

Sun and sky: Does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading?

Schofield, Andrew; Rock, PB; Georgeson, MA

DOI:

[10.1016/j.visres.2011.09.004](https://doi.org/10.1016/j.visres.2011.09.004)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Schofield, A, Rock, PB & Georgeson, MA 2011, 'Sun and sky: Does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading?', *Vision Research*, vol. 51, no. 21-22, pp. 2317-2330. <https://doi.org/10.1016/j.visres.2011.09.004>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

NOTICE: this is the author's version of a work that was accepted for publication in Vision Research. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in Vision Research, Vol 51, Issues 21-2, DOI: 10.1016/j.visres.2011.09.004.

This version of the article is subject to the terms of a Creative Commons Attribution Non-Commercial No-Derivatives license.

Checked June 2015

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

1 Sun and Sky: Does human vision assume a mixture of point and diffuse illumination when
2 interpreting shape-from-shading?

3

4 Andrew J Schofield¹, Paul B Rock¹, Mark A Georgeson²

5

6 1. School of Psychology, University of Birmingham, Edgbaston, Birmingham, UK, B15 2TT
7 2. School of Life and Health Sciences, Aston University, Aston Triangle, Birmingham, UK, B4
8 7ET UK

9

10 AJS: a.j.schofield@bham.ac.uk,

11 MAG: m.a.georgeson@aston.ac.uk

12

13

14 Corresponding author: Andrew. J. Schofield.

15 Email: a.j.schofield@bham.ac.uk.

16

17

18 This author post-print is provided under a Creative Commons: Attribution-NonCommercial-
19 ShareAlike Licence Copyright Andrew J Schofield, University of Birmingham

20

21 Supplementary file appended at end of document.

22 **Abstract**
23 People readily perceive smooth luminance variations as being due to the shading produced
24 by undulations of a 3-D surface (shape-from-shading). In doing so, the visual system must
25 simultaneously estimate the shape of the surface and the nature of the illumination.
26 Remarkably, shape-from-shading operates even when both these properties are unknown
27 and neither can be estimated directly from the image. In such circumstances humans are
28 thought to adopt a default illumination model. A widely held view is that the default illuminant
29 is a point source located above the observer's head. However, some have argued instead
30 that the default illuminant is a diffuse source. We now present evidence that humans may
31 adopt a flexible illumination model that includes both diffuse and point source elements. Our
32 models estimates a direction for the point source and then weights the contribution of this
33 source according to a bias function. For most people the preferred illuminant direction is
34 overhead with a strong diffuse component.
35
36 Keywords: shading, illumination, lighting-from-above, dark-is-deep.
37
38

39 **1. Introduction**

40 *1.1 Background*

41 It is well known that humans can discern the shape of a surface from the pattern of shading
42 produced when it is illuminated – shape-from-shading – even when there are no other cues
43 to shape in the image (Christou & Koenderink, 1997; Erens, Kappers & Koenderink, 1993;
44 Kleffner & Ramachandran, 1992; Langer & Bülthoff, 2000; Ramachandran, 1988; Todd &
45 Mingolla, 1983; Tyler, 1998). Note however that shape-from-shading is not always veridical
46 (Pentland 1988; Zhang, Tsai, Cryer, & Shah, 1999). To interpret shape-from-shading we
47 must simultaneously estimate the shape of the surface, its reflectance properties, and the
48 nature and direction of the illuminant – a task which is inherently ambiguous (D’Zmura, 1991;
49 Belhumeur, Kriegman, & Yuille, 1999). Nonetheless a number of different cues enable
50 humans to estimate the direction of illumination for a scene (Cavanagh & Leclerc, 1989;
51 Erens et al., 1993; Gerhard & Maloney, 2010; Koenderink, van Doorn, Kappers, te Pas, &
52 Pont, 2003; Koenderink, van Doorn, & Pont, 2004; Koenderink, van Doorn, & Pont, 2007; Liu
53 & Todd, 2004; Norman, Todd, & Orban, 2004; Pentland, 1982; Todd & Mingolla, 1983) and
54 although such estimates are not always accurate when the scene is well articulated we can
55 estimate the light field with considerable accuracy (Koenderink, Pont, van Doorn, Kappers &
56 Todd, 2007). In the absence of cues to lighting we may assume a default light source. There
57 is debate, however, about the nature of the default light source. Several studies have
58 suggested that humans assume a single spatially limited (point) light source located
59 approximately overhead (Adams, Graf & Ernst, 2004; Mamassian & Goutcher, 2001;
60 Ramachandram, 1992; Sun & Perona, 1998): lighting-from-above. In contrast, Langer &
61 Bülthoff (2000) showed that humans can, if required, interpret shape-from-shading to be
62 consistent with a diffuse, multidirectional light source. Tyler (1998) argues that diffuse
63 illumination is the primary default assumption for highly reduced scenes.

