

Does the brain know the physics of specular reflection?

Andrew Blake* & Heinrich Bülthoff†‡

* Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK

† Center for Biological Information Processing, Department of Brain and Cognitive Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

‡ Present address: Department of Cognitive and Linguistic Sciences, Brown University, Box 1978, Providence, Rhode Island 02912, USA

IMAGES of artificial and natural scenes typically contain many highlights generated by mirror-like reflection from glossy surfaces. Until recently, computational models of visual processes have tended to regard highlights as obscuring the structure of the underlying scene. The truth is that, on the contrary, highlights are rich in local geometric information. Here we report that the three-dimensional appearance of a highlight on a computer-simulated stereoscopic curved surface affects observers' judgment of surface gloss. We also show that the 3-D appearance of a highlight affects the perception of surface curvature—that is, it can force an ambiguous convex-concave figure to change state. We thus conclude that human visual analysis seems to employ a physical model of the interaction of light with curved surfaces, a model firmly based on ray optics and differential geometry.

According to ray optics, the virtual image of a light source—a specularity—is located in depth behind a glossy convex surface and, generally, in front of a concave surface (Fig. 1). But what the eye actually registers is not depth but 'relative disparity' (Fig. 2). The ray-optic 'specular stereo' model^{1,2} (Fig. 2) shows that the relative disparity of a specularity and the shape of the underlying surface are directly related. Suppose that the specular stereo model were fully utilized by the visual system: then the

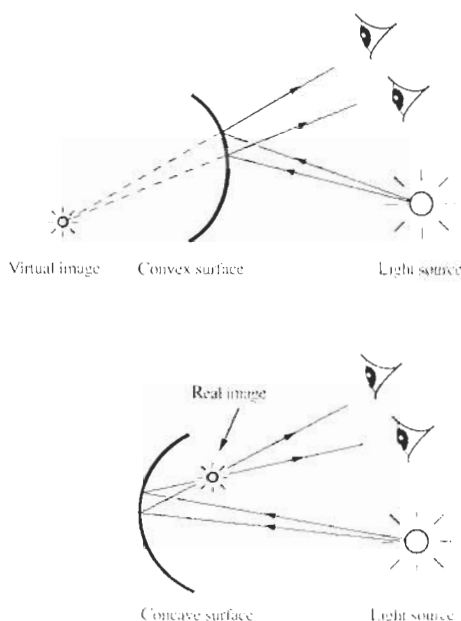


FIG. 1 The basic principle of the specular stereo model: unlike surface points such as scratches, specularities are not constrained to lie on surfaces. They appear behind a convex reflective surface but (generally) in front of a concave one. In fact things are more complicated than that: if the glossy surface is hyperbolic, specularities can appear either behind or in front of the surface^{2,10}. Generally their apparent depth varies as the surface is rotated about the line of sight. Moreover, specularities do not obey the 'epipolar' constraint¹¹ of stereoscopic vision, an important feature of the specular stereo model¹ which further complicates the idea of perceived depth¹².

relative disparity of a specularity would be consistent only with certain values of local surface curvature, if the position of the light source was known. Even without knowledge of the light-source position, the relative disparity of a specularity still constrains curvature. No convex surface can generate a convergent (–) relative disparity; a concave surface does not generate a divergent (+) one (except under certain conditions¹, unusual in practice). We performed experiments aimed to test whether the human visual system exploits such constraints.

An adjustment task was devised in which the subject interactively varied the disparity of a specularity to maximize the perceived glossiness of a stereoscopically displayed surface. Images of glossy textured curved surfaces were generated with a computer graphics workstation (Symbolics) and displayed on a high-resolution colour monitor with a stereo viewing system (Stereo-optic)³. The texture is of sufficient density to furnish strong cues for curvature from edge-based stereo. Simulation of surface gloss⁴ causes a specularity to appear superimposed on the texture, as in Fig. 3a. When the specular relative disparities are consistent with the geometry of the surface that is visible to edge-based stereo, the whole surface (not just the vicinity of the specularity⁵) appears glossy as in Fig. 3a. However, when relative disparity of a simulated specularity is inconsistent the surface ceases to look glossy (Fig. 3b).

Five test surfaces were used in the adjustment task—a convex sphere, two convex ellipsoids and two concave ellipsoids. In each case, naive subjects were asked to adjust the disparity of the simulated specularity so as to achieve the most realistic-looking glossy surface.

