as much as variations in the actual depicted material.

The statistical structure of illumination

The light maps employed in the present experiment were selected based on a subjective evaluation of the overall distribution of illuminant directions, but it would be useful if these distributions could be quantified in some way. One possible method to achieve this is to decompose the spherical function of the illumination environment by the sum of its spherical harmonics ([Doerschner, Boyaci, & Maloney, 2007](#); [Mury, Pont, & Koenderink, 2007](#)), which is analogous to a Fourier analysis of planar patterns. The power of the spherical harmonics at each order characterizes the angular distributions of the illumination at that order. The zero order component of a spherical harmonic series has a single coefficient that represents the intensity of spherically diffuse ambient light. The first order has three coefficients, each of which represents a light dipole with a source in one direction and a sink in the opposite direction. The relative directions of




Figure 9. The average confidence rating for all the images with high PBPs as a function of the percentage of pixels with high local contrast. The solid curves in each graph show the best fits to the data using logistic regression.

the three dipoles are all orthogonal to one another. Each subsequent order in the series adds two additional coefficients that represent larger and larger groupings of sources and sinks at finer and finer scales.

Figure 10 shows the relative power of the different components of a spherical harmonic series (up to the




Figure 10. The relative power weightings of the different components of a spherical harmonic analysis (up to the 10th order) for each of the five light maps used in the present experiment.

	Bias Index	Diffuseness	Brilliance	Diffuseness2
Exhibit hall	0.13	0.74	0.79	0.72
Atrium	1.63	0.36	0.46	0.47
Esplanade	1.66	0.6	0.55	0.53
Room	4.16	0.74	0.53	0.38
Snowfield	4.51	0.17	0.51	0.21
		$R^2 = 0.18$	$R^2 = 0.40$	$R^2 = 0.88$

Table 1. The bias index for each light map used in the present experiment, and three possible measures of the spherical harmonic series for each map. The bottom row shows the correlations (R^2) between each measure and the bias index.

10th order) for each of the five light maps used in the present experiment. Each point on these curves represents the root mean squared power for all the coefficients at a given order. Note that the energy drops off quickly with increasing order and that it mostly levels out near zero by the seventh order. Zhang et al. (2019) have described two possible metrics for describing these distributions: A diffuseness metric (Xia, Pont, & Heynderickx, 2017) is defined as the ratio of the power of the first order harmonic relative to the power of the zero order; and a brilliance metric is defined as the ratio between the sum of the higher harmonics greater than or equal to the third order, relative to the sum of all orders. The values of these metrics (computed to the 30th order for brilliance) for each of the five light maps used in the present study are shown in Table 1, together with the bias index for each map that was computed from the observers' confidence ratings. This table also shows that there is a relatively small correlation between each of these metrics and the bias index, resulting in R^2 values of 0.18 and 0.40, respectively.

A closer examination of Figure 10 reveals that most of the variance among these particular light maps




Figure 11. High pass (left) and low pass (right) versions of the atrium (top) and esplanade (bottom) light maps from Figure 4.

occurs at orders zero and two. Could the relation between those orders provide a better account of the observers' biases? The fourth column in Table 1 shows another possible metric labeled *diffuseness2*. Whereas *diffuseness* is the first order component divided by the zero order, *diffuseness2* is the second order component divided by the zero order. Note that the correlation of that measure with the bias index produced an R^2 of 0.88, more than twice the values obtained for the *brilliance* or *diffuseness* metrics. Because we have no theoretical explanation to justify this particular metric, it is possible that its high correlation with the bias index could be an accidental property of these particular light maps. Nevertheless, among the wide variety of measures we have considered, this is the only one that provides a good fit to the empirical data.

Todd and Norman (2018) have recently demonstrated that it is possible to alter the apparent material of an object between metal and shiny white by selectively filtering the pattern of illumination. To follow up on their observations, we created a new set of light maps by filtering some of the ones shown in Figure 4, each of which had a spatial resolution of 4800×2400 pixels. Four of these new maps are shown in Figure 11. The ones in the right column were created in HDRshop using an 800 pixel wide Gaussian blur filter on the images of the atrium (top) and esplanade (bottom). This produces low pass filtered images that only contain the lower frequency components of the original light maps shown in Figure 4. The images in the left column of Figure 11 were created using a 10-pixel-wide Gaussian blur filter on the images of the atrium (top) and esplanade (bottom), and then subtracting the blurred images from the originals. This

produces high pass filtered images that only contain the higher frequency components of the original light maps shown in Figure 4.

