Chapter 16
The Statistics of Shape, Reflectance, and Lighting in Real-World Scenes

Richard F. Murray

16.1 Introduction

Visual perception is a statistical problem *par excellence*. If the goal of vision is to give a reliable reconstruction of the scenes and objects in the external world from 2D images on the retinas, then in one sense it is an impossible problem: there is simply not enough information in retinal images alone to infer what is being seen. Put differently, there are many combinations of lighting, surface shapes, and surface colors that could have given rise to any particular 2D retinal image. Accordingly, a visual system can only reconstruct a 3D scene if it has criteria for choosing a particular 3D interpretation of a 2D retinal image out of the wide range of interpretations that are physically consistent with the image.

The generalized bas-relief (GBR) ambiguity illustrates this problem vividly [9]. Suppose we have a Lambertian object, with an arbitrary shape and an arbitrary surface reflectance pattern, under arbitrary lighting. The GBR ambiguity shows that we can drastically change the shape and reflectance pattern of the object, without changing the retinal image that it generates. Specifically, we can compress and shear the object along the viewer's line of sight; then adjust the lighting so that the positions of cast shadows on the object are unchanged; and finally adjust the surface reflectance at each point on the object (now with a new surface normal and new lighting conditions) so that it creates the same image luminance as before. The GBR ambiguity shows constructively that images are not ambiguous just in special cases, or to some small degree, but that they are consistently and deeply ambiguous.

R.F. Murray (✉)
Centre for Vision Research, York University, Toronto, Canada
e-mail: rfm@yorku.ca

S.J. Dickinson, Z. Pizlo (eds.), *Shape Perception in Human and Computer Vision*, Advances in Computer Vision and Pattern Recognition,
DOI 10.1007/978-1-4471-5195-1_16, © Springer-Verlag London 2013

Furthermore, the GBR ambiguity is a lower limit; images are much more ambiguous than the GBR ambiguity alone suggests.[1] The strategy of distorting the shape of an object, and then adjusting lighting conditions and reflectance so that the resulting image is unchanged, is obviously a very general one. We can make many kinds of local, nonlinear deformations in an object's shape, and compensate for their effect on the image by adjusting lighting and reflectance accordingly. At the extreme, we can relight and repaint practically any scene so that it generates the same image as practically any other scene. Thus a visual system cannot simply live with the ambiguity in 2D images. The ambiguity is too great, and without some way of at least partly overcoming it, visual stimuli are all but useless.

The objects created by GBR and GBR-like transformations of real objects are, in some sense, odd. A highly transformed object (e.g., greatly stretched or compressed along the viewer's line of sight) does not usually correspond to our percept of what is shown in the image, and it surprises us that these distorted objects create the same images as the more familiar, untransformed objects. But in what sense are the transformed objects odd? The answer I will explore in this chapter is that there are statistical regularities in the shapes, reflectance patterns, and lighting conditions of real world scenes that allow observers to rule out implausible interpretations, and thereby overcome image ambiguity. On this view, transformed scenes are odd because they have shapes, reflectance patterns, or lighting conditions that are unlikely to occur in the real world.

This answer seems so natural, even obvious, that it is worth pointing out that it is not the only possible answer. The generic viewpoint assumption, for instance, is a reasonable principle that suggests we should prefer 3D interpretations of 2D images that are stable across small changes in viewpoint, instead of interpretations that assume the scene is being viewed from a tightly constrained 'accidental' viewpoint [17]. This principle makes only weak assumptions about shape, reflectance, and lighting (e.g., that lighting is equally likely from all directions), and yet it gives a criterion for preferring some 3D interpretations over others. The generic viewpoint assumption resolves the GBR ambiguity by preferring planar interpretations over interpretations with depth [38], so it is not, by itself, a good way of completely overcoming image ambiguity. Nevertheless, it demonstrates that approaches other than relying on the most obvious scene statistics are possible.

These issues have long been understood in broad terms, and yet remarkably little is known about exactly what statistical regularities in real world scenes can support perception of shape and reflectance, or which of these regularities human observers rely on. Here I selectively review and evaluate recent work on these problems.

