

Specular reflections and the estimation of shape from binocular disparity

Alexander A. Murry^a, Andrew E. Welchman^{a,1}, Andrew Blake^b, and Roland W. Fleming^c

^aSchool of Psychology, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom; ^bMicrosoft Research Ltd., Cambridge CB3 0FB, United Kingdom; and ^cDepartment of Psychology, University of Giessen, 35394 Giessen, Germany

Edited by Brian A. Wandell, Stanford University, Stanford, CA, and approved December 19, 2012 (received for review August 3, 2012)

Binocular stereopsis is a powerful visual depth cue. To exploit it, the brain matches features from the two eyes' views and measures their interocular disparity. This works well for matte surfaces because disparities indicate true surface locations. However, specular (glossy) surfaces are problematic because highlights and reflections are displaced from the true surface in depth, leading to information that conflicts with other cues to 3D shape. Here, we address the question of how the visual system identifies the disparity information created by specular reflections. One possibility is that the brain uses monocular cues to identify that a surface is specular and modifies its interpretation of the disparities accordingly. However, by characterizing the behavior of specular disparities we show that the disparity signals themselves provide key information ("intrinsic markers") that enable potentially misleading disparities to be identified and rejected. We presented participants with binocular views of specular objects and asked them to report perceived depths by adjusting probe dots. For simple surfaces—which do not exhibit intrinsic indicators that the disparities are "wrong"—participants incorrectly treat disparities at face value, leading to erroneous judgments. When surfaces are more complex we find the visual system also errs where the signals are reliable, but rejects and interpolates across areas with large vertical disparities and horizontal disparity gradients. This suggests a general mechanism in which the visual system assesses the origin and utility of sensory signals based on intrinsic markers of their reliability.

psychophysics | perception | gloss | texture | computational analysis

Shiny objects such as sports cars, jewelry, and consumer electronics can be beautiful to look at. However, such objects pose a difficult challenge to the visual system: if all (or most) of the light reaching the eye comes from the reflections of other nearby objects, how does the viewer discern the object itself? This problem becomes more acute when viewing with two eyes. Unlike shading or texture markings, the positions of reflections relative to a specular (shiny) surface depend on the observer's viewpoint. This means that when the surface is viewed binocularly (i.e., from two viewpoints at the same time), corresponding reflections fall on different surface locations. In consequence, the binocular disparities created by specular reflections indicate depth positions displaced from the object's physical surface (1, 2) and the 3D shape specified by disparity can be radically different from the true shape of the object. For special cases, such as an ideal planar mirror, the visual system could not, even in principle, estimate the true depths of the surface from the reflections. However, for more complex shapes, such as a polished metal kettle, we rarely encounter problems judging shape. Most models of biological vision place heavy weight on binocular disparity cues, whereas artificial systems often rely almost exclusively upon them. How does human vision recover the depth of these specular objects?

We suggest that the brain's treatment of specular reflections is likely to exploit general mechanisms for assessing the origin, and utility, of sensory signals. Because specular reflections are a naturally occurring situation in which the visual system is faced with potentially large discrepancies between different depth cues (i.e., disparity, shading, texture), they present a valuable opportunity to

gain insights into how the brain derives robust estimates from noisy, unreliable, or inconsistent information. In particular, there are two key problems: (i) how the brain discerns the cause of a given signal (e.g., does a disparity originate from a surface marking or a specular reflection?) and (ii) how the brain determines whether the information is trustworthy [e.g., how are statistically "optimal" cue weights chosen (3–6)?]. Here we exploit the natural discrepancies between depth cues that arise for specular objects to address these questions.

Broadly speaking, there are two general approaches that the visual system might use to identify and overcome the spurious binocular information from specular reflections. First, it has been suggested that the brain "knows" the physics of specular reflections (1). On this basis, once a disparity is identified as originating from specular reflection, the brain could apply specific computations to infer the true surface from the reflection. The process of identifying a surface as specular could exploit ancillary markers (4) (i.e., nonstereoscopic information that the surface is shiny), which alter the interpretation of disparities. In particular, there are believed to be several monocular cues that indicate surface specularity such as (i) the distribution of image intensities [surface highlights, lowlights, and other signals (7–9)], (ii) the elongation of image features (10, 11), and (iii) patterns of motion (12) and color (13). Thus, ancillary markers could indicate that a specular reflection model applies and therefore that disparities should be interpreted using "knowledge of the physics" of specular reflection. This approach is analogous to processes that alter the interpretation of scene lightness based on scene layout (14–16).

Alternatively the brain might exploit intrinsic markers—that is, characteristic properties of the signals themselves, such as their magnitude or distribution—to temper the use of disparity. We reason that rather than explicitly "knowing" the detailed, quantitative physics of specular reflections (i.e., having dedicated mechanisms for correctly interpreting disparity fields from specular reflections), the visual system may be able to detect when disparity signals are substantially abnormal and therefore reject them as untrustworthy.