Figure 12 shows three images of a boy's bust made of polished chrome with different patterns of illumination. The image in the middle panel was created with the atrium light map shown in Figure 4. The one on the right was created using the low pass filtered version of the atrium, and the one on the left was created using the high pass filtered version. Note how these variations in illumination have a huge effect on the appearance of the surface materials. The one created with the original atrium light map appears as metal. The one created with the low pass version appears as shiny white (see also Todd & Norman, 2018), and the one created with the high pass version is perceived as shiny black. It is important to point out that the relative intensity of the three light maps were adjusted so that the maximum luminance would be the same in all three images. It is also important to note, however, that the images all have different percentages of bright pixels. From left to right the PBP measures are 17%, 81%, and 99%, respectively.

These examples provide some anecdotal evidence that the higher frequency components of the illumination field are a possible factor that biases metallic surfaces to appear shiny black, shiny white, or metallic, but a closer examination reveals that is an oversimplification. Consider the three images in Figure 13 of a boy's bust made of polished chrome illuminated with different versions of the esplanade light map (see Figures 4 and 11). As in Figure 12, when the object is illuminated by the unfiltered light map (middle panel), it appears metallic (PBP = 79%), and when it is illuminated by



Figure 12. Images of a boy's bust made of polished chrome illuminated using the atrium light map (center), a high pass filtered version of that map (left), and a low pass filtered version (right).




Figure 13. Images of a boy's bust made of polished chrome illuminated using the esplanade light map (center), a high pass filtered version of that map (left), and a low pass filtered version (right).

the low pass version (right panel), it appears as shiny white (PBP = 99%). However, when it is illuminated by a high pass version of the esplanade (left panel), it appears arguably even more metallic (PBP = 85%) than the image produced with the unfiltered original.

The high pass versions of the atrium and the esplanade were created using exactly the same filtering process, but they have quite different effects on the perceptual appearance of the depicted chrome material. It is interesting to note in the left column of




Figure 14. Images of a red dielectric material (left) and a red metal material (right).