---

[1] Belhumeur et al. showed that image ambiguity is limited to the GBR ambiguity for an observer who has images of an object under all possible distant point lighting conditions. This is important for understanding the limits of methods such as photometric stereo, but the ambiguity is much greater when the observer sees an object under just one lighting condition. This has sometimes not been understood, e.g., Todd [36] suggests that work on the GBR ambiguity shows that the ambiguity of 2D images is highly constrained.

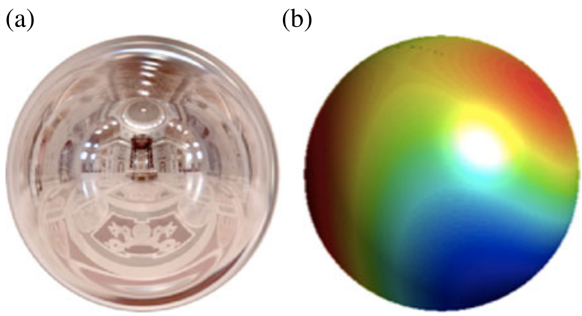

Fig. 16.1 A light probe is a spherical representation of the illumination incident from all directions at a single point in space, i.e., an omnidirectional luminance snapshot. (**a**) A spherical, globe-like representation of a high-resolution light probe assembled from photographs of a mirrored sphere (from [13]). This light probe includes color as well as luminance. (**b**) A similar representation of a low-resolution light probe measured with a multidirectional photometer. *White* represents high luminance, and *black* represents low luminance (from [22])

I will use "scene" to mean a 3D arrangement of surfaces, reflectance patterns, and lights, and "image" to mean the 2D retinal luminance pattern that a scene gives rise to.

16.2 Lighting

16.2.1 Lighting: Scene Statistics

The direction, diffuseness, and complexity of lighting can have an enormous effect on the appearance of a scene. Many studies have examined real world lighting conditions, and have found that despite the great variability in lighting, there are also strong regularities. Most interesting for our purpose are the relatively few studies of 'light probes', omnidirectional snapshots of the pattern of light incident from all directions at a single point in space (Fig. 16.1). A light probe captures the lighting that would illuminate an object at the light probe's measurement location, so an understanding of the statistical regularities in real world light probes is useful for understanding the relationship between 3D scenes and 2D images.

Dror, Willsky, and Adelson [14] examined around 100 high-resolution light probes, and found that they had some of the same statistical properties as conventional images: a pink-noise-like power spectrum, kurtotic wavelet coefficient distributions, and statistical dependencies between wavelet coefficients at adjacent scales, orientations, and positions. They also found that, unlike conventional images, the luminance distribution of light probes peaks at low luminances, with a few very high luminance values due to small, bright sources such as the sun.

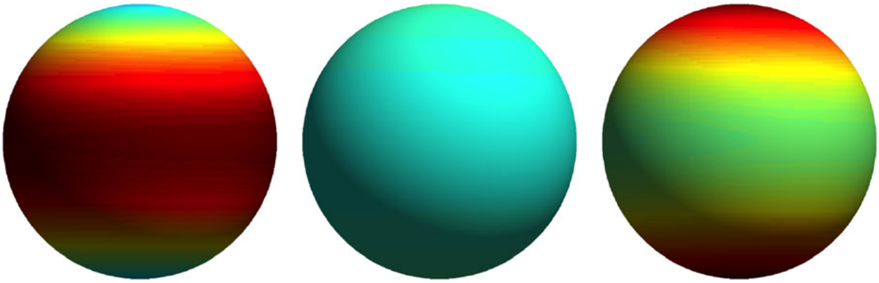

Fig. 16.2 GBR transformations change a diffuse light probe (*center*) so that its luminance is concentrated either along a great circle (*left*) or in two opposite directions (*right*)

Mury, Pont, and Koenderink [28] used a similar approach, but they paid special attention to the low-pass components of natural lighting that are relevant to shading of convex Lambertian objects [8, 32]. They found that although high frequency components of light probes vary rapidly as one moves through a scene, the low frequency components are much more stable. They examined a few different types of scenes, such as open-sky scenes and forests, and showed that the pattern of changes in low-frequency lighting structure through a scene is largely determined by the scene's coarse geometry.