To understand the roles played by ancillary cues and intrinsic markers, we first analyze the disparities produced by specular reflections to identify candidate signals that could act as intrinsic markers. Then, we use a custom stimulus generation method to pit monocular and binocular information against one another. In particular, we compare 3D shape judgments between renderings that have normal specular reflections with renderings in which we "paint" the reflections onto the object's surface so that they have

Author contributions: A.A.M., A.E.W., A.B., and R.W.F. designed research; A.A.M. and A.E.W. performed research; A.A.M., A.E.W., and R.W.F. analyzed data; and A.A.M., A.E.W., A.B., and R.W.F. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: a.e.welchman@bham.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1212417110/-DCSupplemental.

the same stereoscopic depths as the surface itself. The two types of renderings produce images in which the nonstereoscopic information is almost indistinguishable, but which differ critically in the stereoscopic behavior of the reflections. Thus, ancillary markers (i.e., the lustrous appearance of the surface derived from nonstereoscopic cues) should be very similar in the two conditions, whereas intrinsic markers (i.e., diagnostic properties of the disparity signals themselves) are quite different. By comparing observers' depth judgments for the two classes of stimuli, we sought to determine the relative influence of ancillary and intrinsic sources of information for perceptual estimates of shape.

Results

Physical Analysis of Binocular Specular Reflections. To understand the properties of binocular specular reflection, consider a sphere reflecting its surrounding environment (Fig. 1). A given feature in the environment is reflected into the two eyes by different locations on the sphere's surface, giving rise to a binocular disparity. By considering light information originating from all possible illumination directions, we can map out a surface defined by these disparities. This surface generally lies some distance away from the object's physical surface and we refer to it hereafter as the "virtual surface."

The location of the virtual surface is influenced in large part by the curvature of the object: it generally falls behind the true surface for convex surfaces and in front for concave surfaces (2). Moreover, the lower the curvature, the further the virtual surface is from the true surface. Thus, small changes in object shape or orientation can lead to large changes in the disparity field (Fig. 2A), highlighting the difficulty faced by the viewer in estimating the true shape of the object. Finally, because the viewing vectors can be skewed, they do not always intersect, creating large vertical disparities (see *Experimental Procedures* and Fig. S1 for a detailed explanation). If the brain took the disparities produced by specular surfaces at face value, its estimate of shape would be completely incorrect.

To test whether the visual system uses ancillary cues (i.e., monocular signals that indicate that a specular object is being viewed) in modifying the interpretation of disparities, we developed a stimulus rendering procedure in which we effectively paint the reflections onto the surface of the object (Fig. 2B). This creates stereopairs in which the monocular images are almost identical to the specular case, but the disparities indicate the true physical surface positions (Fig. 2C). We use this painted case to illustrate that these physical-surface-based disparities are quite different from the virtual surface. Consider the simple near-spherical object depicted in Fig. 3A (we describe such objects as "muffins"; in this case the object's deviation from a sphere is so

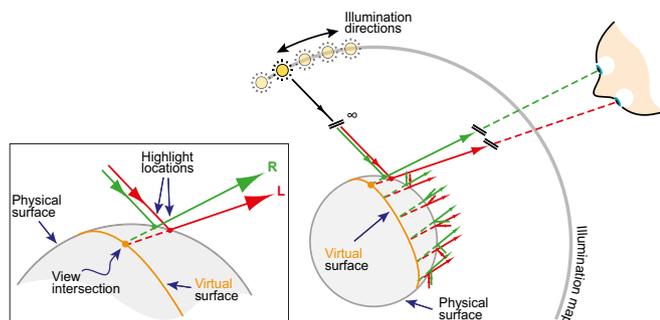


Fig. 1. Physical properties of binocular specular reflections. Reflected features (highlights) appear at different locations in the two eyes' views. Corresponding view rays intersect at locations that are not on the physical surface. Considering light arriving from all different view directions and identifying the matching reflected ray vectors, traces out the "virtual surface."

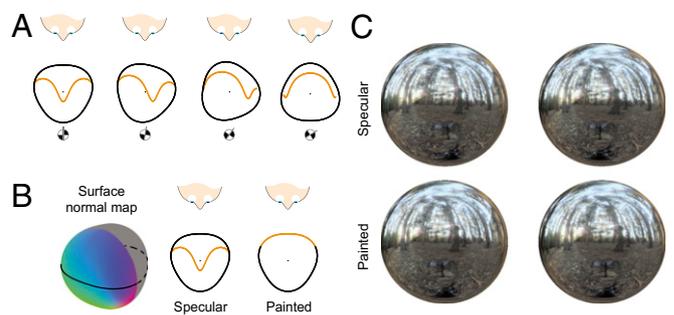


Fig. 2. Properties of the virtual surface and "painted" vs. specular surface rendering. (A) Virtual surface is highly sensitive to both surface and viewing geometry. It can be concave even though the physical surface is convex, and small object rotations lead to large changes in its shape. (B) Near-spherical "muffin" whose surface normals are illustrated using a red-green-blue color representation. Cartoons depict the disparity-defined profile for a slice through the equator of the shape: One for the object rendered with a specular surface, the other for the "painted" case where the illumination map is essentially stuck onto the physical surface. (C) Stereopairs (cross-eyed fusion) for both the painted and specular views of the object from B.

slight that it is practically indistinguishable). If the object is painted (Fig. 3B), viewing the shape binocularly gives rise to a shape-specific distribution of horizontal and vertical disparities across the image (17). Now consider the same shape with a specular surface (Fig. 3C). The pattern of horizontal disparities is quite different (i.e., the shape of the virtual manifold differs from that of the surface, as in Fig. 2A) and generally contains greater extremes of horizontal disparity gradient. Moreover, note that although the distribution of vertical disparities in the image is qualitatively similar (i.e., a "cross" centered on fixation at zero elevation and azimuth in Fig. 3C), the magnitude of the vertical disparities increases dramatically. Thus, even for a relatively simple shape (smooth, convex, and almost spherical), specular reflections can give rise to binocular signals that have large magnitudes.

