Vision Research 115 (2015) 199–208

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Material properties derived from three-dimensional shape representations

Phillip J. Marlow, Barton L. Anderson

Department of Psychology, Griffith Taylor (A19), University of Sydney, Sydney, NSW 2006, Australia

ARTICLE INFO

Article history:
Received 30 April 2015
Received in revised form 4 May 2015
Available online 14 May 2015

Keywords:
Material perception
Shading
Gloss
Three-dimensional shape
Inverse optics
Binocular vision

ABSTRACT

Retinal image structure is due to a complex mixture of physical sources that includes the surface's 3D shape, light-reflectance and transmittance properties, and the light field. The visual system can somehow discriminate between these different sources of image structure and recover information about the objects and surfaces in the scene. There has been significant debate about the nature of the representations that are used to derive surface reflectance properties such as specularity (gloss). Specularity could be derived either directly from 2D image properties or by exploiting information that can only be derived from representations in which 3D shape has been made explicit. We recently provided evidence that 3D shape information can play a critical role in the perception of material specularity, but the shape manipulation in our prior study also significantly changed 2D image properties (Marlow, Todorović, & Anderson, 2015). Here, we held fixed all monocularly visible 2D image properties and manipulated 3D shape stereoscopically. When binocularly fused, the depicted 3D shapes induced striking transformations in the surfaces' apparent material properties, which vary from matte to 'metallic'. Our psychophysical measurements of perceived specularity reveal that 3D shape information can play a critical role in material perception for both singly-curved surfaces and more complex geometries that curve in two directions. These results provide strong evidence that the perception of material specularity can depend on physical constraints derived from representations in which three-dimensional shape has been made explicit.

1. Introduction

Retinal image structure arises from the interaction between a surface's three-dimensional (3D) shape, its reflectance and transmittance properties, and the surrounding light field. Perceptual experience reveals that our visual system somehow extracts these distinct contributions to image structure, but there is no consensus about how this computational feat is accomplished. Any local image structure can be produced by an infinite number of different combinations of shape, reflectance, and illumination, which suggests that some additional information is required to determine the particular combination responsible for a given image. The inability to compute a unique inverse suggests that the visual system solves this problem probabilistically, exploiting constraints on the likelihood of the possible scene interpretations. One of the main theoretical and empirical tasks is to identify the regularities that the visual system uses to infer scene structure, and to characterize the representational space over which these constraints are defined.

The majority of work has focused on regularities that can be computed directly from 2D images, which have been applied in a variety of different domains. Some early work used image statistics to theoretically motivate efficient coding schemes of subcortical and early cortical areas (Field, 1987; Olshausen & Field, 1996; Srinivasan, Laughlin, & Dubs, 1982). Image statistics have also been used to understand a variety of mid-level visual processes, such as contour completion (Geisler & Perry, 2009; Geisler, Perry, Super & Gallogly, 2001), and the computation of surface reflectance (Fleming & Bülthoff, 2005; Giesel & Zaidi, 2013; Liu et al., 2010), such as lightness (Motoyoshi et al., 2007; Sharan et al., 2008) and gloss (Arce-Lopera et al., 2012; DelPozo & Savarese, 2007; Fleming, Dror, & Adelson, 2003; Motoyoshi & Matoba, 2012; Motoyoshi et al., 2007; Nishida & Shinya, 1998). However, it has been shown that the perception of gloss depends on the spatial organization of specular image structure, which cannot be derived from image statistics that fail to capture that organization (Anderson & Kim, 2009; Kim & Anderson, 2010). In particular, we have previously argued that perceived gloss depends on two photogeometric

E-mail addresses: phillip.marlow@sydney.edu.au (P.J. Marlow), barton.anderson@sydney.edu.au (B.L. Anderson)

http://dx.doi.org/10.1016/j.visres.2015.05.003
0042-6989/

constraints: *Orientation congruence* and *brightness congruence* (Kim, Marlow, & Anderson, 2011; Marlow, Kim, & Anderson, 2011). *Orientation congruence* refers to the fact that local orientations of specular reflections tend to run parallel to local orientations of diffuse shading gradients (Anderson & Kim, 2009; Beck & Prazdny, 1981; Kim, Marlow, & Anderson, 2011, 2012; Marlow, Kim, & Anderson, 2011). *Brightness congruence* refers to the fact that specular highlights are typically located close to the brightest regions of diffuse shading (Marlow, Kim, & Anderson, 2011). The luminance maxima on a glossy surface need to satisfy these constraints in order to appear as specular reflections; highlights that violate these constraints appear as light pigment on a matte surface or a disconnected overlay (e.g., Beck & Prazdny, 1981; Todd, Normal, & Mingolla, 2004). In principle, both of these constraints could be computed directly from images. For example, orientation congruency may be derived from orientationally selective image filters (or relatedly orientation fields), and brightness congruency may be derived from the position of a highlight relative to the luminance maxima and minima of the surrounding luminance gradient.