Mury, Pont, and Koenderink [29, 30] built a multidirectional photometer to measure low-pass light probes. Mury et al. [30] measured light probes in several environments, and consistent with their previous work [28], they found that a lighting model based on a coarse description of the scene layout accounted for the structure of the measured light probes.

Morgenstern [22, 23] used a multidirectional photometer to measure several hundred low-pass light probes in diverse environments. He examined the diffuseness of natural lighting, that is, the extent to which light comes mainly from one direction, as on a sunny day, or from all directions, as on a cloudy day. He found that the diffuseness levels in different environments (e.g., sunny, cloudy, indoors), span a fairly limited range, and that across all environments the diffuseness levels cluster in the lower-middle region of the physically possible range of diffuseness.

Morgenstern also found that some low-pass lighting patterns were much more likely than others. He showed that natural low-pass lighting can be approximated reasonably well using a classic computer graphics lighting model, in which light is the sum of a distant point source and uniform ambient source (e.g., [31]). Consistent with this, he found that ring-like lighting patterns, where light is weak in two opposite directions and strong along the great circle halfway between them, are rare. Interestingly, GBR transformations can create just such ring-like lighting patterns (Fig. 16.2): some GBR transformations shift uniform (i.e., very diffuse) light distributions so that they are concentrated either along a great circle of directions, or in two directly opposed directions. This is one sense in which GBR-transformed scenes are unusual, and it gives one possible criterion for choosing the most likely 3D interpretations of 2D images.

16.2.2 Lighting: Psychophysics

Whether an observer's 3D interpretation of an image is accurate often depends on whether they have an accurate estimate of the scene's lighting. For human observers, this estimate is a compromise between the observer's prior on lighting conditions, and cues to lighting conditions in individual scenes. We know more about the human visual system's priors on lighting than about any other perceptual prior.

Metzger [21] first suggested that in order to perceive the 3D shapes depicted in 2D images, we rely on an assumption that light comes from above and slightly to the left, an assumption that has come to be known as the light-from-above prior[2] Metzger based this suggestion on informal observations of the appearance of images at different orientations, and more quantitative psychophysics has supported his notion of a preferred lighting direction above and to the left [35]. (One cannot say more *careful* psychophysics, since it is remarkable that Metzger was able to conclude that we prefer a lighting direction to the left of vertical, based only on his own qualitative observations.) Interestingly, there are large individual differences in the precise direction of the prior [2], and the direction of the prior can be modified by just an hour or two of experience in an environment with oblique lighting [1].

Even if light comes from some direction on average (e.g., directly above, or above and to the left), it does not come from that direction in every scene, leading to the question of what happens when lighting direction cues like shading and shadows indicate a lighting direction different from the direction suggested by the light-from-above prior. Morgenstern, Murray, and Harris [24] showed that instead of the prior overriding lighting direction cues or vice versa, information about lighting direction from the prior and from lighting direction cues is combined, so that the perceived lighting direction that guides shape from shading is a compromise between the prior and direction cues. They also found that the light-from-above prior is remarkably weak, in the sense that even very faint cues to lighting direction have a greater effect in this compromise. This suggests that the light-from-above prior has little influence in everyday perception.

Recent psychophysical work suggests that human vision also relies on assumptions about the diffuseness of lighting. Boyaci et al. [11] and Bloj et al. [10] examined lightness constancy as a function of the orientation of the test patch being judged. They found that constancy was quite good when the patch was within $\pm 60^\circ$ of frontoparallel to the dominant light source, but that reflectances were consistently underestimated when the test patch was more oblique than this. They noted that this unusual pattern of success and failure in judging surface reflectance is what one would expect from observers who overestimate lighting diffuseness in the scene

[2]Brewster [12] is often credited with discovering the light-from-above prior. In fact, he mostly elaborated Rittenhouse's [33] observation that we perceive ambiguous shaded patterns as having a 3D shape that is consistent with whatever we believe about the lighting direction in the scene being viewed. Neither Rittenhouse nor Brewster suggested that we have a default assumption that light comes from overhead.

being viewed. Boyaci et al. and Bloj et al. calculated the level of assumed diffuseness that would explain each observer's performance, and they found that observers consistently behaved as if lighting was much more diffuse than it actually was in the experimental apparatus. This suggests that observers' diffuseness estimates may have been biased by a prior for high levels of diffuseness. It is also possible, though, that observers simply always overestimate lighting diffuseness; it would be interesting to see whether experiments similar to Boyaci et al.'s and Bloj et al.'s, but with very diffuse light, find that observers underestimate diffuseness, as would occur if their errors were due to a lighting prior that favored high but not maximal levels of diffuseness.