Now consider viewing a more complex 3D shape (Fig. 4) that is globally convex, but contains local concavities (we describe such objects as "potatoes"). These surfaces give rise to discontinuities in the virtual surface (around surface inflection points) with some isolated portions in front of the true surface and with the majority behind it. This results in very large horizontal disparity gradients that often exceed human fusion limits (18, 19). Further, there are locations where there are potentially one-to-many matches (i.e., multiple disparities along a line of sight), and other locations for which no disparity is defined for a given cyclopean direction (i.e., holes in the disparity field). Moreover, the distribution of vertical disparities across the image becomes extremely unusual: large vertical disparities are experienced near the point of fixation, which does not occur for Lambertian objects.

Based on this analysis, we have identified potential intrinsic markers that indicate specular reflection (i.e., horizontal disparity outliers, horizontal disparity gradients, and large vertical disparities) and shown that the strength of these markers varies with the 3D shape. Moreover, our painted rendering method allows us to test the role of ancillary markers by presenting two different classes of stimuli (painted vs. specular) that are practically indistinguishable when viewed monocularly (Fig. 2C).

Behavioral Measures of 3D Shape Perception. We tested human participants to determine whether the visual system correctly interprets disparities when estimating the shape of specular objects [i.e., the brain "knows the physics of specular reflections" (1)] or whether it is biased toward the erroneous depth indicated by the virtual surface. Importantly, the intrinsic disparity markers

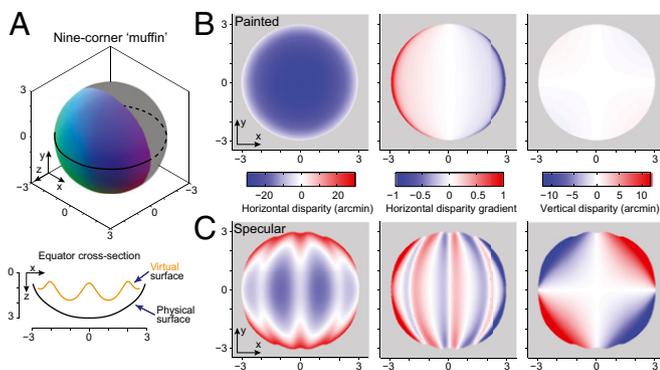


Fig. 3. Example disparity fields of a muffin object. (A) Three-quarters view of a near-spherical nine-cornered muffin (corners are practically invisible). The muffin's surface normals are represented using red-green-blue. A line is projected around the object's equator and the depth profile of the physical and virtual surfaces along this line is shown in cross-section. (B) Maps of horizontal disparity, horizontal disparity gradients, and vertical disparity for a painted nine-cornered muffin. The object is viewed along the depth axis; x , y image locations are in centimeters. The red-blue color code indicates the magnitude of each quantity (color bars are scaled for each column). (C) Pattern of disparities in the virtual manifold produced when viewing the nine-cornered muffin with a specular surface (color scale matched to B).

are weak in magnitude for the muffins, but stronger for the potatoes. Thus, if the brain relies on intrinsic markers, then the perceived depth profiles should follow the virtual surface for the muffins, even though the surface is clearly specular. By contrast, for the potatoes, the brain should reject the disparities wherever the intrinsic markers indicate the disparities are incoherent.

We presented subjects with computer-generated images of painted or perfectly specular (i.e., mirrored) surfaces (Fig. 5A) and asked them to report perceived surface shape by adjusting the binocular disparity of small probes until they appeared to lie on the surface. In both cases, the monocular appearance specified purely specular surface reflection properties, without any diffuse component or texture markings. Thus, we are able to isolate specular disparities from other shape cues in a manner analogous to the use of random dot stereograms to isolate standard disparity signals. As a baseline measure, we also created stimuli whose surfaces combined specular reflectance with both diffuse and texture components, thereby creating disparity layers corresponding to the physical surface as well as the virtual manifold.

In the first experiment, we tested four simple convex shapes (muffins) in three different orientations relative to the viewer. Subjects adjusted 11 probes arranged along a horizontal raster across the midline of the shape. These shapes and probe locations were carefully selected so that although the virtual disparities were “incorrect,” they did not exhibit the large gradients, vertical disparities, or other effects that could indicate that they were unreliable. Nevertheless, the shapes indicated by the disparities were dramatically inconsistent with the true (near-spherical) shapes of the objects, as signaled by monocular cues (11) (Fig. 3C).

When fusing the painted and specular stereopairs, viewers typically appreciate marked differences between the two: the specular shape has reduced amplitude and a quite different topography (Figs. 2C, 5A, and 6A). Our observers reported considerable differences between the apparent depths of the painted and specular muffin objects (Fig. 5B). For the painted stimuli (blue squares), subjects were highly accurate at placing the probes at the true depths of the surfaces, suggesting that ancillary markers of specular reflection are unlikely to play a strong role in the interpretation of binocular disparity signals. For the specular stimuli (orange circles) the subjects' settings lie very close to the virtual surface, rather than the true physical surface.