Results for seven subjects with the convex sphere are given in Fig. 3c. They show good agreement with the model. If the

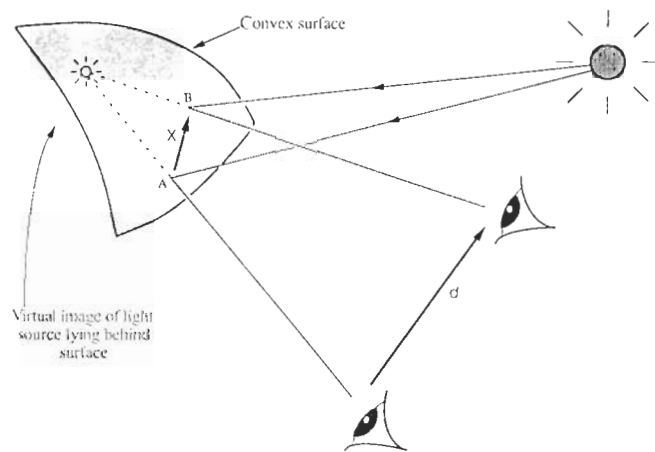


FIG. 2 Ray optics establishes a direct relationship between surface shape and the stereoscopic 'relative disparity' of a specularity^{1,2,12,13}. Disparity is the difference between angular positions of a particular feature in left and right stereoscopic views. The relative disparity of a specularity is the difference between its disparity and that of a 'surface' point lying along a common line of sight. It is a projection onto image coordinates of the displacement vector x between the points of reflection A and B at which light rays strike the surface. The displacement vector x , in turn, is related linearly to baseline vector d , the coefficients of the relation being a function solely of surface and viewer geometry². If the visual system 'knows' the physics of specular reflection and the approximate light source position, then the observed relative disparity of a specularity can be consistent only with certain values of local surface curvature. If, for example, the surface is strongly curved, then for a given baseline d , the displacement x will be small, whereas for a plane surface x tends to be larger. The relative directions of x and d also depend on surface geometry. For a specularity on a curved surface, with given surface orientation and baseline d , displacement x is in one direction if the surface is convex but in exactly the opposite direction if the surface is reversed to become concave.


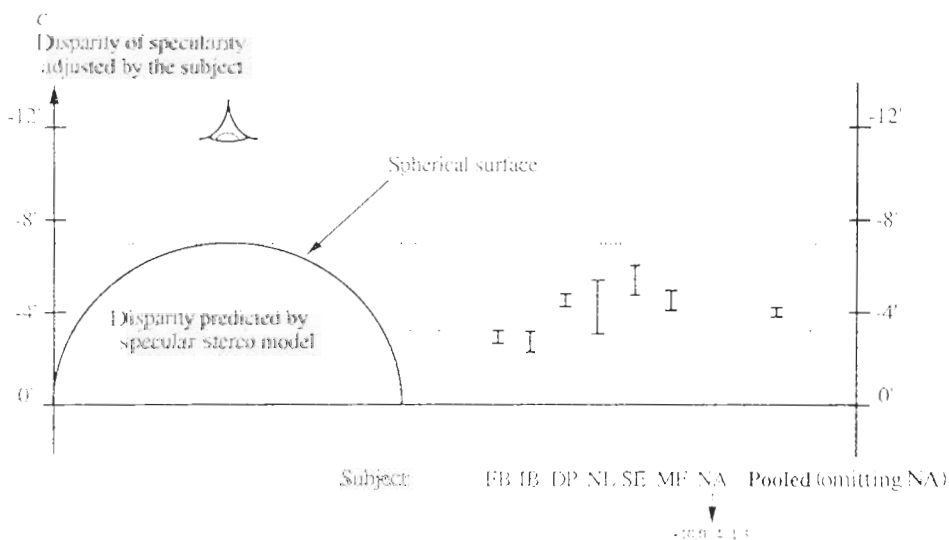


FIG. 3 The perception of surface properties can be changed by moving a simulated specularity, in depth, relative to the surface. The surface of the sphere (a) (uncrossed stereo view) looks glossy because the highlight is in the correct position behind the surface. If the highlight is in front of the surface (b)—excessively convergent (–) relative disparity—surface quality is reported to be matt and opaque, with a puff of cloud in front of it. With excessively divergent (+) relative disparity, subjects usually report that the surface looks transparent, with a light source behind it. When the relative disparity is zero the simulated specularity looks like a powdery surface patch. In an informal 2AFC experiment, 11 out of 12 naive observers who were asked which of *a* and *b* was the “polished” surface chose *a*. Seven naive observers performed an adjustment task in which they varied the relative disparity of a simulated highlight on a spherical surface. *c*. Results, giving mean and s.e.m. of responses, show a considerable degree of consistency with the specular stereo model. Four subjects (F.B., I.B., N.L., S.E.) made adjustments consistent with the model (99% significance). Two others (D.P., M.F.) showed a small bias towards zero disparity, that is, the specularity is displaced slightly towards the surface. The remaining subject (N.A.) seemed to be performing a different task, placing the specularity very far behind the surface. The pooled results show a small but significant (99%) bias of about 0.5' towards zero disparity. Pooled results show that subjects do not, on average, place the specularity on the surface.



