# Contours produced by internal specular interreflections provide visual information for the perception of glass materials

**James T. Todd**

Ohio State University ✉

**J. Farley Norman**

Western Kentucky University ✉

**Two experiments are reported that investigated how the perceptual identification of glass is influenced by banding contours formed by internal specular interreflections within glass materials. Observers made material categorization judgments for images depicting glass, chrome, shiny black and shiny white objects, and for contour drawings that were created by edge filtering images of glass, chrome or textured objects. Observers rated each stimulus by adjusting four sliders to indicate their confidence that the depicted material was glass, metal, shiny black, or something else, and these adjustments were constrained so that the sum of all four settings was always 100%. The results revealed that the rendered images were all categorized correctly with a high level of confidence. The contour drawings of glass and textured materials were also categorized correctly with a high level of confidence. However, the contour drawings of chrome materials were miscategorized as glass, with an average confidence rating that was significantly lower than those obtained for the glass contours. It is hypothesized that these different contour types are perceptually distinguished from one another based on how they align with the pattern of surface curvature on an object and the smoothness of the contours.**

## Introduction

The interreflection of light on surfaces is a largely neglected topic in the study of visual perception. Although there have been a few computational analyses (Koenderink & van Doorn, 1983; Nayar, Ikeuchi & Kanade, 1991; Langer, 1999) and some early psychophysical investigations (Gilchrist & Jacobsen, 1984; Bloj, Kersten & Hurlbert, 1999; Madison, Thompson, Kersten, Shirley, & Smits, 2001; Doerschner, Boyaci, & Maloney, 2004), the effects of indirect illumination are still widely misunderstood (Todd, Egan & Kallie, 2015). Moreover, the limited research that has been performed on this topic has focused almost exclusively on diffuse reflections on surfaces that scatter light uniformly in all directions.

The first researchers to consider specular interreflections on shiny surfaces were Pont & Koenderink (2002, 2005). In an effort to derive a theoretical bidirectional reflectance distribution function for rough pitted materials, they performed a detailed analysis of how light behaves inside a shiny concave hemispherical surface. Consider the six images shown in Figure 1, which show a single hemispherical pit composed of a polished silver material with varying numbers of surface interreflections. The image in the top left panel shows the visible structure that is produced by a single bounce of direct illumination. Note that the image of the surrounding scene is inverted, and that it is only visible in the central portion of the surface where light can be directly reflected toward the point of observation. The image in the top middle panel shows the visible structure that emerges after one additional indirect bounce. Note that this structure contains two circular bands: A large inner one where the surrounding scene appears upright, and a smaller outer one in which the scene appears inverted. As is shown in the remaining panels, each additional bounce produces two additional bands of visible structure that get progressively smaller as they approach the outer edge of the surface.

Although this banding behavior may appear at first blush to be counterintuitive, it can easily be explained using geometrical optics. It is important to keep in mind that the simulated surface in Figure 1 is perfectly smooth. Thus, at any given point on the surface, there is only a single direction of illumination that will reflect light toward the point of observation, and that direction can be determined using backward ray tracing from the eye. Figure 2 shows three such points labeled A, B and C on the lower half of a concave hemisphere, and the optical paths (colored black, green, and red) that would allow those points to reflect light toward the point of observation. Point A is in the region where direct reflections of the environment are visible, as shown by
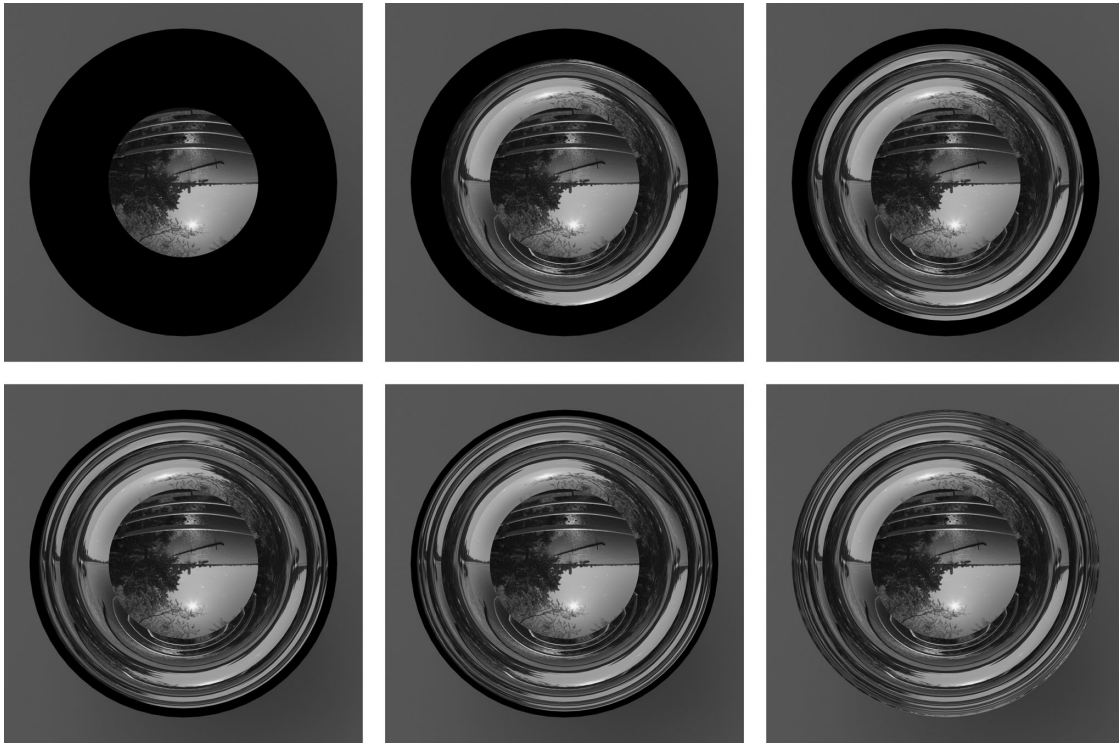
Figure 1. A concave hemispherical pit rendered with different numbers of reflective bounces. The top row from left to right shows 1, 2 and 3 bounces, respectively. The bottom row shows 4, 5, and 100 bounces. Note that if this were a real photograph, the camera would be visible in the reflections.
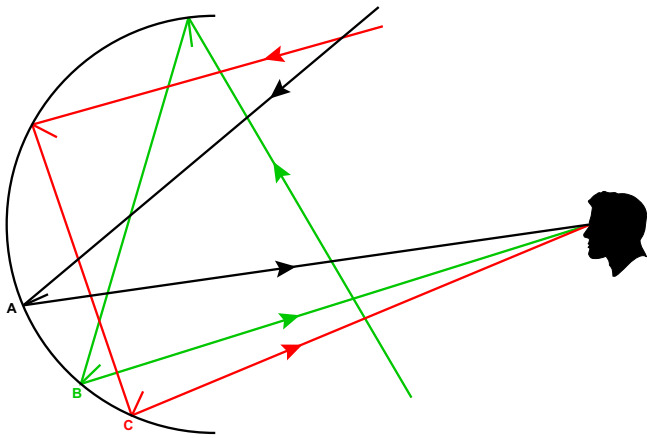


Figure 2. The light paths that reach the point of observation from three points on a concave hemispherical surface.

the path colored black. Note that the light that reaches the eye from that point comes from the upper part of the environment so that the image near A is inverted. Point B is located further out in the periphery. It can only reflect light toward the eye after an additional indirect bounce, as shown by the path colored green. Note how that path originates in the lower part of the environment so that the image near B is upright. Finally, point C is located even farther out in the periphery, and it also requires an additional indirect bounce to reflect

light toward the eye as shown by the path colored red. However, note in that case that the path originates in the upper part of the environment, so that the image near C is again inverted. By moving farther and farther into the periphery, the number of required bounces to reflect light toward the eye increases, and the alternating bands of upright and inverted images repeat with higher and higher frequency.

To summarize briefly, the interreflective bands within a concave hemisphere are organized in a hierarchical manner based on simple geometrical optics. At one level of structure, each band is defined by the number of reflective bounces required for light from the surrounding environment to reach the point of observation. At a more subordinate level, each band is subdivided into two concentric regions: One where the visible image is upright, and another where the visible image is inverted.