64

65 The lighting-from-above assumption seems to explain a range of illusions – known
66 collectively as the crater illusion – where, in monocular viewing, perceived surface shape
67 flips from convex to concave when the image is rotated through 180° (Brewster, 1826; Hess,
68 1950; Ramachandran, 1992; Rittenhouse, 1786; von Fieandt, 1949). Lighting-from-above
69 makes clear predictions about the relationship between shape and luminance. For a
70 Lambertian surface, luminance at any point will be proportional to the cosine of the angle
71 between the surface normal and the line joining the point to the light source. Parts of the
72 surface that point towards the light source will have the highest luminance.

73

74 Diffuse illumination (such as used by Langer & Bülthoff, 2000) represents the situation on a
75 cloudy day where a horizontal surface is illuminated about equally from all parts of a

76 hemispherical 'sky' this being (for surfaces) equivalent to a fully spherical illumination field or
77 Ganzfeld (through this paper we use the terms diffuse or fully diffuse to mean this type of
78 lighting). Diffuse illumination also leads to clear predictions about the relationship between
79 shape and luminance. The luminance at any point on a diffusely illuminated Lambertian
80 surface is approximately proportional to the size of the aperture formed by the rest of the
81 surface (Langer & Zucker, 1997; Stewart & Langer, 1997; Tyler, 1998). Points down a slope,
82 in a pit or in a ravine 'see' less of the sky and are hence dark (the dark-is-deep rule).
83 However, at the very bottom of a ravine or pit the surface points directly towards the un-
84 obscured sky producing a small localized luminance peak (see Langer & Bülthoff, 2000).

85
86 Although both the lighting-from-above and diffuse illumination models have some ecological
87 validity, neither correspond well to everyday lighting conditions. Humans are generally
88 immersed in an illumination field that is highly diffuse but biased towards the sky because of
89 the location of the sun (or room lights) and the relatively low reflectance values of ground-
90 level objects. For real, outdoor situations the illumination reaching a point from any given
91 direction decreases monotonically with decreasing elevation except for a local dip around
92 horizontal (Dror, Willsky & Adelson, 2004; Teller, Antone, Bosse, Coorge, Jethwa, &
93 Masters, (2001); see also Mury, Pont & Koenderink, 2009). It is likely then that people
94 assume a default illumination model that resembles everyday experience and that therefore,
95 when the nature of the illuminant is uncertain, they assume that objects are lit by a light field
96 that is quite diffuse but with a directional component.. We test this hypothesis here.

97

98 *1.2 Choice of stimuli*

99 Langer and Bülthoff (2000) presented observers with images of complex undulating surfaces
100 rendered under either point source or diffuse lighting. The resulting depth judgements show
101 that humans are able to switch between point and diffuse light interpretations depending on
102 cues in the stimulus presented. This suggests that the default illumination model assumed
103 by human vision can only be exposed when the stimuli are ambiguous with respect to
104 illumination. We therefore avoid complex rendered stimuli and present instead simple stimuli
105 which we show are likely to be ambiguous with respect to illumination. In this we follow the
106 lead of Sun & Perona (1998) and Mamassian & Gouche (2001) who presented stimuli
107 where the direction of the illumination was ambiguous. In our case, however, it is the nature
108 of the illumination (diffuse vs point) that is uncertain. We also need test stimuli that are
109 expected to produce quantitatively different results depending on the nature of the assumed
110 illumination.