This suggests that the brain interprets the disparity signals “at face value,” as if they indicated true surface locations, rather than a virtual manifold behind the physical surface. Fig. 5C summarizes the average data from all subjects and shapes, confirming that for the painted condition, settings were very close to the physical surface, whereas for the specular condition, they were much closer to the virtual surface prediction.

In our second experiment, we tested more complex potato shapes, containing both convexities and concavities, and measured probe responses along three horizontal raster lines. Recall that for these stimuli, the disparity field is much more complex than for the muffins (Fig. 4). If the visual system is unable to estimate the depth profiles of specular surfaces when the surface geometry is simple, we might expect that it would perform as badly, or worse, with more complex shapes.

Surprisingly, we found that subjects' settings for potatoes did not always fall on the virtual surface, and indeed there were some locations where the settings lie closer to the true surface than to the virtual surface (Fig. 6B). In some portions of the shape, observers' settings conform to the virtual surface; however, at other locations they do not. How does the visual system determine whether or not to rely on the disparity information at a given location? One possibility is that the unusual properties of specular disparities outlined above act as an intrinsic marker that the underlying binocular information should therefore be rejected when estimating 3D shape. To test this idea, we applied a “disparity detectability constraint (DDC),” using performance limits of the human visual system on outliers, horizontal disparity gradients, and the magnitude of vertical disparities (19–21) to identify and remove unreliable portions of the virtual surface. We then considered participants' estimated depth judgments by separating settings from the “specular” condition into two classes: (i) those that come from reliable portions of the virtual surface, and (ii) those that come from portions that are identified as unreliable by the DDC criteria. Fig. 6C shows that for the reliable portions of the virtual surface (orange dots) settings are clustered close to the virtual prediction, as occurred with the muffin stimuli (Fig. 5C). In other words, whenever the disparity signals are reliable, the brain tends to treat them at face value, as

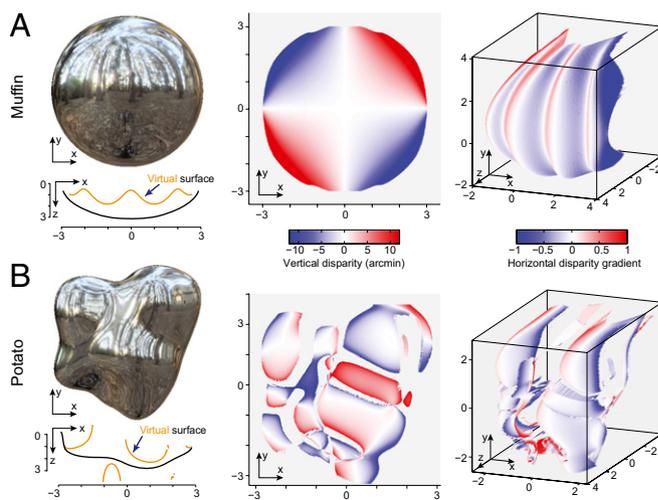
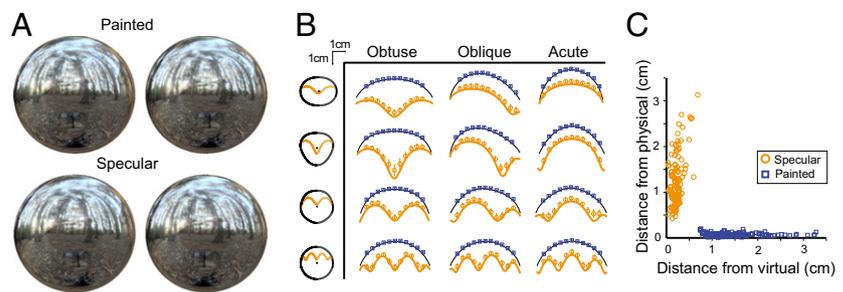


Fig. 4. Illustrations of virtual surface profile of a muffin and potato object. Stimulus examples, the front view, and side view of the virtual surface for a specular muffin (A) and potato (B). Below the stimulus, a depth profile depicts the physical and virtual surfaces as they vary across a horizontal slice through the shape. For the *Center* and *Right* columns, the color codes represent vertical disparities (*Center*) and horizontal disparity gradients (*Right*). Note that vertical disparities can be very large for the potato object.

Fig. 5. Example stimuli and results for muffin objects. (A) Stereopairs for cross-fusing. When viewed monocularly, painted and specular stereopairs are almost indistinguishable. When fused, the painted stereopair appears more matte and more spherical. (B) Data from the probe adjustment task. Each row is a single shape viewed at three orientations relative to the viewer. Black line, physical surface; orange line, virtual surface. Settings are shown for painted (blue squares) and specular (orange circles) conditions. Error bars show SEM across observers. (C) Scatterplot of settings in terms of distance from the physical (y axis) and virtual (x axis) surfaces. Each datum represents the mean setting (across repetitions) of an individual observer for a particular location on one of the shapes.



if they indicated the actual position of the surface (which, of course they do not). By contrast, where the disparity signal is atypical (Fig. 6C, gray circles), the settings are broadly distributed and are in many cases closer to the physical surface than the “good” virtual surface disparities.