There are a couple of other factors that can influence the extent of this banding. One is the relative depth of a pit. This is demonstrated in Figure 3, which shows spherical pits of a silver material with three different depths. As the pit becomes more and more compressed (moving from left to right), the central region with visible direct reflections expands, and the banded region is pushed farther into the periphery. Another important factor is the specular reflectance of the surface material. Figure 4 shows hemispherical pits

Figure 3. A concave ellipsoidal pit with different depth to width ratios. From left to right, the depth relative to the horizontal radius is 1, .75 and .5, respectively.
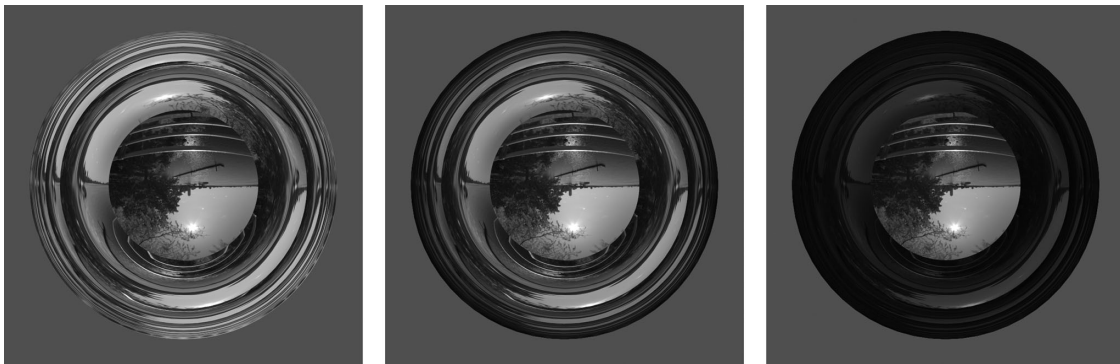


Figure 4. A concave hemispherical pit composed of different materials. From left to right, the depicted materials are silver, chrome and obsidian (i.e., volcanic glass).

composed of silver, chrome and obsidian (i.e., volcanic glass). Silver reflects almost 100% of the illumination at all incident angles; chrome reflects approximately 50% of the illumination; and obsidian only reflects about 5% at most incident angles. The illumination intensity in these images has been adjusted so that their central regions with visible direct reflections all have the same luminance. However, as the surface reflectance is systematically reduced, the visible banded regions do not extend as far into the periphery of the surface, and they have reduced contrast. This is because a greater proportion of energy is absorbed on each bounce.

Although the visible bands produced by specular interreflections in concave surface regions were first discovered and analyzed by Pont & Koenderink (2002, 2005), a similar observation was later described by Todd and Norman (2019) for glass materials, but we did not initially make the connection with Pont and

Koenderink's research. Glass is an interesting material in this regard. When light hits a boundary from air to glass, almost all of its energy is transmitted except at high incident angles. Conversely, when light hits a boundary from glass to air at any incident angle above 41°, 100% of its energy will be reflected—a phenomenon that is referred to as total internal reflection. As a result of this high level of internal reflectance, glass materials can exhibit the same type of visible banding structure as is demonstrated in Figures 1 to 4 for shiny metals. It is important to keep in mind, however, that this behavior occurs internally within an object such that all signs of curvature on its surface are reversed relative to how they appear from the outside. In other words, what appear to be ridges and bumps from the outside are actually valleys and pits from the inside.

The left column of Figure 5 shows two glass objects with relatively complex patterns of concavity and
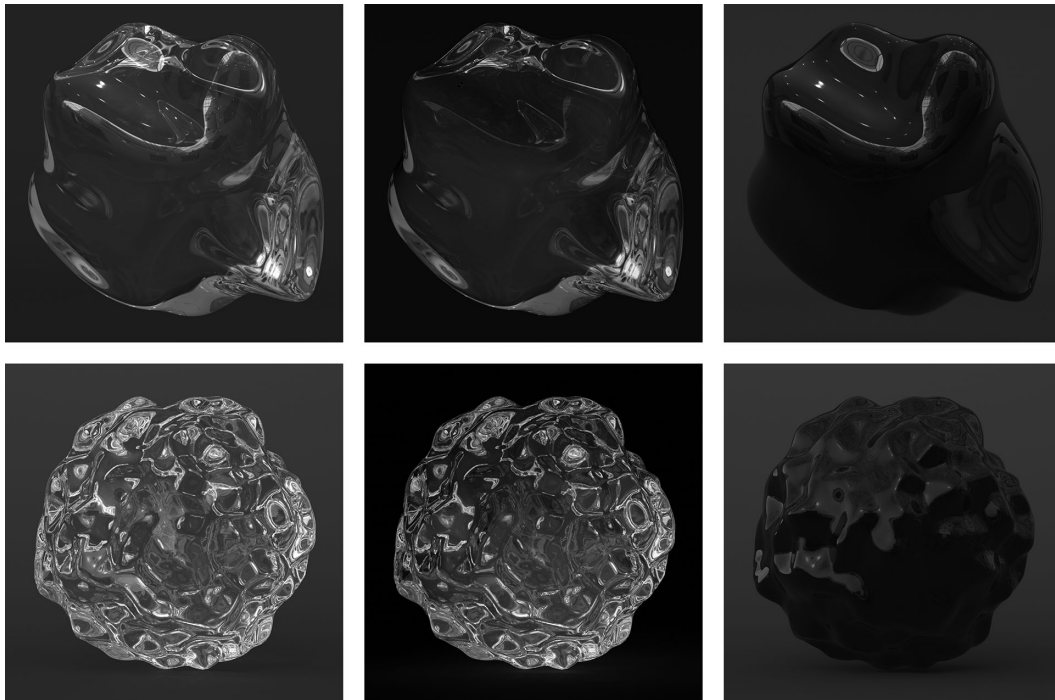
Figure 5. Two glass objects with different rendering constraints. The left column shows the combined effects of both reflected and transmitted light. The middle column shows only the transmitted light and the right column shows only the reflected light.

convexity. The middle column shows only the effects of light that is transmitted inside the objects, and the right column shows only the effects of light that is reflected off their outside surfaces. Note that the images of transmitted light are perceived as a glass material, whereas the images of the outer surface reflections are perceived as a shiny black material like obsidian (see Todd & Norman, 2019). Let us first consider the object in the top row that has a deep internal concavity on its lower right. Note that this region has a complex pattern of light and dark bands similar to the ones shown in the left panel of Figure 3. There are also two shallower internal pits at the top of this object, whose outer edges are marked by thin bands similar to the ones in the right panel of Figure 3. The object in the bottom row contains a dense pattern of internal pits, and many of those are also lined with a pattern of light and dark bands. By manipulating the number of possible indirect bounces (as in Figure 1), Todd and Norman (2019) showed that these structures only emerge over multiple internal reflections.

Is it possible that these banded contours provide useful information for the perceptual identification of glass materials? To address that issue, Todd and Norman (2019) created a few displays in which contours were presented in isolation without any other variations in gray scale. This was achieved by applying a Sobel edge filter to images of glass objects using the Photoshop CC 2019 find edges tool. An example of this is shown in Figure 6. The original image is shown

in the left panel, and the transformed contour pattern is shown in the right panel. Observers categorized both of these objects as glass with 100% confidence. It is important to keep in mind that edge filtered images do not completely isolate internal banding contours from surface interreflections. They also produce edges from occlusion boundaries and specular highlights on the external boundary. However, those latter features are the same or similar for all types of shiny surfaces and are not therefore diagnostic for the perceptual identification of glass.

Todd and Norman (2019) noticed that the banding contours in glass materials can be quite complex. They often create swirling patterns that resemble the turbulent flow of fluid materials. Note in the right panel of Figure 6 that these swirling patterns are especially evident in the base of the object where the trunk is attached to the leg, but there are other less prominent examples in the middle of the back, the top of the ear, and the front of the face. Todd and Norman described these features as flow eddies, and they speculated that the swirling patterns of contours could provide a useful source of information for the perceptual identification of glass materials.

The research described in the present article was designed to explore this finding in much greater detail than the simple demonstration of Todd and Norman (2019). Its goals were twofold: First, to determine whether observers can reliably distinguish the contours on glass objects from those that occur on other types

Figure 6. A glass image of an elephant, and a contour drawing of that image created using an edge filter.

of materials; and second, to determine whether objects defined by those contours are reliably identified as glass.

# Experiment 1

## Methods

The methods and procedure were similar to those used by Todd and Norman (2019) and Norman, Todd, and Phillips (2020). Observers made material categorization judgments for 72 possible stimuli. Half of the stimuli depicted images of metal, glass, shiny black and shiny white objects. These were included to provide a baseline of performance for evaluating observers' judgments of the remaining stimuli that depicted edge filtered contour drawings of the metal and glass materials.

### Stimuli

There are several practical problems that needed to be addressed in creating the stimuli for this experiment. Although patterns of interreflective contours were the primary focus of this research, it was important that those contours be generated from images of metal and glass materials that are clearly recognizable. The first problem we faced is that the appearance of these materials is strongly affected by the pattern of illumination (Todd & Norman, 2018; Todd & Norman, 2019; Norman et al., 2020). For, example, if the illumination is too sparsely distributed, metal surfaces can appear as shiny black, and, if the illumination is too diffuse, metal surfaces can appear as shiny or matte white. The appearance of glass is also particularly sensitive to the way it is illuminated (Hunter, Biver, & Fuqua, 2007; Todd & Norman, 2019). With nonoptimal lighting, the image of a glass object may contain weird looking patches of black and white, or the object may

disappear altogether as if it were invisible. This problem is compounded by the fact that the ideal illumination for one material may not work well for others (Zhang, de Ridder, & Pont, 2015). For example, Figure 7 shows three images of a bumpy sphere composed of shiny white, metal and glass materials. The lighting in these scenes includes two large area lights, one that illuminates the surface from the front left, and another that illuminates it from the back right. This is a common pattern of lighting in photography for opaque dielectric materials because it provides good contrast and clear definition of an object's boundaries. However, this does not work at all for metal or glass. Note in Figure 7 that the metal object in the middle panel appears as a shiny black material, and the glass object in the right panel appears as a metal or shiny black material.