visual judgements were not made according to the specular stereo model, it might be expected that subjects would attempt to place the simulated specularity on the surface. That hypothesis fails at the 99% significance level. The sign of the relative disparity of the specularity after adjustment is also correct—subjects placed the specularity behind, not in front of, the surface. Pooled results (Fig. 3c) show just a small bias towards zero disparity. It amounts to a reduction of 10–25% (99% confidence interval) in the adjusted depth of the specularity behind the surface, compared with the relative depth predicted by the model. Results are almost as good for the other two

convex ellipsoids. The sign of the relative disparity was correct for all but one out of six trials (three subjects, two different ellipsoids). Subjects did not place the specularity on the surface, and the sign of relative disparity was correct. Similar results could not be obtained for concave surfaces, however. Two concave surfaces were used, exactly the inverses of the convex ellipsoids used previously. Out of six different naive subjects, four placed the specularity on the surface and two behind it. The model predicts, of course, that the specularity should be in front of the surface. One possible explanation for the failure of the model to predict performance for concave surfaces is as


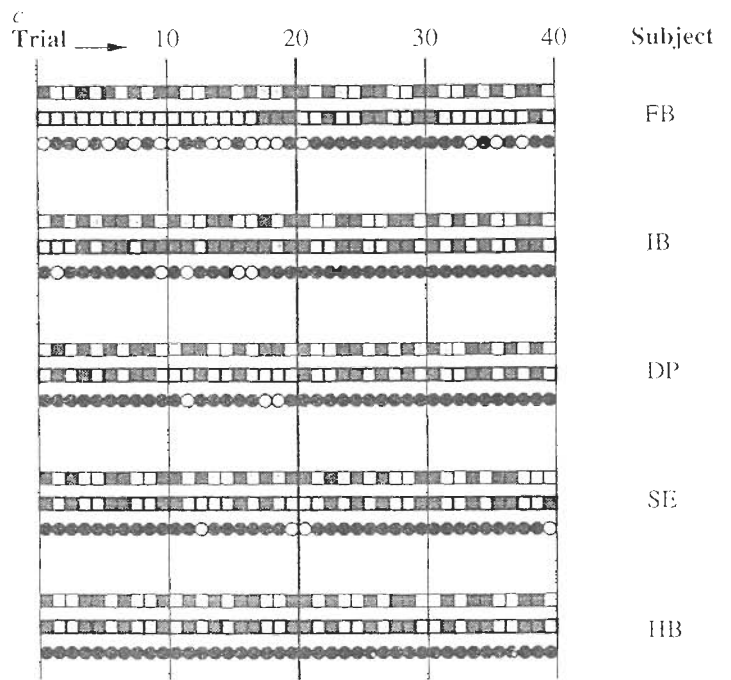


FIG. 4. The perception of surface curvature can be affected by the position of a specular highlight. As a demonstration that the human visual system 'knows' the physics of ray optics we used an image of a surface whose 3-D interpretation can flip between two states (convex, concave). If a highlight is added to the image the 3-D interpretation, in a monocular view, is biased so that the inner part of the surface tends to be convex (like a ball in a saucer). A simulated specularly superimposed on a stereoscopic view of the surface can be either behind (a) or in front of (b) the textured surface (uncrossed stereo view). The textured surface itself is at zero disparity. The two views were presented in random order, with a stereo viewing system, using 5-, 10- or 15-s exposures separated by random-dot masking frame. Subjects made a 2AFC between a convex and concave interpretation. After a short training period (20 exposures without feedback) they predominantly made choices which conformed to the predictions of the model (c). For each of five subjects (four were naïve about the purpose of the experiment), the last 20 responses were biased towards those predicted by the model, at the 99% significance level. (The effect is weakened in the printed images because of the limited contrast range of the print-medium, compared with that of the CRT monitor used for the experiment: it is essential, for the strongest effect, that the highlight looks like a real reflection of a light source.



Specularity position: behind □ : in front ■
 Subject's response: convex □ : concave ■
 Specular stereo prediction: right ● : wrong ○

follows. Currently available graphics tools do not simulate the effects of shadowing by extended light-sources or of mutual illumination⁶, both of which are crucial for faithful rendering of concave surfaces. Neglect of these effects produces surfaces which are concave stereoscopically but, monocularly, tend to appear convex—a cue conflict³. Hence subjects reported that it was difficult to find any adjustment for which a concave surface looked convincingly glossy.