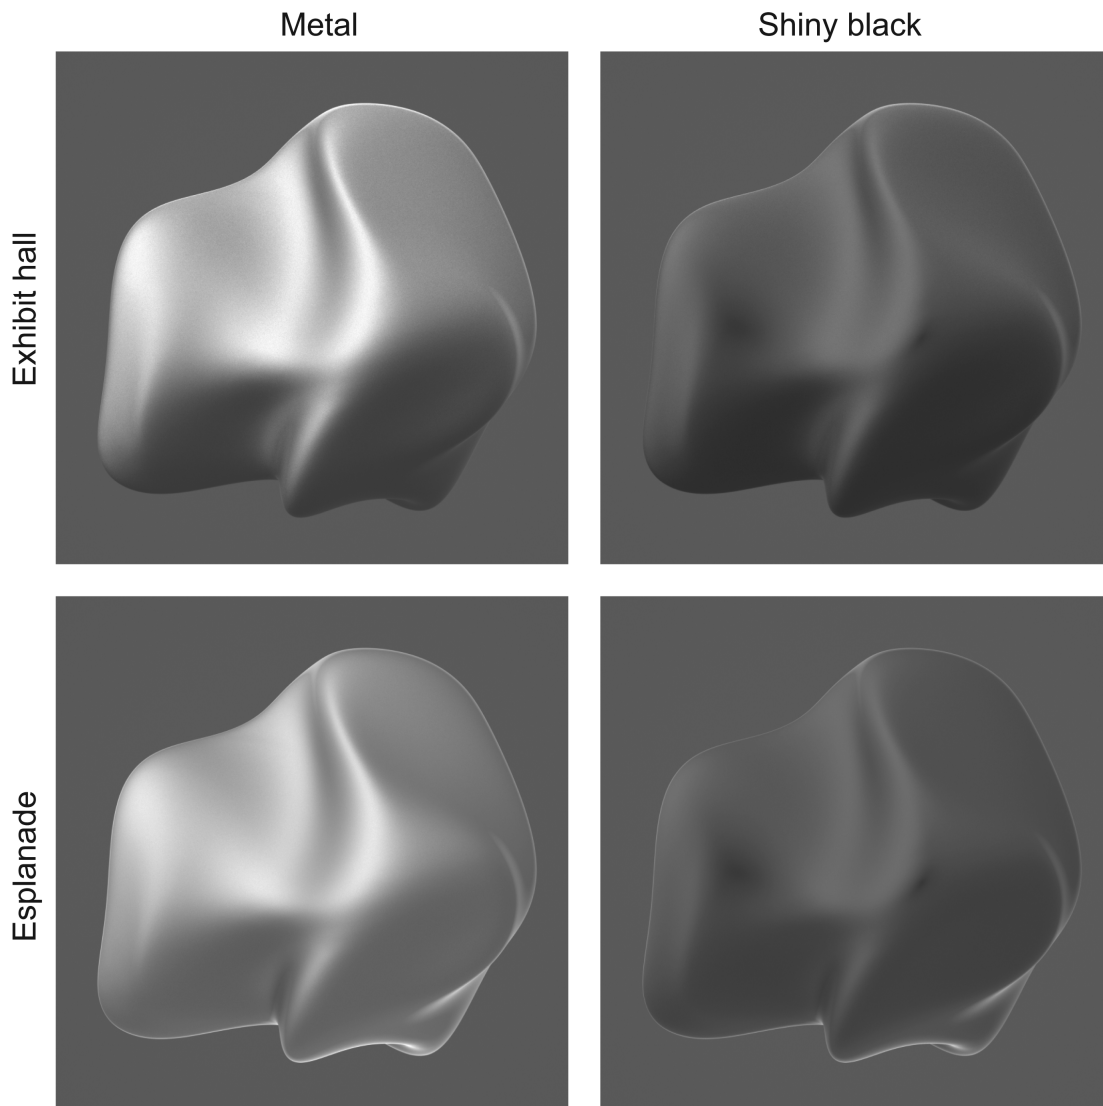


Figure 15. Images of a rough metal material (left) and a rough shiny black material (right) illuminated by the exhibit hall and esplanade light maps.

Figure 11, that the high frequency energy in the atrium light map is primarily localized in a limited arc of directions near the top, whereas the high frequency energy in the esplanade light map is more broadly distributed over a wider range of directions. If the distinctions between metal, shiny black, and shiny white are based on the PBP (and local contrast) then it is the distribution of illumination directions rather than spatial frequency per se that is the most important aspect of the illumination field for influencing the appearance of shiny materials.

It turns out that these distributional differences are also captured by the diffuseness2 metric. The value obtained for the high pass version of the atrium light map is 0.90, nearly double the value of 0.47 obtained for the unfiltered version. In contrast, the value obtained for the high pass version of the esplanade light map is only 0.48, which is slightly less than the value of

0.53 that was obtained for the unfiltered version. If high values of the diffuseness2 metric bias observers to perceive specular materials as shiny black rather than metal, this could explain why the high pass atrium light map produces images that appear shiny black (see left panel of Figure 12), whereas the high pass esplanade light map produces images that appear as metal (see left panel of Figure 13).

Effects of color, roughness, and background surfaces

Color

There are several other issues that deserve to be considered involving the generality of these findings. It is important to keep in mind that all of the stimuli in this



Figure 16. Three images of a distorted sphere illuminated by the exhibit hall (left), the atrium (middle), and the snowfield (right), with a visible background to provide more contextual information about each scene.

experiment were achromatic, but that is not typically the case for real objects in the natural environment. For example, consider the red dielectric material depicted in the left panel of Figure 14. This type of material has two distinct types of reflection: part of the illumination is transmitted a very short distance into the material, where some of it is scattered back toward the point of observation. This is referred to as diffuse reflection, and its color is determined primarily by the material. Another part of the illumination reflects directly off the object's surface. This is referred to as specular reflection, and its color is determined exclusively by the color of the illumination. The presence of these two different components in combination provides a potentially powerful source of information to indicate that one is looking at a shiny dielectric material. Note that shiny black materials constitute a special case, in which that information is unavailable, and the identity of a material as a dielectric must be determined by other means.

The right panel of Figure 14 shows a red metallic material with the same base color as the one in the left panel. For metals, 100% of the transmitted light is absorbed, so there are no diffuse reflections. However, unlike dielectrics, the color of the material can influence the colors of its specular reflections. That is why the color of the object in the right panel appears red, even though the illumination is completely desaturated. This likely provides useful information for the identification of colored metals, such as copper or gold, but there are other common metals, such as silver, aluminum, or chrome, whose colors are mostly desaturated. The identification of those materials as metal must depend on other factors, like the PBP and local contrast.