Hunter, Biver, and Fuqua (2007) have observed that glass objects have "grayed the hair and wasted the time of more photographers than any other substance." This is because most patterns of illumination do not produce compelling images of glass materials, and that is also true for metals. If you pick a light map at random to illuminate a scene, the odds of it producing recognizable glass or metal materials are rather small. Our solution to this problem was to use a light map of the Charles River esplanade from the sIBL archive that has been shown in previous research to allow accurate categorizations of all the materials used in the present experiment (Todd & Norman, 2019; Norman et al., 2020).

Another important issue that needed to be addressed concerns the multiple scale spatial structure of internal and external surface interreflections. Note in Figure 1 how the spatial frequency of the surface reflections becomes greater and greater in more peripheral regions of the hemisphere, and a similar effect is also observed within the glass images of Figure 5. For even modestly complex surface geometries, these high frequency variations in luminance may not be resolved
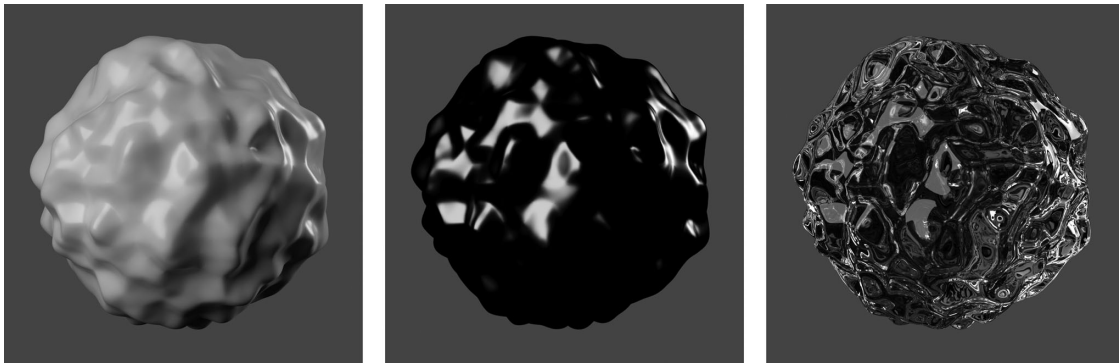
Figure 7. A bumpy sphere illuminated by two large area lights. The depicted materials include shiny white (left), metal (middle), and glass (right).
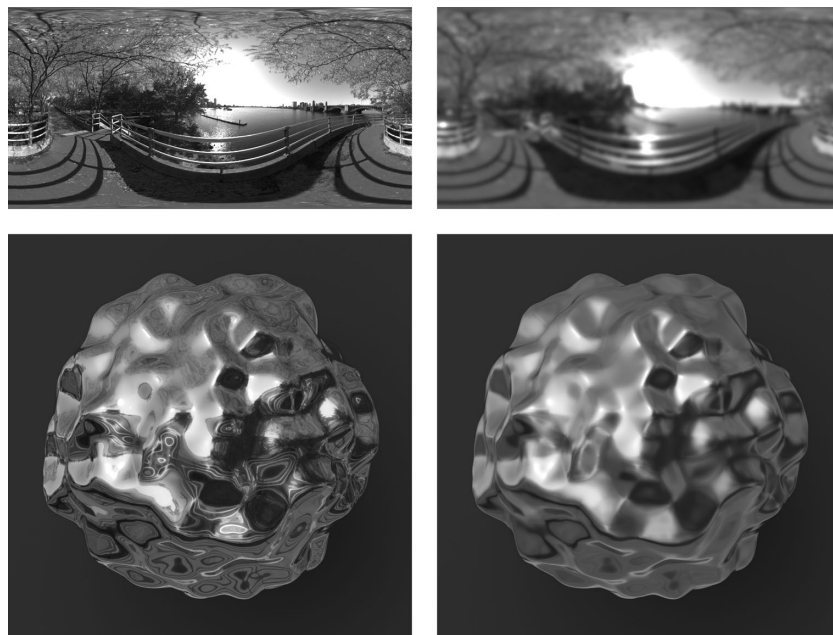


Figure 8. The top row contains blurred (right) and unblurred (left) light maps of the Charles River esplanade. The bottom row shows images of a bumpy sphere rendered with these light maps.

adequately by the renderer, and this can produce a noisy appearance in the resulting image. This can be mitigated to some extent by rendering an image at a very high spatial resolution, but that dramatically increases rendering times for glass materials. The alternative is to somehow blur the high frequency structure. This can be achieved in several possible ways. One is to blur (i.e., antialias) the final image, which can create a blurry appearance. Another is to add a small amount of roughness to the surface material. This works well with shiny opaque materials (Mooney & Anderson, 2014; Todd & Norman, 2018), but it radically changes the appearance of glass (Todd & Norman, 2019). A third alternative is to blur the light map, so that its high frequency components are

smoothed. This is functionally equivalent to adding roughness to opaque surfaces, but it also works well with glass (see Todd & Norman, 2019). For example, the top left panel of Figure 8 shows a light map of the Charles River esplanade. The top right panel shows a blurred version of that map. The original esplanade light map had a spatial resolution of 3200 × 1600 pixels, and we transformed it using a Gaussian blur filter with a radius of 15 pixels. The images in the lower two panels of Figure 8 show bumpy chrome spheres illuminated by the light map directly above them. To our eyes, the blurred one on the right looks much more natural than the unblurred one on the left, and that is the one we employed for all of the stimuli in the present experiment. Although it is possible to produce
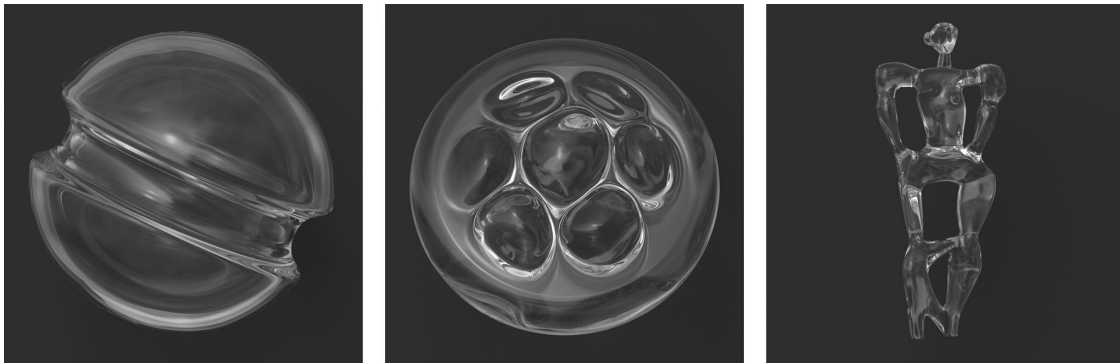
Figure 9. Three images of glass objects used in the present experiments.

polished chrome materials that produce perfect mirror reflections, they typically become smudged by dust or water deposits when exposed to the elements of a natural environment.

The rendered images used in this experiment were created using Maxwell Renderer 4 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. Although the quality of the images it produces is quite high, this comes at a substantial cost in rendering time, especially for materials that involve transparency or translucency. The images were rendered on a computer cluster with 64 cores.

There were nine different stimulus objects used in the experiment that were all presented in front of a dark gray background plane. All of the objects were approximately 10 cm in height, and their widths varied from 5 to 10 cm. Each object was rendered using a simulated camera at a distance of 55 cm with a 171 mm lens that had a 12° field of view, and an F-stop of 40 so that there was a large depth of field. The simulated materials included a glass material with a complex IOR of (1.5, 0), a chrome material with a complex IOR of (3.2, 3.3), a shiny black material (i.e., obsidian), whose reflections were identical to glass, but without any light transmission, and a shiny white material with a linear combination of diffuse and specular components. Figure 9 shows three of the objects with a glass material, and Figure 10 shows the remaining six objects with metal, shiny black, and shiny white materials. The experiment included all 36 possible combinations of nine objects with four materials.

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 gamma 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 map was 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.

An additional 36 stimuli were created by applying a Sobel edge filter to all of the images of glass and chrome materials using the Photoshop CC 2019 find edges tool. Because we suspected that the contrast polarity might have a significant effect on observers' perceptions, we created two versions of these contour images, in which the contours could be presented as either white on black or black on white. Figure 11 shows white on black contour images for three of the objects in both the glass and metal conditions. Figure 12 shows black on white contour images for three other objects. Note that the white on black images of glass contours appear vaguely similar to the images produced using the dark field method that is popular in photography, in which objects are illuminated from the top and sides against a black background (see Todd & Norman 2019). Similarly, the black on white contours appear vaguely similar to images produced using the bright field method, in which objects are illuminated with diffuse light from behind.