The conclusion of this experiment is that, at least for simulated convex surfaces, the human visual system models the physics of specular reflection well enough to predict relative disparity effects. The visual system 'expects'—correctly—that a specularity lies behind, not on or in front of, a convex surface. Note that results reported here are for variation of horizontal disparity. Subjects also adjust vertical disparities⁷ on convex surfaces, in accordance with the specular stereo model⁸.

Our second experiment was complementary to our first and addressed the question of whether the visual system can accommodate variations in specular relative disparities by changing its hypothesis of surface curvature, rather than its hypothesis of glossiness.

We devised the stimulus shown in Fig. 4—a stereo textured variant of an ambiguous (reversible) shaded surface⁹. The texture elements all have zero disparity, consistent with a fronto-parallel planar surface. Nonetheless, monocular shading/texture cues are not entirely overridden³, so that subjects could usually see both convex (like a ball in a saucer) and concave (like a dog bowl) interpretations. A superimposed specularity, with either convergent or divergent relative disparity, strongly influenced the interpretation. As the specular stereo model predicts, convergent relative disparity of a specularity biases the subjects' interpretation away from convex. Similarly, divergent relative disparity biases interpretations away from concave. The prediction was tested by a two-alternative forced choice (2AFC) between a convex or a concave surface interpretation, when subjects were presented with simulated specularities, divergent and convergent ($\pm 5^\circ$), in random sequence. Note that subjects were asked whether the surface appeared convex or concave, not whether the specularity was behind or in front of the surface. Figure 4c shows how the surface inter-

pretation predicted by the model develops gradually with repeated exposures. Although initially subjects tended to be locked into either a convex or a concave interpretation, after about 20 exposures they were usually picking, quite reliably, the interpretation that was consistent with the sign of relative disparity of the specularity. Note that the change in position of the specularity is contrary to that of the surface—when the specularity is furthest away (divergent relative disparity) the centre of the surface is nearest to the viewer (convex), and vice versa. Therefore, any explanation in terms of a pulling effect exerted by the specularity on the surface is excluded.

Thus, naive observers, asked where a specularity appears to be in relation to the surface that generated it, usually reply that it seems to lie on the surface. The first experiment that we have described attempted to show that the early visual system 'knows' better, choosing interpretations that are broadly consistent with the physics of specular reflection. The second experiment demonstrated that the stereoscopic appearance of a specularity influences perceived curvature, again in a way that is consistent with the specular stereo model. This influence was detectable despite the presence of other strong shape cues. Exactly how these cues, often in mutual conflict, are combined³ and resolved into a stable percept is unknown. □

Received 18 May; accepted 17 October 1989.

1. Blake, A. *Proc. 9th Int. Joint Conf. on Artif. Intell.* 973-976 (1985).
2. Blake, A. & Breilstaff, G. J. *Proc. 2nd Int. Conf. comp. Vis.* 394-403 (Institute of Electrical and Electronics Engineers, Washington DC, 1988).
3. Bülthoff, H. H. & Mallot, H. A. *J. opt. Soc. Am.* **5**, 1749-1758 (1988).
4. Phong, B. T. *Commun. Ass. comput. Mach.* **18**, 311-317 (1975).
5. Beck, J. *Surface Color Perception* (Cornell University Press, New York, 1972).
6. Greenberg, D. P. *Science* **244**, 166-173 (1989).
7. Mayhew, J. E. W. & Longuet-Higgins, H. C. *Nature* **297**, 376-379 (1982).
8. Blake, A. & Bülthoff, H. H. *Artificial Intelligence Lab Memo* 1149 (Massachusetts Institute of Technology, Cambridge, Massachusetts, 1989).
9. Ramachandran, V. *Nature* **331**, 163-166 (1988).
10. Koenderink, J. J. & van Doorn, A. J. *Optica Acta* **27**, 7, 981-996 (1976).
11. Mayhew, J. E. W. & Frisby, J. P. *Artif. Intel. J.* **17**, 349-385 (1981).
12. Koenderink, J. J. & van Doorn, A. J. *Biol. Cybern.* **21**, 29-35 (1976).
13. Zisserman, A., Giblin, P. & Blake, A. *Image Vis. Comput.* **7**, 38-42 (1989).

ACKNOWLEDGEMENTS. We acknowledge the support of a grant from the Office of Naval Research, Cognitive and Neural Sciences Division (H.H.B.) and the Fellowship of Engineering, Exeter College, Oxford, University of Oxford and the Science and Engineering Research Council (A.B.). We thank N. Gryzwacz, D. Kersten, D. Plummer, T. Poggio and S. Ulfman for discussions.