Schofield, Rock, and Georgeson [34] used very different methods than Bloj and Boyaci, and reached similar conclusions. They noted that observers tend to see sinusoidal luminance gratings as corrugated surfaces that are sinusoidal in depth, and that the phase difference between the luminance grating and the perceived depth grating changes with the orientation of the luminance grating. They showed that this is what one would expect from an observer who has a prior on lighting direction, and they showed that, under certain modelling assumptions, the magnitude of the phase change across luminance grating orientations is a signature of the observer's assumptions about lighting diffuseness: a smaller phase change corresponds to an assumption of more diffuse light. They calculated the level of diffuseness that explained each observer's shape percepts, and their results were quantitatively very similar to Bloj et al.'s. Furthermore, Morgenstern [22, 23] showed that the diffuseness levels that Bloj et al. and Schofield et al. arrived at match the range of lighting diffuseness found in real world environments. This suggests that just as observers have a prior on lighting direction that matches the average lighting direction in real world scenes, they also have a prior on lighting diffuseness that matches real world lighting. Boyaci et al.'s observers seem to have been guided by much higher levels of diffuseness, but Morgenstern argues that Boyaci et al.'s results were probably biased by partial failures of lightness constancy in perceiving computer-generated stimuli.

Fleming, Dror, and Adelson [15] found that human vision also relies on higher-order properties of natural lighting, beyond its direction and diffuseness. They examined material perception under real and synthetic lighting, and they found that natural lighting is important for accurate perception of material properties, in particular for perception of gloss. They concluded that one of the most important properties of natural lighting, for human observers, is that it contains point-like light sources and spatially extended light sources with edges.

16.3 Shape and Reflectance

16.3.1 Shape and Reflectance: Scene Statistics

Torreão [37] developed a Markov random field (MRF) approach to shape from shading, that used mathematically convenient assumptions about object shape to arrive

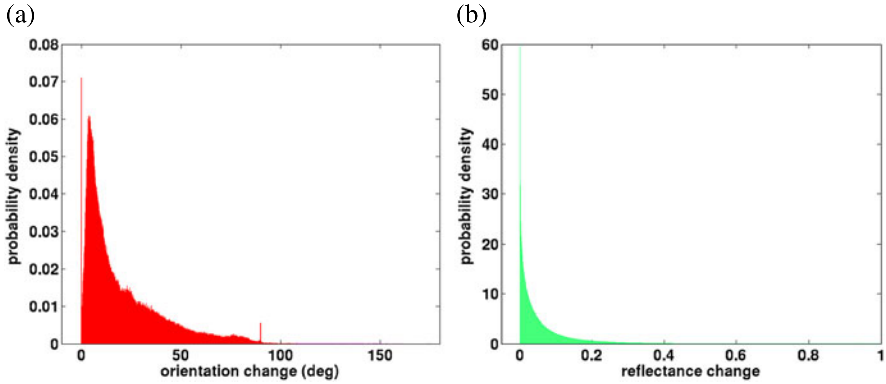

Fig. 16.3 (**a**) Orientation and (**b**) reflectance changes across real-world objects [26]. These histograms were created by viewing digital scans of real objects of various sizes through a virtual 100×100 grid, and measuring the changes in orientation or reflectance between adjacent grid elements. Here reflectance is measured on the interval [0,1]

at maximum a priori (MAP) shape interpretations of shaded images. Freeman, Pasztor, and Carmichael [18] used a similar approach to the problem of recovering both shape and reflectance from shading, but with some important innovations. For our purpose, the most interesting development was that instead of using convenient assumptions about shape and an assumption of uniform reflectance, Freeman et al.'s algorithm learned the statistical distribution of shape and reflectance patches in a computer-rendered virtual world. The algorithm built up a library of surface patches from the virtual world, with both shape and reflectance information represented, and interpreted new images by assembling a grid of surface patches from the library that best accounted for the luminance patterns in the shaded image.