### 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) using an Apple 27-inch LED Cinema Display (2560 × 1440 pixel resolution). The monitor was located at a 60 cm viewing distance. The luminous 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

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
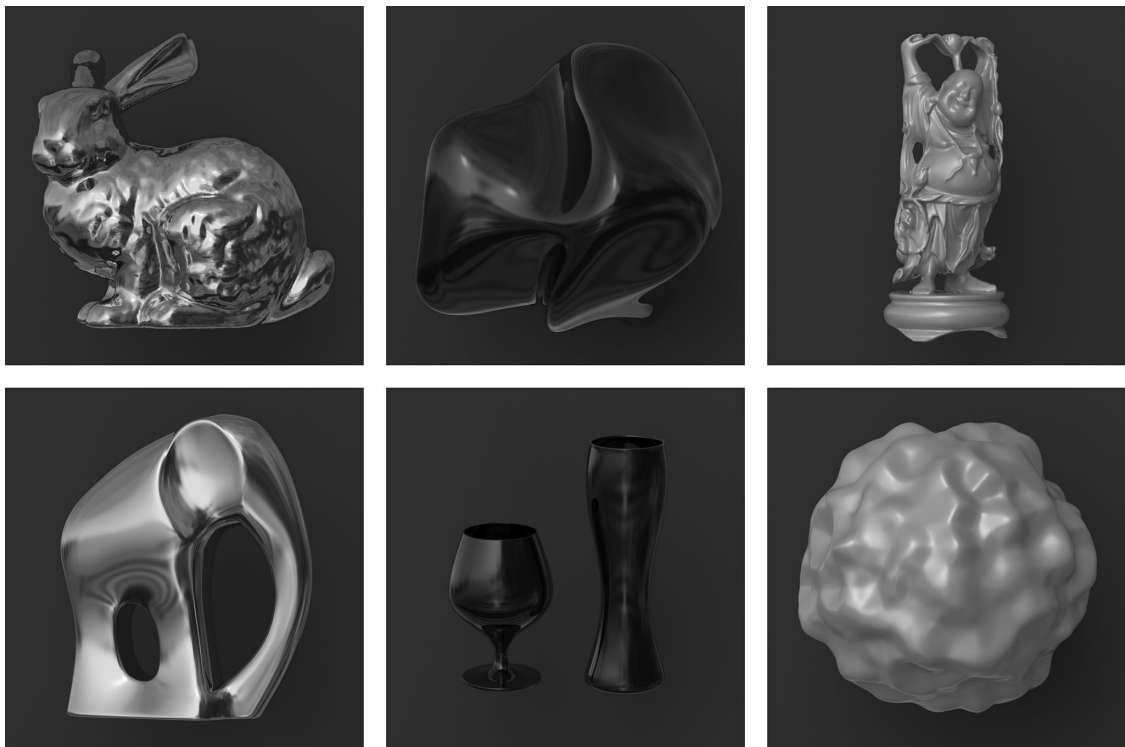
Figure 10. Six of the objects used in the present experiments composed of metal, shiny black and shiny white materials.
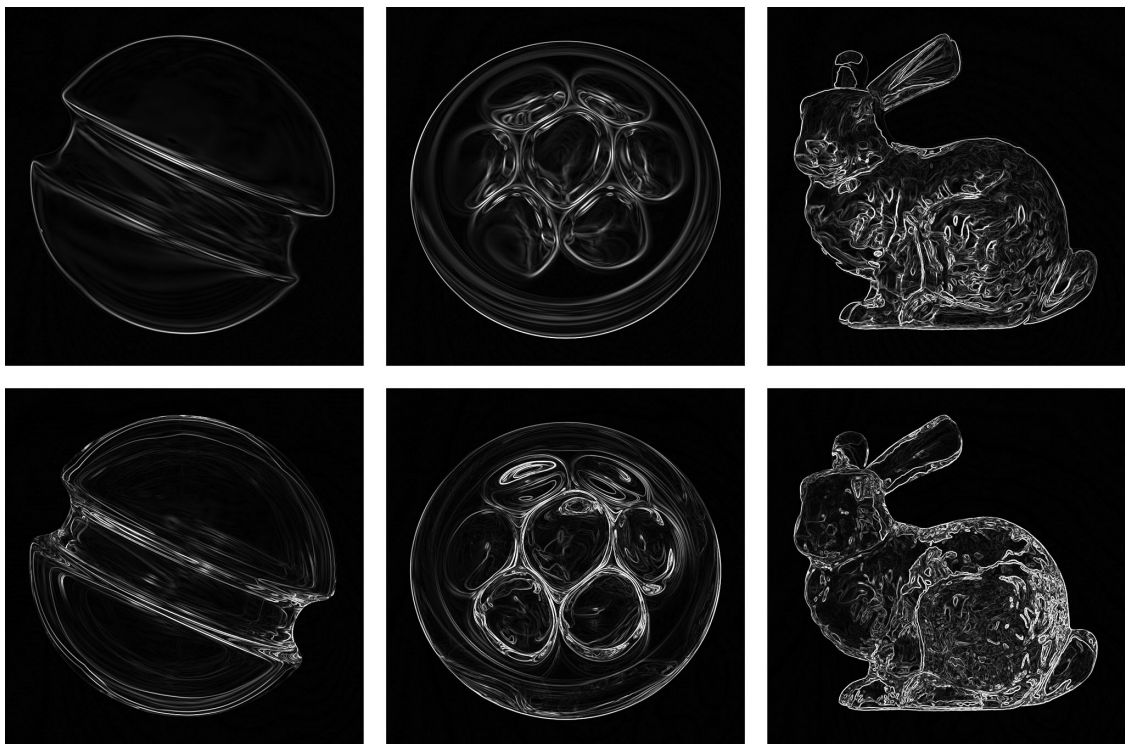


Figure 11. White on black contour drawings of three objects used in Experiment 1. The top row depicts chrome and the bottom row depicts glass.
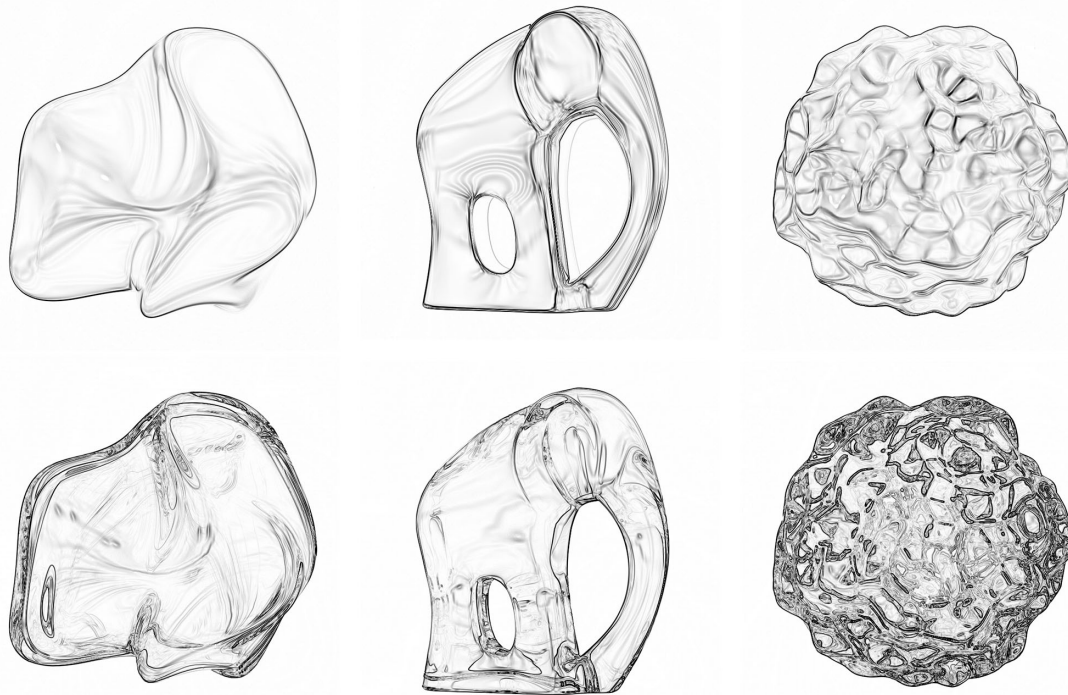
Figure 12. Black on white contour drawings of three objects used in Experiment 1. The top row depicts chrome and the bottom row depicts glass.

mouse. Each of the sliders represented a different category labeled glass, metal, shiny black 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, and that these confidence ratings should always sum to 100%. Any set of responses that did not satisfy this criterion would prevent the program from advancing to the next trial. Because of the digital readout of their settings, observers had no difficulty conforming to this instruction. We knew from our earlier research that images of glass can occasionally be misinterpreted as a metal or shiny black material (see Todd & Norman, 2019). That is why we incorporated those categories in the response settings. However, we did not want the observers to feel forced to only consider glass, metal or shiny black when evaluating each stimulus, so we also included a "something else" response option, and also added the shiny white stimuli in an effort to force them to give that category a high rating on a subset of the trials.

### Observers

The displays were judged by one of the authors (JFN) and nine other observers who were completely naïve 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 72 stimuli, and all observers participated in two sessions. At the beginning of each session, the details of the response task were explained, and observers were shown real physical examples of glass, metal and shiny black materials. However, they were also instructed that other types of materials would be presented as well, and that those should be categorized as something else.