What can account for this perhaps counterintuitive result that settings are closer to the physical surface when specular disparities are atypically large or incoherent? One possibility is that the visual system could interpolate across the gaps caused by these signals, basing its estimates on the more reliable information flanking the unreliable regions and shape information from monocular cues. To test this idea, we filtered out unreliable disparities and interpolated across the resulting gaps in the virtual surface using Bezier curve fits as a simple way of imposing a smoothness constraint (22) on the interpolated surface. These prediction curves are qualitatively quite similar to the subjects’ settings (Fig. 6B, black dashed line), although note these fits are not intended as a quantitative or biological model of spatial interpolation—other interpolation methods may fit subjects’ settings more closely. The important idea is that the brain appears to use some kind of spatial interpolation to deal with missing, inconsistent, or otherwise untrustworthy disparity signals. Such a strategy would be applicable not only to specular surfaces, but to other “bad” disparity signals as well, such as in refractive media (e.g., a heat haze), or when retinal or eye-movement noise leads to spurious vertical disparity signals, or where contrast is locally too low for disparity to be measured reliably.

Whereas many objects encountered in the natural world have a specular component, purely specular surfaces are relatively rare—most materials have some combination of specular and diffuse reflection. In the main experiments, we used purely specular surfaces to isolate the information provided by specular disparities. However, it is interesting to ask how well the brain estimates shape when additional cues are present. We therefore obtained settings for objects that had partially specular surfaces (i.e., combinations of shading, texture, and specular reflections). These objects provided information about the physical location

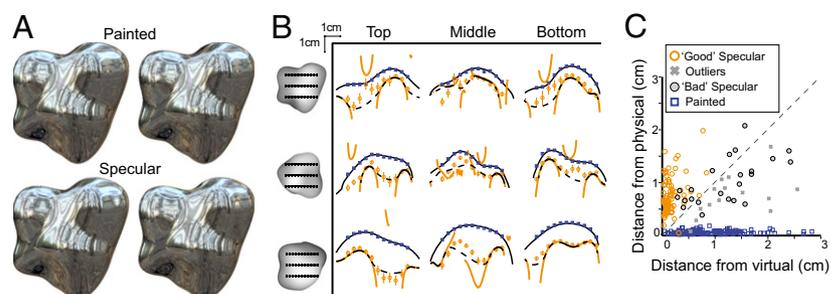
of the surface from the binocular disparities associated with surface texture and shading, as well as information about the virtual surface overlaid in a physically realistic manner. Unsurprisingly, we found that when a textured surface component was visible, observers’ settings lay on the physical surface of the object (Fig. S2). This suggests that when segmenting the two potential surface locations, observers select the one that is most coherent as the one likely to represent the true surface location. As the relative strength of the surface markings change (e.g., high spatial frequency texture marks become visible relative to low spatial frequency shading signals), the observer’s impression of the surface shape is likely to change.

Discussion

Previous work has suggested that specular highlights can aid 3D shape perception (23, 24), especially when combined with other cues (25, 26). However, the presence of a specular highlight does not always influence shape judgments (27) and these signals may sometimes be ignored (28). Here, by isolating the specular disparity cue, we have identified the specific image quantities that the brain could use to reject potentially misleading disparities.

It is important to clarify that the extreme values of vertical disparities and horizontal disparity gradients that are rejected by the disparity detectability constraint are unresolvable and therefore probably not encoded at all by the visual system. These portions also tend to be flanked by regions that are fusible but still contain unusual values. Our computational analysis shows that for complex shapes (potatoes), fusible areas appear to be isolated regions surrounded by unresolvable areas. These fusible “islands” correspond to regions of local convexity and concavity, which are isolated from each other by inflection contours where the virtual surface depths go to infinity. At the borders of these regions the vertical disparities and horizontal disparity gradients reach their maximum, beyond which disparities become unavailable. Recent studies show that the visual system is sensitive to rapid changes of sign in the vertical disparities (29). It is possible that a similar mechanism could help the visual system to identify regions of

Fig. 6. Example stimuli and results for potato objects. (A) Stereopairs for cross-fusing. (B) Data from the probe adjustment task. Each row is a single shape; columns show three rasters through the object (locations illustrated on gray-scale depictions next to the axes). Black line, physical surface; orange line, virtual surface. Settings are shown for painted (blue squares) and specular (orange circles) conditions. Error bars show SEM across subjects. (C) Relative proximity of settings to the physical and virtual surfaces. Data from the specular surfaces were separated into three groups: “good” vs. “bad” specular were identified using the disparity detectability constraint; “outliers” corresponded to locations that were below the DDC, but were outliers in relation to the surrounding points. Each datum indicates average probe settings for a given location on a given shape for a single observer.



unfusable disparities. Thus, the transitions from fusible to unfusable regions are not random, but have specific binocular properties that indicate the underlying signals are unreliable.

If the visual system rejects unreliable disparity signals in the way we have suggested, it is interesting to ask what happens when the stimulus contains only reliable or only unreliable signals. In the main experiments we tested this by comparing simple and complex shapes. However, another approach would be to isolate those locations within a given object that contain reliable or unreliable signals, respectively. In Fig. S3 we show what happens when the unreliable or reliable portions are selectively removed from the image, by hiding them behind an occluding surface. Observe that when the reliable portions of the specular reflection are the only portions visible, the surface appears to be smooth, coherent, and reliable, much like the muffins (Fig. 5A). By contrast, when the reliable portions are occluded, the remaining surface regions appear incoherent and difficult to interpret as a surface. This suggests, again, that the visual system uses a spatially localized measure of the trustworthiness of disparity signals—derived from the disparity signals themselves—rather than an ancillary marker based on the global monocular appearance of the material.