Roughness

In a previous experiment by Todd and Norman (2018) we examined how the appearance of metal and shininess of chrome objects is influenced by

variations of surface roughness. When the illumination had a broad range of directions, observers' metal and shininess ratings dropped rapidly with increasing roughness. For roughness values of 60 or higher, the material no longer appeared metallic or shiny, and was perceived instead as matte. Figure 15 shows images of a chrome material (left column) and a purely specular black dielectric material (right column), both with a roughness of 60. Note that the appearance of metal and shiny black have been completely eliminated. Both materials appear matte, although the metal one appears lighter than the black dielectric. The objects in the top row were illuminated using the exhibit hall light map, whereas the ones in the bottom row were illuminated by the esplanade light map. Note that there were significant differences in the metal and shiny black confidence ratings for these light maps when the materials were depicted with low roughness (see Figure 6), but that effect appears to be eliminated when the objects are depicted with high roughness.

Background surfaces

Adams et al. (2018) have recently demonstrated that constancy of gloss perception over variations of tone mapping is significantly improved if depicted objects are presented against a background of a natural scene, as opposed to a neutral gray background as in the present experiment. We were curious if contextual information about the lighting might also improve the categorization of metal and shiny black materials, so we created a set of images of a polished metal deformed sphere with the exhibit hall, atrium, and snowfield light maps, in which the same maps also provided a background scene. The resulting images are shown in Figure 16. To our eyes, the one on the left appears as shiny black, the one in the middle appears as metal, and the one on the right is ambiguous between metal and shiny white. These examples suggest that scene context may not noticeably improve the material constancy of

purely specular surfaces. However, it is possible that the visual structure of a surrounding scene might be more informative for material constancy if the observer is embedded in the scene, rather than viewing it in a picture. That is an issue that will remain to be addressed in future research.

Conclusions

This research has examined how the pattern and intensity of illumination can influence the perceptual categorization of shiny materials. The results reveal that metal and shiny black dielectric materials are easily confused. With broad distributions of illumination from multiple directions and higher illumination intensities, observers are biased to perceive shiny black surfaces as metal. Conversely, with narrow distributions of illumination from a small number of directions and lower illumination intensities, observers are biased to perceive metal surfaces as shiny black. Analyses of the image structure of the experimental stimuli reveal several possible measures that are highly correlated with observers' categorization judgments, and an analysis of the illumination fields reveals other measures that could potentially predict how much they bias the perceptual distinction between metal and shiny black materials.

Keywords: illumination, material constancy, surfaces

Acknowledgments

Supported by a grant from the National Science Foundation (BCS-1849418).

Commercial relationships: none.

Corresponding author: James T. Todd.

Email: todd.44@osu.edu.

Address: Department of Psychology, The Ohio State University, Columbus, OH, USA.

References

- Adams, W. J., Kucukoglu, G., Landy, M. S., & Mantiuk, R. K. (2018). Naturally glossy: Gloss perception, illumination statistics, and tone mapping. *Journal of Vision*, *18*(13):4, 1–16, doi:<https://doi.org/10.1167/18.13.4>.
- Doerschner, K., Boyaci, H., & Maloney, L. T. (2007). Testing limits on matte surface color perception in three-dimensional scenes with complex light fields. *Vision Research*, *47*, 3409–3423.
- Doerschner, K., Boyaci, H., & Maloney, L. T. (2010). Estimating the glossiness transfer function induced by illumination change and testing its transitivity. *Journal of Vision*, *10*(8), 1–9, doi:<https://doi.org/10.1167/10.4.8>.
- Fleming, R. W., & Bühlhoff, H. (2005). Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception*, *2*(3), 346–382.
- Marlow, P., & Anderson, B. L. (2013). Generative constraints on image cues for perceived gloss. *Journal of Vision*, *13*(14):2, 1–23, <https://doi.org/10.1167/13.14.2>.
- Marlow, P. J., & Anderson, B. L. (2015). Material properties derived from three-dimensional shape representations. *Vision Research*, *115*, 199–208.
- Marlow, P. J., Kim, J., & Anderson, B. L. (2017). The perception and misperception of surface opacity. *Proceedings of the National Academy of Sciences*, *114*(52), 13840–13845.
- Mooney, S. W. J., & Anderson, B. L. (2014). Specular image structure modulates the perception of three dimensional shape. *Current Biology*, *24*, 2737–2742.
- Mury, A. A., Pont, S. C., & Koenderink, J. J. (2007). Light field constancy within natural scenes. *Applied Optics*, *46*(29), 7308–7316.
- Nishida, S., & Shinya, M. (1998). Use of image-based information in judgments of surface-reflectance properties. *Journal of the Optical Society of America A*, *15*, 2951–2965.
- Olkkonen, M., & Brainard, D. H. (2010). Perceived glossiness and lightness under real-world illumination. *Journal of Vision*, *10*(9):5, 1–19, <https://doi.org/10.1167/10.9.5>.
- Olkkonen, M., & Brainard, D. H. (2011). Joint effects of illumination geometry and object shape in the perception of surface reflectance. *i-Perception*, *2*, 1014–1034.
- Pont, S. C., & te Pas, S. F. (2006). Material illumination ambiguities and the perception of solid objects. *Perception*, *35*, 1331–1350.
- Sharan, L., Rosenholtz, R., & Adelson, E. H. (2009). Material perception: What can you see in a brief glance? *Journal of Vision*, *9*(8):784, doi:<https://doi.org/10.1167/9.8.784>.
- Sharan, L., Rosenholtz, R., & Adelson, E. H. (2014). Accuracy and speed of material categorization in real-world images. *Journal of Vision*, *14*(9):12, 1–24, doi:<https://doi.org/10.1167/14.9.12>.
- Todd, J. T., & Norman, J. F. (2018). The visual perception of metal. *Journal of Vision*, *18*(3):9, 1–17. <https://doi.org/10.1167/18.3.9>.