### Results

Table 1 shows the average confidence ratings for all of the different response categories in all of the different stimulus conditions. Note that the observers' judgments of the original rendered images were almost perfectly accurate: The chrome materials were categorized as metal with an average confidence rating of 90%; the shiny black materials were categorized as shiny black with an average confidence rating of 94%; the glass materials were categorized as glass with an average confidence rating of 96%; and the shiny white materials were categorized as something else with an average confidence rating of 94%.

| Object | Shaded glass | | | | Shaded metal | | | | Shaded black | | | | Shaded white | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Response category | | | | Response category | | | | Response category | | | | Response category | | | |
| | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other |
| Yoyo | 85 | 2 | 0 | 13 | 1 | 85 | 0 | 14 | 0 | 1 | 96 | 3 | 0 | 4 | 0 | 96 |
| Pepper | 92 | 0 | 0 | 8 | 1 | 87 | 0 | 12 | 0 | 1 | 98 | 0 | 0 | 10 | 0 | 90 |
| Elephant | 98 | 0 | 0 | 2 | 2 | 95 | 0 | 4 | 0 | 3 | 96 | 1 | 1 | 3 | 0 | 96 |
| Cobble | 94 | 0 | 1 | 5 | 0 | 93 | 0 | 7 | 2 | 7 | 86 | 5 | 4 | 10 | 0 | 87 |
| Glasses | 100 | 0 | 0 | 1 | 0 | 100 | 0 | 1 | 0 | 4 | 96 | 0 | 1 | 6 | 0 | 93 |
| Bumpy | 97 | 0 | 0 | 3 | 2 | 85 | 0 | 13 | 1 | 9 | 90 | 0 | 1 | 8 | 0 | 91 |
| Bunny | 99 | 0 | 0 | 2 | 9 | 84 | 2 | 6 | 0 | 1 | 97 | 2 | 0 | 4 | 0 | 96 |
| Woman | 99 | 0 | 0 | 2 | 5 | 90 | 0 | 5 | 1 | 4 | 95 | 2 | 0 | 4 | 0 | 96 |
| Buddha | 99 | 0 | 0 | 2 | 4 | 92 | 0 | 4 | 0 | 3 | 96 | 1 | 0 | 3 | 0 | 98 |
| **Average** | **96** | **0** | **0** | **4** | **3** | **90** | **0** | **7** | **0** | **4** | **94** | **1** | **1** | **6** | **0** | **94** |

| Object | Glass contours black background | | | | Glass contours white background | | | | Metal contours black background | | | | Metal contours white background | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other |
| Yoyo | 83 | 0 | 9 | 8 | 94 | 3 | 0 | 4 | 9 | 2 | 83 | 6 | 71 | 11 | 0 | 18 |
| Pepper | 89 | 0 | 10 | 2 | 97 | 2 | 0 | 1 | 8 | 0 | 89 | 3 | 71 | 10 | 0 | 19 |
| Elephant | 94 | 0 | 6 | 1 | 94 | 5 | 0 | 1 | 46 | 2 | 48 | 4 | 77 | 21 | 0 | 2 |
| Cobble | 79 | 1 | 6 | 14 | 85 | 6 | 0 | 9 | 31 | 0 | 66 | 3 | 72 | 10 | 0 | 18 |
| Glasses | 95 | 0 | 0 | 5 | 89 | 0 | 5 | 6 | 71 | 4 | 24 | 1 | 75 | 24 | 0 | 1 |
| Bumpy | 93 | 1 | 4 | 2 | 75 | 16 | 0 | 10 | 35 | 2 | 55 | 8 | 79 | 15 | 0 | 7 |
| Bunny | 86 | 4 | 7 | 3 | 81 | 7 | 0 | 12 | 85 | 1 | 8 | 6 | 83 | 7 | 0 | 10 |
| Woman | 95 | 0 | 2 | 3 | 88 | 7 | 1 | 4 | 84 | 4 | 10 | 2 | 80 | 17 | 0 | 2 |
| Buddha | 90 | 2 | 4 | 4 | 55 | 18 | 3 | 25 | 81 | 7 | 7 | 5 | 69 | 20 | 2 | 9 |
| **Average** | **89** | **1** | **5** | **5** | **84** | **7** | **1** | **8** | **50** | **2** | **43** | **4** | **75** | **15** | **0** | **10** |

Table 1. The average confidence ratings for each stimulus and response category of Experiment 1. All of the cells in this table have standard errors that are less than 3%. The rows of each stimulus category do not all sum to 100 due to rounding error.

Of course, the main focus of the experiment concerned the contour patterns. It is important to keep in mind that observers were given no instructions about those. In the absence of other evidence, it might be reasonable to expect that the contour drawings might not appear to have any recognizable material at all. We would expect in that case that the observers would categorize those stimuli as something else with a high degree of confidence, just as they rated the shiny white materials that did not have a designated response category—but that is not what occurred. Indeed, the average "something else" confidence rating for all of the contour drawings was only 7%.

One possible factor that could influence observers' judgments of the contour stimuli is edge density. Because the contour patterns were mostly binary, an approximate measure of edge density can be obtained by computing the mean luminance for each black on white contour pattern, excluding the background. On average, the glass contours had 5% higher edge density than the chrome contours, but that difference was dwarfed by the variations between the different objects, whose edge densities ranged from 9% to 39%. The objects listed in Table 1 are ordered with respect to their average edge density for the glass and chrome conditions.

Let us first consider the ratings obtained for the contour patterns that were produced by edge filtering the images of glass objects. The most highly rated response category was glass in both the white on black and black on white conditions, with average confidence ratings of 89% and 84%, respectively, and the second most highly rated category was something else, with average confidence ratings of 5% and 8%.

For the contour patterns produced by edge filtering the images of chrome objects, there was a noticeable effect of contrast polarity, which did not occur for the glass contours. For the stimuli with white on black contours, the most highly rated response category was glass, but the average confidence rating in that case was only 50%. The second most highly rated category was shiny black, with an average confidence rating of 43%. It is also interesting to note that the glass ratings for the white on black chrome contours were highly correlated with edge density (r = .80). The low density objects were much more likely to be rated as shiny black than the high density ones. For the stimuli with black on white contours, the most highly rated response category was also glass, with an average confidence rating of 75%, and the second most highly rated category was metal, with an average confidence rating of 15%. For the black on white contours there was a negligible correlation between the glass ratings and edge density (r = 0.08).

A statistical analysis of these data is complicated by the fact that the variance of observers' ratings is far from uniform across all of the different conditions. When the average confidence rating is near 0% or 100% the variance is negligible, but it becomes much larger when the average rating is in the middle of that range. In order to spread out the judgments at the tails of the distribution, we performed a logit transformation on the data. To prevent the transform from going to infinity, any rating of 100% was converted to 99.5% and any rating of 0% was converted to 0.5%.

The transformed glass confidence ratings were then analyzed using a 2 (materials) by 2 (contrast polarity) analysis of variance. The results revealed that there was a significant effect of the depicted material ($F(1, 9) = 17.67$, $p < 0.002$), such that the contours from glass materials received higher confidence ratings than contours from chrome materials. There was no significant main effect of contour polarity ($F(1, 9) = 3.01$, $p > 0.1$), but there was a significant interaction ($F(1, 9) = 22.09$, $p < .001$). We also performed $t$-tests to compare observers' ratings of the glass objects to those obtained in the contour conditions, collapsed over variations in contrast polarity. The average rating of the glass images was 9% higher than the ratings obtained for the glass contour conditions, and this effect was marginally significant ($t(9) = 2.84$, $p < .05$). However, the average rating of the glass images was 33% higher than the ratings obtained for the chrome contour conditions, and this effect had a much higher level of significance ($t(9) = 5.63$, $p < .0005$).

These findings indicate that much of the visual information in our glass images is also present in contour drawings that are generated from those images, and that observers can identify those drawings as glass with only slightly less confidence than they exhibit for fully rendered images. The results also demonstrate that observers could reliably distinguish the contours obtained from images of glass and chrome—at least to some extent. It is important to keep in mind, however, that all of the chrome contour patterns were categorized primarily as glass. Observers' judgments of those patterns only differed from their judgments of glass in terms of their reported level of confidence.

So what is the information on which these confidence ratings are based? In order to address this question it is useful to pay careful attention to Figures 11 and 12. One similarity between glass and metal contours is that they tend to cluster in regions of high slant or curvature. The primary difference is that the glass contours have a more complex chaotic structure. Todd and Norman (2019) referred to these patterns as flow eddies because they resemble the turbulent flow of fluid materials. We suspect that both of these factors are important sources of information for observers' judgments. The clustering of the contours with respect to the surface geometry is what makes the glass and chrome contours appear similar, and the chaotic structure of the glass contours is what makes them appear different.
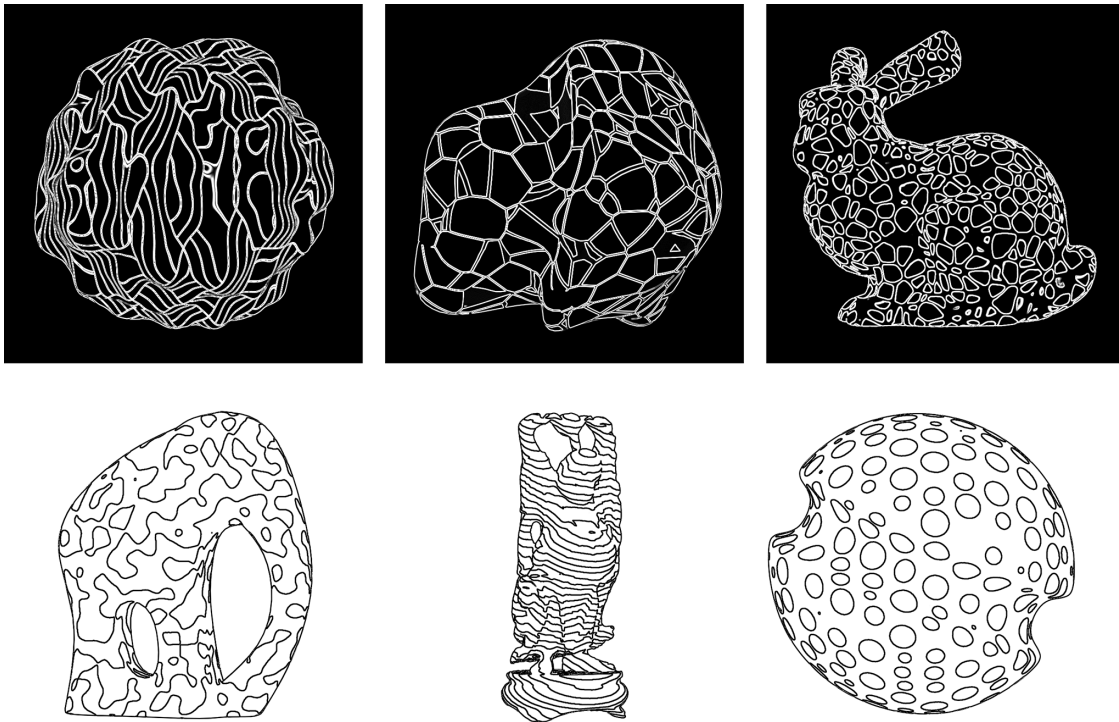
Figure 13. Contour drawings of six objects rendered with 3D volumetric textures from Experiment 2.