Murray [26] examined the shape and reflectance distribution in real objects. He used 3D digital scans of 80 natural and man-made objects. From random viewpoints, he examined how surface orientation and reflectance changed across the surface of the objects. The histograms of surface and reflectance changes were highly regular (Fig. 16.3). One of the most noteworthy findings was that reflectance changes were much more narrowly peaked around zero than surface orientation changes, suggesting that a rational visual system will tend to attribute shading changes to surface orientation changes instead of reflectance changes whenever possible.

Barron and Malik [6, 7] developed an MRF algorithm to recover shape and reflectance from shading, but unlike Torreão and Freeman et al. they incorporated shape and reflectance statistics of real objects. They used parametric models of shape and reflectance gradients, and they fit these models to digital scans of ten real world objects. They found that the MAP estimates from the resulting algorithm were able to recover shape and reflectance from shaded images, and that scene statistics learned from a set of training objects worked well with a new set of test objects.

These studies support the notion that natural scene statistics can overcome the ambiguity inherent in 2D images. Murray, and Attewell and Baddeley [5], show that scene statistics are sufficiently stable that they can be measured with a reasonable number of samples. Torreão, Freeman et al., and Barron and Malik show that probabilistic algorithms that incorporate assumptions about scene statistics can recover shape and reflectance from shaded images. One shortcoming of work to date is that it has not given us a good theoretical understanding of what statistical regularities in 3D scenes are important for recovering shape and reflectance. For instance, is it enough to assume in some way, as Murray suggests, that reflectance changes are rare compared to shape changes, or is the specific parametric form that Barron and Malik assume for reflectance changes also important? Is there any role for correlations between shape and reflectance? Researchers have developed algorithms that learn to infer shape and reflectance, demonstrating that scene statistics can overcome image ambiguity, but there has been less progress in determining what the key properties of natural scenes are, that such algorithms learn.

One possibility that has been overlooked in previous work is that measuring the precise distribution of shape and orientation in real world scenes might be less important than using statistical properties that one can predict from first principles. For instance, in GBR and GBR-like transformations, an object's shape, reflectance, and lighting are put through changes that cancel one another precisely, e.g., if the shape and lighting transformations result in a lower image luminance for a given surface patch, then the surface patch's reflectance is increased in order to undo the change in image luminance. This introduces statistical dependencies between surface orientation, illuminance, and reflectance. Some of these dependencies might be very unnatural. In natural objects, for instance, we expect illuminance and reflectance to be statistically independent across surface patches. Such almost-a-priori constraints may be useful for finding correct 3D interpretations of shaded images.

16.3.2 Shape and Reflectance: Psychophysics

Very little is known about the assumptions that human observers make about shape and reflectance. One of the most successful approaches to lightness perception is Gilchrist's anchoring theory, a set of rules for predicting human lightness percepts under a wide range of conditions [19, 20]. Gilchrist [19] argues that Bayesian theories of lightness perception are unlikely to be successful, because they are normative theories, and will not account for the systematic errors in human lightness perception that are observed empirically and that form an important part of anchoring theory.

Adelson [3] proposes an alternative, Bayesian theory of lightness perception. He suggests that human observers assume that real world reflectances follow some statistical distribution, and that observers also assume that reflectances r and image luminances l are related by an affine transformation, $l = mr + b$. From observed luminances l, observers make statistical estimates of the reflectances r and the lighting condition parameters m and b.

Recent work in my laboratory has shown that a Bayesian theory along these lines accounts for much of anchoring theory [27]. The Bayesian theory assumes that

(a) reflectances follow a broad, asymmetric normal distribution, (b) lighting consists of multiplicative and additive components, so luminance and reflectance are related by $l = mr + b$, and (c) the proportion of additive light $b/(m + b)$ tends to be low. This simple theory predicts many of the rules of anchoring theory, thereby showing that some systematic errors in human lightness perception are actually rational consequences of simple assumptions about lighting and reflectance.