- Wiebel, C., Valsecchi, M., & Gegenfurtner, K. R. (2013). The speed and accuracy of material recognition in natural images. *Attention, Perception, & Psychophysics*, 75(5), 954–966.
- Xia, L., Pont, S. C., & Heynderickx, I. (2017). Light diffuseness metric part 1: Theory. *Lighting Research & Technology*, 49(4), 411–427.
- Xiao, B., Walter, B., Gkioulekas, I., Zickler, T., Adelson, E., & Bala, K. (2014). Looking against the light: how perception of translucency depends on lighting direction. *Journal of Vision*, 14(3):17, 1–22, doi: 10.1167/14.3.17.
- Zhang, F., de Ridder, H., & Pont, S. C. (2015). The influence of lighting on visual perception of material qualities. *Proceedings of SPIE 9394, Human Vision and Electronic Imaging XX, 93940Q*, <https://doi.org/10.1117/12.2085021>.
- Zhang, F., de Ridder, H., Barla, P., & Pont, S. (2019). A systematic approach to testing and predicting light-material interactions. *Journal of Vision*, 19(4):11, 1–22, <https://doi.org/10.1167/19.4.11>.

Appendix

The parameterization of surface reflectance

The manner in which a material interacts with light is determined by its IOR, which is a complex number ($n - ki$). The imaginary coefficient (k) is also known as the extinction coefficient. For most of the materials encountered in nature (i.e., dielectrics), the value of k is vanishingly small, but that is not the case for metals. To better appreciate how n and k can influence reflections on metal surfaces it is useful to consider the three images shown in Figure A1. The left panel of this figure depicts an aluminum material with an IOR of (1.2, 7.0), which is perceived as a shiny metal. The image in the middle panel was generated with exactly the same parameters as the one on the left, except that the value

of k was set to zero. This changes the appearance to a black dielectric material. However, by raising the value of n to 50, and keeping k at zero, the appearance of a shiny metal can be restored, as shown in the right panel. The pattern of shading in that case is similar (but not identical) to the image in the left panel.

The complex IOR is used by the Fresnel equations to compute the surface reflectance at all possible incident angles and directions of polarization. Some renderers, such as Maxwell, Renderman, or Arnold, allow the user to input values of n and k to characterize a material. This is the best way to achieve physical accuracy because the appropriate values can be obtained from a material handbook or at the web site <https://refractiveindex.info/>. Other renderers employ a different approach. By specifying the reflectance and color at incident angles of 0° and 90° (also referred to as facing and grazing angles), it is possible to compute the values of n and k from that information. This is just a reparameterization of the same space, which some users find more intuitive because it focuses on the end result of the computation.

Still other renderers only allow the user to specify the real coefficient (n) of the IOR, or a facing angle reflectance, and the value of k is assumed to be zero. This approach allows the program to use an approximation of the Fresnel equations, which can greatly reduce the required computations. Metals can only be simulated in these renderers by using an unnaturally high value of n , as shown in the right panel of Figure A1. Another variation of this approach is to incorporate a metalness parameter. Although the term metalness does not occur in physics, it is used in renderers to make the user interface more intuitive. The metalness parameter acts as a switch between metal and dielectric materials. When set to zero, a low value of n is used that is typical of dielectric materials, and the color of the specular reflections is exclusively determined by the color of the illumination. When set to one, a high value of n is used, and the material color is allowed to influence the specular reflections. The IOR parameter in these renderers may only be used for refractions.




Figure A1. Three images of a distorted sphere with different complex IORs. From left to right the IORs are (1.2, 7.0), (1.2, 0), and (50, 0).