# Experiment 2

Because all of the contour patterns in Experiment 1 were primarily identified as glass, one possible explanation of the results is that observers are inherently biased to interpret any contour pattern as glass regardless of how they are oriented with respect to the surface geometry. Experiment 2 was designed to test that hypothesis using contour patterns that are completely unrelated to those that arise from specular reflections and interreflections.

## Methods

The stimuli in Experiment 2 were the same as those used in the previous study, except that the chrome contour patterns were replaced by a new set of contour patterns that were created by rendering each object with a 3D volumetric texture. Each of the nine objects was rendered with a different texture and then converted into a contour pattern using the Photoshop find edges tool with both types of contrast polarity. Six of these texture-object combinations are shown in Figure 13 as white on black in the top row and black on white in the bottom row. In all other respects, the methods were identical to those described for Experiment 1.

Eight new naïve observers were recruited, and they each participated in just one experimental session.

## Results

Table 2 shows the average confidence ratings for all of the different response categories in all of the different stimulus conditions. As in Experiment 1, the observers' judgments of the rendered images were quite accurate. The chrome materials were categorized as metal with an average confidence rating of 97%; the shiny black materials were categorized as shiny black with an average confidence rating of 90%; the glass materials were categorized as glass with an average confidence rating of 96%; and the shiny white materials were categorized as something else with an average confidence rating of 84%.

Unlike Experiment 1, however, there was no ambiguity at all between the two types of contour patterns. The glass contours were categorized as glass with an average confidence rating of 91% for both of the different contrast polarities. Conversely, the volumetric texture patterns were categorized as something else with an average confidence rating of 97%.

It is also interesting to note in evaluating these data that six of the eight main conditions were identical to those used in Experiment 1 with different subjects. Thus this makes it possible to measure the test-retest

|  | Shaded glass | | | | Shaded metal | | | | Shaded black | | | | Shaded white | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
|  | Response category | | | | Response category | | | | Response category | | | | Response category | | | |
| Object | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other |
| Yoyo | 96 | 3 | 0 | 1 | 0 | 91 | 0 | 9 | 10 | 1 | 87 | 3 | 0 | 14 | 0 | 86 |
| Pepper | 88 | 0 | 0 | 13 | 0 | 94 | 0 | 6 | 3 | 6 | 91 | 0 | 0 | 22 | 0 | 78 |
| Elephant | 88 | 0 | 0 | 13 | 0 | 99 | 0 | 1 | 0 | 15 | 85 | 0 | 0 | 9 | 0 | 91 |
| Cobble | 100 | 0 | 0 | 0 | 0 | 96 | 0 | 4 | 1 | 11 | 88 | 0 | 3 | 22 | 0 | 75 |
| Glasses | 100 | 0 | 0 | 0 | 0 | 98 | 3 | 0 | 0 | 4 | 96 | 0 | 0 | 9 | 0 | 91 |
| Bumpy | 99 | 0 | 0 | 1 | 0 | 98 | 0 | 3 | 1 | 16 | 83 | 0 | 0 | 19 | 0 | 81 |
| Bunny | 93 | 0 | 0 | 8 | 1 | 96 | 0 | 3 | 1 | 4 | 95 | 0 | 10 | 6 | 0 | 84 |
| Woman | 100 | 0 | 0 | 0 | 0 | 100 | 0 | 0 | 1 | 13 | 87 | 0 | 0 | 20 | 0 | 80 |
| Buddha | 100 | 0 | 0 | 0 | 1 | 99 | 0 | 0 | 0 | 4 | 96 | 0 | 0 | 9 | 0 | 91 |
| **Average** | **96** | **0** | **0** | **4** | **0** | **97** | **0** | **3** | **2** | **8** | **90** | **0** | **1** | **14** | **0** | **84** |

|  | Glass contours black background | | | | Glass contours white background | | | | Texture contours black background | | | | Texture contours white background | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Object | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other | Glass | Metal | Black | Other |
| Yoyo | 91 | 0 | 6 | 3 | 93 | 5 | 0 | 3 | 0 | 0 | 6 | 94 | 1 | 0 | 0 | 99 |
| Pepper | 89 | 3 | 8 | 0 | 89 | 9 | 0 | 3 | 0 | 0 | 8 | 93 | 4 | 1 | 0 | 95 |
| Elephant | 90 | 0 | 10 | 0 | 96 | 3 | 0 | 1 | 0 | 0 | 8 | 93 | 3 | 0 | 0 | 98 |
| Cobble | 91 | 0 | 9 | 0 | 91 | 0 | 0 | 9 | 0 | 0 | 1 | 99 | 0 | 0 | 0 | 100 |
| Glasses | 99 | 0 | 1 | 0 | 100 | 0 | 0 | 0 | 0 | 0 | 11 | 89 | 3 | 1 | 0 | 96 |
| Bumpy | 72 | 4 | 1 | 23 | 84 | 1 | 0 | 15 | 0 | 1 | 1 | 98 | 0 | 5 | 0 | 95 |
| Bunny | 97 | 1 | 1 | 1 | 94 | 4 | 0 | 3 | 0 | 0 | 1 | 99 | 0 | 0 | 0 | 100 |
| Woman | 98 | 0 | 2 | 0 | 94 | 1 | 0 | 5 | 0 | 0 | 1 | 99 | 0 | 0 | 0 | 100 |
| Buddha | 89 | 1 | 0 | 9 | 80 | 6 | 0 | 14 | 0 | 0 | 0 | 100 | 1 | 0 | 0 | 99 |
| **Average** | **91** | **1** | **4** | **4** | **91** | **3** | **0** | **6** | **0** | **0** | **4** | **96** | **1** | **1** | **0** | **98** |

Table 2. The average confidence ratings for each stimulus and response category of Experiment 2. All of the cells in this table have standard errors that are less than 5%. The rows of each stimulus category do not all sum to 100 due to rounding error.
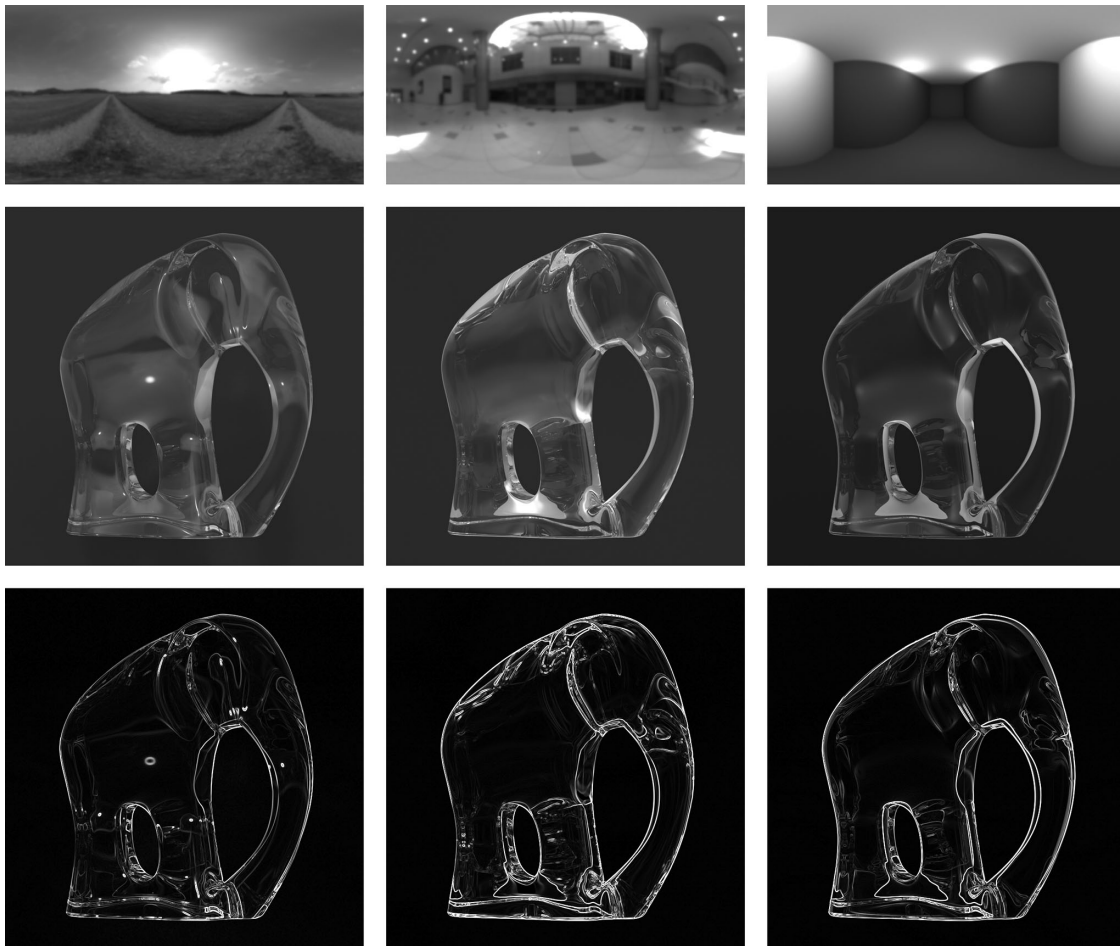
Figure 14. Three different light maps (top row). Images of an elephant rendered from them (middle row), and contour drawings created from those images.