One obstacle to understanding the human visual system's assumptions about real world scenes is that we know little about how the visual system represents scenes. I have spoken, for instance, of the human visual system's assumptions about reflectance, but reflectance is a notion that is most useful in a Lambertian imaging model. The image luminance of a non-Lambertian surface varies as a function of incident lighting and/or viewing direction, so to describe such surfaces we need more information than the proportion of light reflected. Few real world surfaces are truly Lambertian, and this fact along with recent work on material perception (e.g., [4, 25]) suggests that the human visual system does not rely on a purely Lambertian model. Without knowing more about the human visual system's model of surfaces, though, it is difficult to know what properties of real world surfaces we should try to characterize statistically. Similar comments apply to the perceptual representation of surface shape and lighting conditions. For instance, Fleming, Holtmann-Rice, and Bülthoff [16] suggest that the human visual system infers 3D shape from the local orientation field of the retinal image, in which case an important goal for studies of 3D scene statistics would be to examine the statistical relationship between local image orientation and 3D shape in real world scenes.

16.4 Conclusion

There has been important progress on understanding statistical properties of 3D natural scenes and how they guide human vision. Nevertheless, the most fundamental problems are still almost completely open.

What assumptions about 3D scenes guide human perception of shape and reflectance? Consider assumptions about lighting. Human observers certainly have a prior that light comes from above, but Morgenstern et al. [24] have shown that this prior is very weak, and probably unimportant in everyday perception. Bloj et al., Boyaci et al., and Schofield et al. report intriguing evidence that human observers have a prior on lighting diffuseness, but these results are highly model-dependent, and further work is needed before we can say with confidence exactly what assumptions human observers make about diffuseness, and how these assumptions guide visual perception.

Our understanding of assumptions about shape and reflectance is even more tentative. Work in computer vision has shown that assumptions about shape, reflectance, and lighting can be used to estimate 3D scene properties. These studies provide functioning algorithms, but they leave many basic questions unanswered. For instance, what types of image luminance patterns are best attributed to shape patterns, and what types are best attributed to reflectance patterns? Are any general principles possible, such as that changes in reflectance are less likely than changes in surface orientation? Furthermore, what scene properties does the human visual system have priors on, e.g., reflectance, or the gradient of reflectance, or both? And how strong are the various priors that the human visual system relies on, e.g., priors on shape vs. priors on reflectance? It is remarkable, when there is such broad support for the notion that assumptions about natural scenes play a crucial role in perception of 3D scenes, that these fundamental questions are still largely unanswered.

References

1. Adams WJ, Graf EW, Ernst MO (2004) Experience can change the 'light-from-above' prior. Nat Neurosci 7:1057–1058
2. Adams WJ (2007) A common light-prior for visual search, shape, and reflectance judgments. J Vis 7(11):1–7
3. Adelson EH (2000) Lightness perception and lightness illusions. In: Gazzaniga M (ed) The new cognitive neurosciences, 2nd edn. MIT Press, Cambridge, pp 339–351
4. Anderson BL, Kim J (2009) Image statistics do not explain the perception of gloss and lightness. J Vis 9(11):10, 1–17
5. Attewell D, Baddeley RJ (2007) The distribution of reflectances within the visual environment. Vis Res 47:548–554
6. Barron JT, Malik J (2011) High-frequency shape and albedo from shading using natural images. In: IEEE conference on computer vision and pattern recognition, Colorado Springs, Colorado, June 20–25, pp 2521–2528
7. Barron JT, Malik J (2012) Shape, albedo, and illumination from a single image of an unknown object. In: IEEE conference on computer vision and pattern recognition, Providence, Rhode Island, June 16–21, pp 334–341
8. Basri R, Jacobs DW (2002) Lambertian reflectance and linear subspaces. IEEE Trans Pattern Anal Mach Intell 25:218–233
9. Belhumeur PN, Kriegman DJ, Yuille AL (1999) The bas-relief ambiguity. Int J Comput Vis 35:33–44
10. Bloj M, Ripamonti C, Mitha K, Hauck R, Greenwald S, Brainard DH (2004) An equivalent illuminant model for the effect of surface slant on perceived lightness. J Vis 4:735–746
11. Boyaci H, Maloney LT, Hersh S (2003) The effect of perceived surface orientation on perceived surface albedo in binocularly-viewed scenes. J Vis 3:541–553
12. Brewster D (1826) On the optical illusion of the conversion of cameos into intaglios, and of intaglios into cameos, with an account of other analogous phenomena. Edinb J Sci 4:99–108
13. Debevec P (1998) Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: Proceedings of SIGGRAPH 98. Annual conference series, pp 189–198
14. Dror RO, Willsky AS, Adelson EH (2004) Statistical characterization of real-world illumination. J Vis 4:821–837
15. Fleming RW, Dror RO, Adelson EH (2003) Real-world illumination and the perception of surface reflectance properties. J Vis 3:347–368
16. Fleming RW, Holtmann-Rice D, Bülthoff HH (2011) Estimation of 3D shape from image orientations. In: Proceedings of the National Academy of Sciences of the USA, vol 108, pp 20438–20443
17. Freeman WT (1994) The generic viewpoint assumption in a framework for visual perception. Nature 368:542–545
18. Freeman WT, Pasztor EC, Carmichael OT (2000) Learning low-level vision. Int J Comput Vis 40:25–47