Until now, we have considered the role of monocular cues in providing ancillary cues to the material properties of the surface (i.e., specular or Lambertian) rather than as an additional source of information about 3D shape. However, monocular cues, including the occluding contours (30) and the compression of the surrounding environment (11), provide potentially useful information about surface structure. If intrinsic markers temper the use of disparity as we suggest, it is expected that the relative importance of priors and other shape cues will increase for locations in which disparity is unreliable (3, 5, 31, 32). To test this idea experimentally, we made a simple manipulation of swapping the two eyes' views of our stimuli to pit monocular and binocular cues against each other. This manipulation inverts the disparity-defined depth ordering of all points in the image (i.e., turning a convex surface into a concave one), while keeping all other aspects of the display identical. We contrasted observers' 3D shape judgments for painted and specular potatoes using occluding masks to show the reliable and unreliable portions of these shapes. We found that reversing the binocularly specified shape reversed observers' 3D shape perception when painted objects were displayed or when reliable portions of the specular potatoes were visible (Fig. S4). By contrast, reversing the two eyes' views had no effect on the perceived depth structure for the unreliable portions of the shapes; instead, observers' judgments were consistent with the 3D shape specified by monocular shape cues in conjunction with a convexity prior (33, 34). Thus, consistent with the use of intrinsic disparity markers, in locations where disparity signals are less reliable, observers' judgments of shape rely more on other sources of information about 3D shape.

It is important to note that in computing the virtual surface, we made the simplifying assumption that objects reflect an environment at optical infinity, which clearly would not hold in the real world. One consequence of this is that the disparities created by reflected features do not depend solely on the curvature of the surface, but also on the distance of the reflected features from the surface. We tested the effects of illumination distance and found that disparities are dominated by the object's surface curvatures, with the distance of reflected features playing a minor role, even for surfaces with only shallow curvature (Fig. S5). Nevertheless, when reflected scene elements occlude one another, unmatchable features (i.e., Da Vinci occlusion) can occur, with the potential to introduce horizontal and vertical matching offsets. Whereas such discontinuities are encountered in everyday viewing, they complicate the virtual surfaces we compute by introducing additional discontinuities (e.g., the virtual surfaces of our muffin objects would not be smooth, due to additional discontinuities imposed by the scene's 3D structure). How the visual system distinguishes

unmatchable features that are due to occlusion from those that are due to specular reflection is an important unsolved problem. Further, it is still unclear whether it is possible, even in principle, to fully and uniquely infer 3D surface locations from specular disparities. To date, computational work (2, 35–37) suggests that specular disparities provide constraints on surface structure but do not necessarily specify shape uniquely. Interestingly, monocular cues based on compression (11) also provide only constraints on shape, but the constraints are different. A promising topic for future research is whether the intersection of monocular and binocular constraints can be used to uniquely identify surface structure from specular reflections.

Conclusions

Together, our findings suggest that a single general strategy can account for the way the brain handles disparity signals arising from specular surfaces in a wide variety of contexts. The simple, convex stimuli in our experiments do not contain extensive unmatchable regions or large vertical disparities at unexpected locations. Therefore, although the surfaces are clearly specular, the disparity signals themselves do not contain the intrinsic indicators that they are unreliable, causing the visual system to interpret them at face value (and thus mistake the virtual surface for a true surface). This also occurs for portions of the more complex objects, where the disparity signals are reliable.

In contrast, where features are outliers in terms of either horizontal or vertical disparities, this indicates to the brain that the disparity signals are unreliable. In the limit, some regions become unfusable and disparity signals are lost completely. In response, to estimate 3D shape, the visual system relies more on monocular cues or spatially interpolates the estimates of depth from more reliable disparity signals at other locations across the surface. This allows the brain to reject portions of the virtual surface, resulting in estimates that sometimes lie closer to the true surface. Thus, rather than knowing the physics of specular reflection, the brain likely interprets specular objects by applying a general robust strategy that would be useful whenever disparity signals behave abnormally, whether or not the origin of those signals is a specular surface.

The findings also have more general implications for the coding of sensory signals. It is common to think of the reliability of sensory signals as depending primarily on their noise or variance (3–6). However, here we have shown that other aspects of the signals (in this case values that are outside the expected range in certain dimensions) can also play a role.

Experimental Procedures

The five subjects had normal or corrected-to-normal visual acuity and normal stereo vision; one was author A.A.M., others were naïve to the study. They provided written informed consent in line with the ethical approval granted to the study by the University of Birmingham Science, Technology, Engineering and Mathematics ethics committee. Participants viewed stimuli using a dual-display (ViewSonic FB2100X) mirror stereoscope. Viewing distance was 50 cm. Stimulus presentation was controlled by a computer with an NVIDIA Quadro FX4400 graphics card. Screen resolution was 1,600 × 1,200 pixels at 100 Hz. The two displays were matched and linearized using photometric measurements. Head movements were restricted using a chin rest.