reliability of this task. The results revealed that the shared cells in Tables 1 and 2 were almost perfectly correlated (*r* = 0.99).

## General discussion

It has long been known that contour drawings without any smooth shading provide considerable information about the 3D shapes of objects (e.g., DeCarlo, Finkelstein, Rusinkiewicz & Santella, 2003; Judd, Durand & Adelson, 2007). However, the results of the present experiments are the first to show that contour patterns presented in isolation can provide sufficient information to perceptually categorize glass materials. Although it is possible that glass may have a unique status in this regard, we suspect that is not the case. It is important to keep in mind that many materials such as wood or marble are defined primarily by their patterns of texture, which remain relatively invariant when an image is transformed using edge

filtering. Thus it is likely that those materials may also be identifiable from contour drawings.

Our informal observations suggest that edge filtered images of glass objects remain surprisingly invariant over large changes in the pattern of illumination. The top row of Figure 14 shows three light maps that are different from the one used in the present experiments. The middle row shows rendered images of an elephant that were illuminated by each of those light maps, and the bottom row shows edge filtered versions of those images in a white on black format. Although there are clearly noticeable differences among the patterns of shading, the variations among the contour patterns are more subtle. The occlusion contours obviously do not change at all. However, the high frequency contour patterns that are due to internal specular interreflections are more interesting. The precise details of those patterns vary among the different illumination fields, but the locations of those patterns on the object's surface are mostly the same. Whereas the pattern details are mostly influenced by the structure of the surrounding scene, the pattern

positions are mostly determined by the locations of surface concavities along the internal boundary of an object.

The results of the present experiments provide strong evidence that observers can distinguish the different types of contours from one another. This is most clear in Experiment 2 where the glass contours were categorized as glass with 91% confidence, and the texture contours were categorized as something else with 97% confidence. The perceptual distinction is less clear for the glass and chrome contours, especially because they were both categorized primarily as glass. However, the glass confidence ratings were significantly lower in the chrome contour condition, thus indicating that those contours were less perceptually convincing as glass than the contours that were produced from images of glass objects.

It is interesting to note that contour labeling is one of the oldest problems in the field of computer vision (Guzman, 1968; Waltz, 1975). Although there are some effective algorithms for categorizing the contours in line drawings (Malik, 1987), they are generally restricted to drawings that only depict sharp corners and occlusion boundaries, and they do not work well with images that contain other types of structure, such as shadows, specular reflections, or surface textures. Human observers, in contrast, can easily label all of these features.

One important influence on the contour structures in Experiment 1 were the specular highlights arising from direct reflections in all of the different materials we used. However, those structures were identical for the glass, shiny black, and shiny white materials, so they could not have provided the relevant information for distinguishing those materials from one another (see Norman et al., 2020). Reflectance contours on shiny metal objects include all of the same ones that are visible for dielectric materials, but because they have higher contrast, they also include some additional ones that are not visible on less reflective objects.

A more likely source of visual information about glass is provided by banding contours produced by internal specular interreflections. Banding contours can also occur on shiny opaque surfaces, but they align quite differently with the pattern of surface curvature. For transparent glass, these contours surround concave regions on the internal boundary of the surface, whereas for opaque metals, they surround concave regions on the external boundary. Another important difference between glass and metal contours is that glass contours have a more chaotic structure, with local swirling patterns that resemble the turbulent flow of fluid materials. Observers can apparently detect these differences. Even though the chrome contours in Experiment 1 were primarily categorized as glass, the observers' confidence ratings were significantly lower than those obtained for the glass contours. The difference between glass contours and the volumetric

textures is even greater, because the texture contour alignments are completely independent of surface curvature.

Fleming et al. (2004) have argued that specular reflections on shiny surfaces can be analyzed in much the same way as optical texture (e.g., see Egan, Todd, & Phillips, 2011; Fleming, Holtmann-Rice, & Bülthoff, 2011; Todd & Akerstrom, 1987; Todd & Reichel, 1990; Todd & Thaler, 2010). The basic idea is that reflected features of a scene stretch out in directions perpendicular to the surface depth gradients, and are compressed as a cosine function of the magnitudes of those gradients. Similar ideas involving oriented flow gradients have also been proposed for the analysis of shading on matte surfaces (Bretan & Zucker, 1996; Kunsberg, Holtmann-Rice, Alexander, Cholewiak, Fleming, & Zucker , 2018; Kunsberg & Zucker, 2018; Pont & Koenderink, 2008, Pont, van Doorn, Wijntjes, & Koenderink, 2015).

It is especially interesting to note in this regard that the presence of interreflections can severely distort the scalar field of image intensities, relative to what occurs when only direct reflections are considered. In order to explore that phenomenon, we created maps of the isointensity contours in our images, using a representation from Todd, Egan and Phillips (2014) in which bands of image intensity are replaced by homogeneous colors that are separated by small black bands. The flow gradients are also captured in this representation by the spacing between bands. Figure 15 shows the results of that analysis for chrome and glass images of three different objects. The top row shows isointensity maps of the chrome objects. Although these patterns are quite complex, the bands of similar intensity have well defined boundaries. However, that is not the case for the isointensity maps of the glass images in the bottom row. We have tried several different band sizes for this representation, but they all appear as a chaotic mess. There are at least two possible reasons for this. One is that the image resolution of 800 × 800 pixels is too coarse to resolve the high frequency structure of the patterns. Another is that reflections of glass can have a fractal structure that cannot be smoothly resolved at any scale.