19. Gilchrist A (2006) Seeing black and white. Oxford University Press, New York
20. Gilchrist A, Kossyfidis C, Bonato F, Agostini T, Cataliotti J, Li X, Spehar B, Annan V, Economou E (1999) An anchoring theory of lightness perception. Psychol Rev 106:795–834
21. Metzger W (2006) The laws of seeing. The MIT Press, Cambridge. (Original work published 1936)
22. Morgenstern Y (2011) The role of low-pass natural lighting regularities in human visual perception. Ph.D. thesis, York University
23. Morgenstern Y, Murray RF, Geisler WS (submitted) Human vision is attuned to the diffuseness of natural light
24. Morgenstern Y, Murray RF, Harris LR (2011) The human visual system's assumption that light comes from above is weak. Proc Natl Acad Sci USA 108:12551–12553
25. Motoyoshi I, Nishida S, Sharan L, Adelson EH (2007) Image statistics and the perception of surface qualities. Nature 447:206–209
26. Murray RF (2006) Local 3D shape and orientation statistics of natural surfaces. In: Vision sciences society annual meeting, Sarasota, Florida, May 5–10
27. Murray RF (2013) Human lightness perception is guided by simple assumptions about reflectance and lighting. In: Human vision and electronic imaging XVII, San Francisco, California, February 4–7, 2013
28. Mury AA, Pont SC, Koenderink JJ (2007) Light field constancy within natural scenes. Appl Opt 46:7308–7316
29. Mury AA, Pont SC, Koenderink JJ (2009) Representing the light field in finite three-dimensional spaces from sparse discrete samples. Appl Opt 48:450–457
30. Mury AA, Pont SC, Koenderink JJ (2009) Structure of light fields in natural scenes. Appl Opt 48:5386–5395
31. Phong BT (1975) Illumination for computer generated pictures. Commun ACM 18:311–317
32. Ramamoorthi R, Hanrahan P (2001) On the relationship between radiance and irradiance: determining the illumination from images of a convex Lambertian object. J Opt Soc Am A 18:2448–2459
33. Rittenhouse D (1786) Explanation of an optical deception. Trans Am Philos Soc 2:37–42
34. Schofield AJ, Rock PB, Georgeson MA (2011) Sun and sky: does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading? Vis Res 51:2317–2330
35. Sun J, Perona P (1998) Where is the Sun? Nat Neurosci 1:183–184
36. Todd J (2004) The visual perception of 3d shape. Trends Cogn Sci 8:115–121
37. Torreão JRA (1995) Bayesian shape estimation: shape-from-shading and photometric stereo revisited. Mach Vis Appl 8:163–172
38. Yuille AL, Coughlan JM, Konishi S (2003) The generic viewpoint assumption and planar bias. IEEE Trans Pattern Anal Mach Intell 25:775–778