Stimuli were created and rendered in Matlab (The MathWorks, Inc.). We used two sets of objects: simple muffins and complex potatoes. Muffins were created by distorting spheres (radius, $R = 3$ cm) with a sinusoidal wave whose period and amplitude were varied. Period was defined in terms of the number of cycles of the wave within the sphere ($n = 2, 3, 5, \text{ or } 9$), and is intuitively understood in terms of the number of "corners" the object has. Amplitude was varied for each n -cornered muffin so that the resulting object was everywhere convex ($\alpha = 1/8, 1/15, 1/60, \text{ or } 1/220$). There were three rotations of the muffins with respect to the viewer (denoted by ϕ_0). The Cartesian profile of the objects was defined in terms of spherical functions of elevation (θ) and azimuth (ϕ), where

$$\begin{aligned}
X(\theta, \varphi) &= R(1 + \alpha \sin(n\varphi - \varphi_0)) \times \cos \varphi \times \sin \theta \\
Y(\theta, \varphi) &= R(1 + \alpha \sin(n\varphi - \varphi_0)) \times \sin \varphi \times \sin \theta \\
Z(\theta, \varphi) &= R \times \cos \theta.
\end{aligned}$$

Potatoes were spherical functions created by distorting the sphere with a number ($n = 20\text{--}100$) of symmetric and normalized Gaussian bumps ($\sigma_\theta = \pi/12$, $\sigma_\varphi = \pi(12 \sin \theta)$):

$$R_{\text{pot}}(\theta, \varphi) = R_{\text{sphere}} + \sum_{k=1}^{N_{\text{bumps}}} \frac{0.2}{2\pi\sigma_{\theta k}^2} e^{-\frac{(\theta - \theta_k)^2}{2\sigma_{\theta k}^2} - \frac{(\varphi - \varphi_k)^2}{2\sigma_{\varphi k}^2}}.$$

Bump locations were selected randomly across the surface of the sphere (specified by φ_k and θ_k), with the effect that surface had regions of local convexity and local concavity. However, convexities dominated, as is typical for most natural objects. Objects were scaled so that the maximum radius did not exceed 3.5 cm.

Stimuli were rendered in Matlab under natural illumination [Eucalyptus light probe (38)] where the illumination was treated as arriving from infinite distance (7, 39). Specifically, we mapped the illumination map onto the surface of the object using the surface normal vectors and the physical law of specular reflection. Specular stimuli were rendered such that the reflected ray vectors arrive at left and right eyes (i.e., texture maps differed for left and right eyes). For painted stimuli, the reflection process was modeled with respect to the cyclopean point, so the same texture map was applied for left and right views, so that it appeared that the illumination was painted on the surface of the object. Following the mapping process, the objects were rendered using off-axis stereoscopic projection.

Disparity Properties of the Virtual Surface. To compute specular disparities, we traced view rays from each eye to the surface and calculated the reflected rays that point into the environment. Corresponding locations on the surface are those whose reflected rays point at the same location in the world in the left and right eye (i.e., parallel rays, assuming illumination at infinity). For specular surfaces, the view vectors that point at corresponding surface locations from the left and right eye generally do not intersect ("skew rays"). In this case, we impose a match by projecting the two rays into a plane that contains the two eyes and the center of the object, where they do intersect.

We project this intersection point back onto the left and right eyes' view rays and define the virtual surface point as the average of these two positions in 3D space. Thus, correspondence is defined using only the horizontal (epipolar) disparity component, and the vertical (orthopolar) disparity component remains as a measurable residual (see Fig. S1 for further details).

To define the DDC, we used thresholds for vertical disparities of 12 arcmin (20, 21) and thresholds of horizontal disparity gradients of one (19). A disparity was identified as unreliable on the basis of passing either threshold. Interior regions that were flanked on both sides by portions that were excluded by the DDC were treated as outliers and also removed due to their small size (in practice, this tended to remove the small concave portions of the virtual surface in front of the true surface). We considered the use of constraints based on discontinuities in vertical disparity signals (29); however, abrupt changes of vertical disparities in our stimuli were typically bounded by holes in the virtual surface making this criterion unreliable. To interpolate over unavailable or unreliable regions, we applied Bezier curves to down-sampled (to avoid overfitting) data from reliable portions of the virtual surface.

Psychophysical Procedure. In the first experiment, we tested geometrically simple objects (muffins). There were four shapes viewed in three orientations, and two surface material conditions per shape (perfectly Lambertian and perfectly specular), i.e., 24 conditions per subject. Subjects ($n = 5$) indicated perceived depth using a probe adjustment task in which they controlled the depth of 11 probe dots along the horizontal midline of the stimulus such that the probes appeared to lie on the surface of the object (initial depths of the probes were randomized). In the second experiment we used irregular, nonconvex 3D objects (potatoes). There were four shapes and two surface material conditions (perfectly Lambertian and perfectly specular). We used the same probe adjustment task, but this time there were three rasters (11 points each) of probes per shape (one raster was placed along the horizontal midline of the shape, the other two were shifted 1.3 cm up and down in the frontal plane forming a regular grid of 33 probes).