To test these possibilities, we rendered one last image of our glass bunny object at a very high resolution of 8000 × 8000 pixels and a high quality level of 25. The object was illuminated using the unblurred esplanade light map shown in Figure 8. The processing was performed on a dedicated computer cluster with 96 Xeon cores. It took 60 hours to render the scene, and the resulting image file was just under 50 Mb. It is obviously not possible to present that image without down sampling, but we can present local patches to see how the structure is resolved at very small scales. Figure 16 shows three representative patches with a spatial resolution of 800 × 800 pixels. Note that the patterns are still chaotic even at that fine spatial scale, which
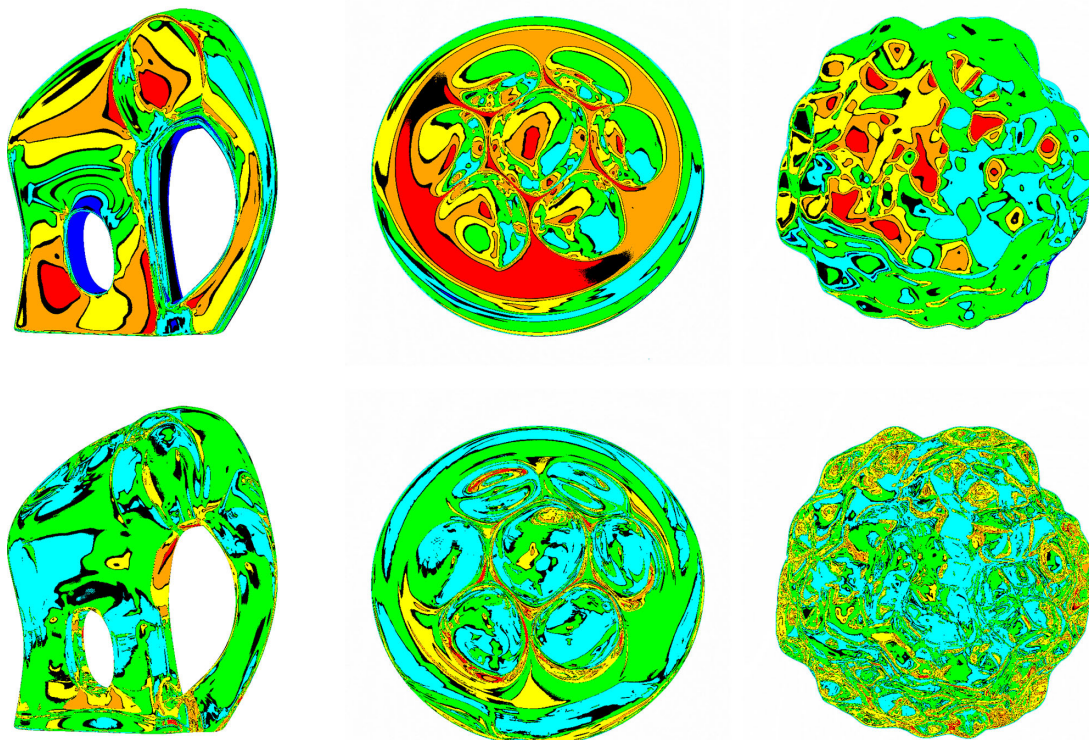
Figure 15. Isointensity maps for three chrome objects (top row) and three glass objects (bottom row).
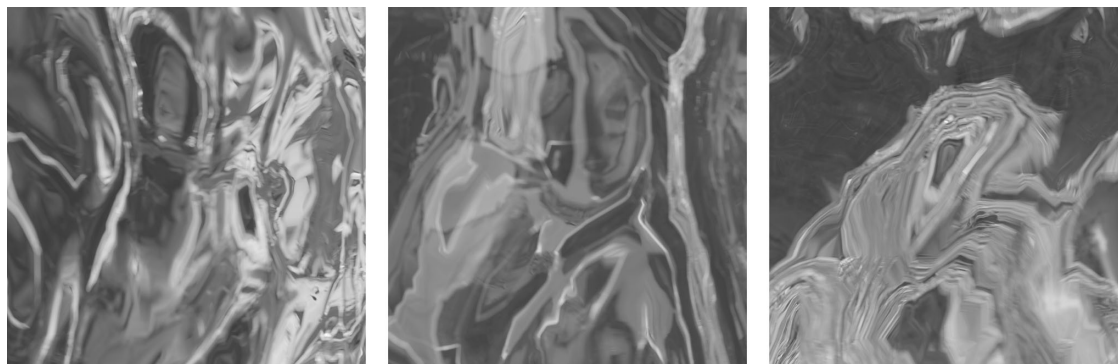


Figure 16. Three 800 × 800 pixel patches from an 8000 × 8000 pixel image of a glass bunny.

suggests reflections of glass may indeed have a fractal structure. It should also be pointed out in this regard that the patterns would have been even more complex if we had not used a neutral gray background surface to occlude the environmental structure behind the object (see Todd & Norman, 2019).

## Conclusions

The present experiments have examined how the perceptual identification of glass is influenced by banding contours formed by internal specular interreflections within glass materials. The results reveal that observers can reliably distinguish internal banding contours from those that arise from the specular reflections of chrome materials or 3D volumetric textures. When edge filtered contours of glass objects are presented without any smooth shading gradients, the depicted objects are reliably identified as glass with a high level of confidence. The edge filtered contours from chrome objects are also identified as glass, but with a significantly lower level of confidence. Finally, the contours created from volumetric textures are never identified as glass and are rated as something else with a high level of confidence. These findings suggest that

observers are sensitive to the smoothness (or chaotic structure) of the contours and how they align with the overall pattern of surface curvature.

*Keywords: contours, interreflections, glass*

## Acknowledgments

## References

Bloj, A., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature, 402*, 877–887.

Breton, P., & Zucker, S. W. (1996). Shadows and shading flow fields. In *Proceedings of CVPR*, San Francisco, CA, 782–789.

DeCarlo, D., Finkelstein, A., Rusinkiewicz, S., & Santella, A. (2003). Suggestive contours for conveying shape. *ACM Transactions on Graphics, 22*(3), 848–855.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2004). Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *Journal of Vision, 4*, 92–105, doi:10.1167/4.2.3.

Egan, E. J. L., Todd, J. T., & Phillips, F. (2011). The perception of 3D shape from planar cut contours. *Journal of Vision, 11*(12):15, 1–13, http://www.journalofvision.org/content/11/12/15., doi:10.1167/11.12.15.

Fleming, R. W., Holtmann-Rice, D., & Bülthoff, H. H. (2011). Estimation of 3D shape from image orientations. *Proceedings of the National Academy of Sciences, 108*(51) 20438–20443

Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision, 4*(9), 798–820.

Gilchrist, A., & Jacobsen, A. (1984). Perception of lightness and illumination in a world of one reflectance. *Perception, 13*, 5–19, doi:10.1068/p130005.

Guzman, A. (1968). Computer recognition of three-dimensional objects in a scene. MIT Tech. Rept. MAC-TR-59.

Hunter, F., Biver, S., & Fuqua, P. (2007). *Light, science & magic: An introduction to photometric lighting*. Amsterdam, the Netherlands: Elsevier.

Judd, T., Durand, F., & Adelson, E. (2007). Apparent ridges for line drawing. *ACM Transactions on Graphics, 26*(3), 19.

Koenderink, J. J., & van Doorn, A. J. (1983). Geometrical modes as a general method to treat diffuse interreflections in radiometry. *Journal of the Optical Society of America, 73*, 843–850.

Kunsberg, B., Holtmann-Rice, D., Alexander, E., Cholewiak, S., Fleming, R. W., & Zucker, S. W. (2018). Colour, contours, shading and shape: flow interactions reveal anchor neighbourhoods. *Interface Focus, 8*, 20180019.

Kunsberg, B., & Zucker, S. (2018). Critical contours: An invariant linking Image flow with salient surface organization. *SIAM Journal on Imaging Sciences, 11*, 1849–1877.

Langer, M. S. (1999). When shadows become interreflections. *International Journal of Computer Vision, 34*, 193–204.

Malik, J. (1987). Interpreting line drawings of curved objects. *International Journal of Computer Vision, 1*, 73–103.

Madison, C., Thompson, W., Kersten, D., Shirley, P., & Smits, B. (2001). Use of interreflection and shadow for surface contact. *Perception & Psychophysics, 63*, 187–194, doi:10.3758/BF03194461.

Mooney, S. W. J., & Anderson, B. L. (2014). Specular image structure modulates the perception of three-dimensional shape. *Current Biology, 24*, 2737–2742.

Nayar, S. K., Ikeuchi, K., & Kanade, T. (1991). Shape from interreflections. *International Journal of Computer Vision, 6*, 173–195.

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, doi:10.1167/jov.20.5.2.

Pont, S. C., & Koenderink, J. J. (2002). Bidirectional reflectance distribution function of specular surfaces with hemispherical pits. *Journal of the Optical Society of America A, 19*, 2456–2466.

Pont, S. C., & Koenderink, J. J. (2005). Reflectance from locally glossy thoroughly pitted surfaces. *Computer Vision and Image Understanding, 98*, 211–222.

Pont, S. C., & Koenderink, J. J. (2008). Shape, surface roughness, and human perception. In M. Mirmehdi, X. Xie, & J. Suri (Eds.), *Handbook of Texture*

*Analysis* (197–222). London, UK: Imperial College Press.

Pont, S. C., van Doorn, A. J., Wijntjes, M. W. A., & Koenderink, J. J. (2015). Texture, illumination, and material perception. *Proc. SPIE 9394, Human Vision and Electronic Imaging XX, 93940E (17 March 2015)*; doi:10.1117/12.2085953940Q., https://doi.org/10.1117/12.2085021.

Todd, J. T., & Akerstrom, R. A. (1987). The perception of three dimensional form from patterns of optical texture. *Journal of Experimental Psychology: Human Perception and Performance, 2*, 242–255.

Todd, J. T., Egan, E. J. L., & Kallie, C. S. (2015). The darker-is-deeper heuristic for the perception of 3D shape from shading: Is it perceptually or ecologically valid? *Journal of Vision, 15*(11):24, 1–9.

Todd, J. T., Egan, J. L., & Phillips, F. (2014). Is the Perception of 3D shape from shading based on assumed reflectance and illumination? *i-Perception, 5*, 497–514.

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.

Todd, J. T., & Norman, J. F. (2019). Reflections on glass. *Journal of Vision, 19*(4):26, 1–21, https://doi.org/10.1167/19.4.26.

Todd, J. T., Norman, J. F., & Mingolla, E. (2004). Lightness constancy in the presence of specular highlights. *Psychological Science, 15*, 33–39.

Todd, J. T., & Reichel, F. D. (1990). Visual perceptions of smoothly curved surfaces from double-projected contour patterns. *Journal of Experimental Psychology: Human Perception and Performance, 16*(3), 665–674.

Todd, J. T., & Thaler, L. (2010). The perception of 3D shape from texture based on directional width gradients. *Journal of Vision, 10* (5):17, 1–13.

Waltz, D. (1975). Understanding the drawings of scenes with shadows. In P.H. Winston (Ed.), *The Psychology of Computer Vision*. New York: McGraw-Hill.

Zhang, F., de Ridder, H., & Pont, S. C. (2015). The influence of lighting on visual perception of material qualities. In *Proceedings of SPIE 9394, Human Vision and Electronic Imaging XX*, *9394*, 93940Q.