The photo-geometric constraints described above provide some insight into how the visual system identifies specular reflections, but they do not explain how the perception of gloss can vary between surfaces when both constraints are satisfied (Doerschner, Boyaci, & Maloney, 2010; Fleming, Dror, & Adelson, 2003; Ho, Landy, & Maloney, 2008; Obein, Knoblauch, & Viénot, 2004; Olkkonen & Brainard, 2010, 2011; Pont & te Pas, 2006; te Pas & Pont, 2005; Vangorp, Laurijssen, & Dutré, 2007; Wendt et al., 2010; Wijntjes & Pont, 2008). We recently argued that perceived gloss varies as a function of image cues that are predictive of a surface's gloss level (Marlow & Anderson, 2013; Marlow, Kim, & Anderson, 2012). In particular, we showed that perceived gloss is modulated by the contrast, sharpness, and 'coverage' of specular image structure. Contrast refers to the difference in luminance between a specular reflection and its surround; sharpness refers to the slope of the luminance gradient at the edge of a specular reflection; and coverage refers to the proportion of a surface that generates visible specular reflections. High levels of physical gloss typically generate higher levels of specular contrast, sharpness, and coverage than do low levels of gloss (Berzhanskaya et al., 2005; Billmeyer & O'Donnell, 1987; Hunter & Harold, 1987; Pellacini, Ferwerda, & Greenberg, 2000), but these image properties can also vary dramatically as a function of a surface's 3D shape or the light field. We have shown that psychophysical measurements of the apparent contrast, sharpness, and coverage of specular reflections can predict how perceived gloss scales across a wide range of 3D shapes, light fields, and physical gloss levels (Marlow & Anderson, 2013; Marlow, Kim, & Anderson, 2012). Similar correlations have been found when the psychophysical measurements are substituted with computational measurements derived directly from the image (Marlow, Kim, & Anderson, 2012; Qi et al., 2014).

The preceding theories suggest that our experience of surface gloss could theoretically be derived from differences in image structure, prior to an explicit representation of a surface's 3D geometry or the light field in which it is embedded. However, the perception of gloss is a property associated with surfaces and materials, and is therefore always accompanied by an experience of 3D shape. Although many studies have speculated that 3D shape representations may play a causal role in the computation of gloss, their data do not provide conclusive evidence in support of this view (e.g., Anderson & Kim, 2009; Beck & Prazdny, 1981; Fleming, Dror, & Adelson, 2003; Ho, Landy, & Maloney, 2008; Kim, Marlow, & Anderson, 2011; Marlow, Kim, & Anderson, 2011; Motoyoshi et al., 2007; Nishida & Shinya, 1998). Studies have shown that perceived gloss depends on the stereoscopic depth of specular reflections, which typically appear behind convex surfaces (Blake & Bülthoff, 1990; Kerrigan & Adams, 2013; Muryy et al., 2013). However, the relevant 3D structure used to derive material properties in these studies is not 3D *shape* representations per se, but rather the difference in perceived depth of the surface's texture and shading relative to the depth of the specular reflections.

In order to assess whether computations of gloss exploit information explicitly derived from 3D representations of shape, image structure must be held constant while the 3D shape associated with that structure is varied. This requires constructing different 3D shapes that generate identical image gradients from two different reflectance functions. Fig. 1A depicts luminance gradients generated by a diffuse (matte) surface and a rough specular surface (such as unpolished metal) that have the same 3D structure embedded in an identical illuminant. The matte surface on the left depicts a Lambertian reflectance function that generates a luminance that varies as a cosine of the angle between the surface normal and the direction of the incident illumination. The luminance projected by the specular surface on the right varies much more rapidly than the Lambertian surface, particularly in the neighborhood of the luminance maximum. The steepness of the specular and diffuse luminance gradients depends on surface roughness parameters, which modulate the 'spread' or 'scatter' of light within the diffuse and specular lobes (Nicodemus, 1965; Oren & Nayar, 1994). If the visual system exploits these three-dimensional constraints to derive reflectance properties, then it should be possible to generate identical image gradients that appear to be associated with different reflectance properties if they are perceived as different three-dimensional shapes.

The two surfaces in Fig. 1B provide our first attempt to test this hypothesis (Marlow, Todorović, & Anderson, 2015). The two figures contain identical luminance gratings bounded by two different sets of bounding contours along the left and right sides of the grating. It has been shown previously that bounding contours affect the perceived 3D shape, lightness, and illumination direction of image gradients (Knill & Kersten, 1991; Ramachandran, 1988; Todorović, 2014; Witkin & Tenenbaum, 1983). In this example, the left surface appears as three large half-cylinders (two convex one concave) illuminated from the front, whereas the right surface appears as three ridges and three valleys illuminated from above. We showed that these two shapes also appear to differ in perceived specular reflectance: The left surface appears matte, whereas the right appears more specular, such as a rough metal (Marlow, Todorović, & Anderson, 2015).

The relationship between perceived shape and perceived reflectance of the images in Fig. 1 suggests that the visual system derives information about reflectance from the rate that luminance gradients vary relative to their 3D surface geometry. However, this is not the only possible interpretation of the perceived material difference of these stimuli; there are image differences that could also account for this result. The reflectance of the surfaces could theoretically be derived from correlations between the luminance gradients and the local orientations of the bounding contours in the 2D images rather than being derived from 3D shape representations. Fig. 1C plots the grating's luminance as a function of the angle of the bounding contour. An angle of zero refers to the angle of the bounding contour adjacent to the luminance maxima in the grating. The graph shows that the rate of change in the luminance of the grating along the bounding contour is slow for the matte surface and rapid for the specular surface. Note that these functions derived directly from the images mimic the functions derived from the matte and specular 3D surfaces shown in Fig. 1A. Thus, the change in perceived material properties in these images could theoretically still be derived from image properties, rather than something computed only once the representation of 3D shape has been made explicit.