ACKNOWLEDGMENTS. We thank K. Doerschner, K. Gegenfurtner, and Z. Kourtzi for their comments on the manuscript. This research was supported by Grants 08459/Z/07/Z and 095183/Z/10/Z from the Wellcome Trust and the joint National Science Foundation–Federal Ministry of Education and Research's Bernstein Program for Computational Neuroscience (FKZ: 01GQ1111).

- Blake A, Bülthoff HH (1990) Does the brain know the physics of specular reflection? *Nature* 343(6254):165–168.
- Blake A, Bülthoff HH (1991) Shape from specularities: Computation and psychophysics. *Philos Trans R Soc Lond B Biol Sci* 331(1260):237–252.
- Kersten D, Mamassian P, Yuille A (2004) Object perception as Bayesian inference. *Annu Rev Psychol* 55:271–304.
- Landy MS, Maloney LT, Johnston EB, Young M (1995) Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Res* 35(3):389–412.
- Knill DC, Pouget A (2004) The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends Neurosci* 27(12):712–719.
- Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415(6870):429–433.
- Fleming RW, Dror RO, Adelson EH (2003) Real-world illumination and the perception of surface reflectance properties. *J Vis* 3(5):347–368.
- Motoyoshi I, Nishida S, Sharan L, Adelson EH (2007) Image statistics and the perception of surface qualities. *Nature* 447(7141):206–209.
- Kim J, Marlow PJ, Anderson BL (2012) The dark side of gloss. *Nat Neurosci* 15(11):1590–1595.
- Beck J, Prazdny S (1981) Highlights and the perception of glossiness. *Percept Psychophys* 30(4):407–410.
- Fleming RW, Torralba A, Adelson EH (2004) Specular reflections and the perception of shape. *J Vis* 4(9):798–820.
- Doerschner K, et al. (2011) Visual motion and the perception of surface material. *Curr Biol* 21(23):2010–2016.
- Nishida S, Motoyoshi I, Maruya K (2011) Luminance-color interactions in surface gloss perception. *J Vis* 11(11):397.
- Anderson BL, Winawer J (2005) Image segmentation and lightness perception. *Nature* 434(7029):79–83.
- Bloj MG, Kersten D, Hurlbert AC (1999) Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature* 402(6764):877–879.
- Gilchrist AL (1977) Perceived lightness depends on perceived spatial arrangement. *Science* 195(4274):185–187.
- Mayhew JE, Longuet-Higgins HC (1982) A computational model of binocular depth perception. *Nature* 297(5865):376–378.
- Tyler CW (1973) Stereoscopic vision: Cortical limitations and a disparity scaling effect. *Science* 181(4096):276–278.
- Burt P, Julesz B (1980) A disparity gradient limit for binocular fusion. *Science* 208(4444):615–617.
- Qin D, Takamatsu M, Nakashima Y (2006) Disparity limit for binocular fusion in fovea. *Opt Rev* 13(1):34–38.
- van Ee R, Schor CM (2000) Unconstrained stereoscopic matching of lines. *Vision Res* 40(2):151–162.
- Marr D, Poggio T (1976) Cooperative computation of stereo disparity. *Science* 194(4262):283–287.
- Todd JT, Mingolla E (1983) Perception of surface curvature and direction of illumination from patterns of shading. *J Exp Psychol Hum Percept Perform* 9(4):583–595.
- Norman JF, Todd JT, Phillips F (1995) The perception of surface orientation from multiple sources of optical information. *Percept Psychophys* 57(5):629–636.
- Todd JT, Norman JF, Koenderink JJ, Kappers AML (1997) Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception* 26(7):807–822.
- Norman JF, Todd JT, Orban GA (2004) Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychol Sci* 15(8):565–570.
- Mingolla E, Todd JT (1986) Perception of solid shape from shading. *Biol Cybern* 53(3):137–151.
- Nefs HT (2008) Three-dimensional object shape from shading and contour disparities. *J Vis* 8(11):11–16.
- Serrano-Pedraza I, Phillipson GP, Read JCA (2010) A specialization for vertical disparity discontinuities. *J Vis* 10(3):1–25.
- Koenderink JJ (1984) What does the occluding contour tell us about solid shape? *Perception* 13(3):321–330.
- Mamassian P, Landy MS (2001) Interaction of visual prior constraints. *Vision Res* 41(20):2653–2668.
- Welchman AE, Lam JM, Bülthoff HH (2008) Bayesian motion estimation accounts for a surprising bias in 3D vision. *Proc Natl Acad Sci USA* 105(33):12087–12092.
- Hill H, Bruce V (1993) Independent effects of lighting, orientation, and stereopsis on the hollow-face illusion. *Perception* 22(8):887–897.
- Langer MS, Bülthoff HH (2001) A prior for global convexity in local shape-from-shading. *Perception* 30(4):403–410.
- Savarese S, Perona P (2002) Local analysis for 3D reconstruction of specular surfaces - Part II. *Lect Notes Comput Sci* 2351:759–774.
- Bhat DN, Nayar SK (1998) Stereo and specular reflection. *Int J Comput Vis* 26(2):91–106.
- Healey G, Binford TO (1988) Local shape from specularity. *Comput Vision Graph* 42(1):62–86.
- Debevec PE (1998) Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Proceedings of SIGGRAPH* 1998:189–198.
- Dror RO (2002) *Surface Reflectance Recognition and Real-World Illumination Statistics* (MIT Artificial Intelligence Laboratory, Cambridge, MA).