111

112 People tend to perceive sinusoidal shading patterns¹ as sinusoidally undulating surfaces
113 (see Pentland, 1988 and Schofield, Hesse, Rock & Georgeson, 2006; we also present
114 control data in the supplementary file to further support this claim using stimuli from our main
115 experiments) despite the fact that such surfaces do not always give rise to sinusoidal
116 luminance profiles when shaded (see supplementary Figures S1 & S2). However, the
117 analysis presented below and in the supplementary file shows that sinusoidal undulations do
118 give rise to approximately sinusoidal shading profiles under point source lighting for a range
119 of surface orientation; Figure 1 shows examples where this is the case.

120

121 **Figure 1 about here (double column)**

122

123 We see from Fig 1 that point source lighting produces either approximately sinusoidal
124 shading profiles with luminance peaks offset from the physical surface peaks by $\frac{1}{4}$
125 wavelength, double crested peaks centred on the $\frac{1}{4}$ wavelength offset, or a frequency
126 doubled signal. The $\frac{1}{4}$ wavelength offset is counter intuitive and we now show that this offset
127 does not vary with either the surface or lighting conditions so long at the shading profile has
128 a single peak. Following Pentland (1988) we approximate the luminance at any point on a
129 Lambertian surface with the following equation:

130

131
$$L \approx \cos \theta + p \cos \tau \sin \theta + q \sin \tau \sin \theta - \cos \theta (p^2 - q^2)/2 \quad \text{eqn1}$$

132

133 where θ is the slant of the light source (the angle that the illuminant vector makes with the z-
134 axis) τ is the tilt of the light source (the angle between the x-axis and the projection of the
135 light source vector onto the surface plane, p is the partial derivative of the surface with
136 respect to x and q its partial derivative with respect to y . Let the surface $z = a \cos(\omega x)$, where
137 ω is the undulation frequency and a is the surface amplitude; hence $p = -a\omega \sin(\omega x)$
138 and $q=0$. We further redefine the lighting angles in terms of the elevation angle
139 ($e = \pi/2 - \theta$) between the light vector and the image plane and direction angle ($d = \tau \mp$
140 $\pi/2$ being the angle between the projection of the light vector into the surface plane and the
141 y-axis (vertical) we define positive changes in d as clockwise rotations (see Fig 1B for a
142 diagram of the lighting geometry). Equation 1 can thus be rewritten as:

143

144
$$L \approx \cos\left(\frac{\pi}{2} - e\right) - a\omega \sin(\omega x) \cos\left(\frac{\pi}{2} - d\right) \sin\left(\frac{\pi}{2} - e\right) - \cos\left(\frac{\pi}{2} - e\right) (a^2 \omega^2 \sin^2(\omega x))/2 \quad \text{eqn2}$$

145

146

¹ Stimuli where luminance is a sinusoidal function of position in the image.

147 In order to find the locations of peaks and troughs in L we need to differentiate eqn 2 with
148 respect to x :

149

150 $\frac{dL}{dx} \approx -a\omega^2 \cos(\omega x) \cos\left(d + \frac{\pi}{2}\right) \sin\left(\frac{\pi}{2} - e\right) - \cos\left(\frac{\pi}{2} - e\right) a^2 \omega^3 \sin(2\omega x)/2$ eqn3

151

152 This has two components one with extrema located at $\pi/2$ and $3\pi/2$ (frequency = ω ; offset
153 = $\pi/2$) and the other with extrema at $0, \pi/2, \pi$, and $3\pi/2$ (frequency = 2ω). The locations
154 of these extrema do not change with surface amplitude or frequency nor with lighting
155 direction however the ratio of the two components does change introducing double crested
156 peaks and ultimately frequency doubling for some lighting conditions in a manner that also
157 depends of a and ω (larger values favour double peaks). Aside from cases where the
158 frequency doubled term dominates completely luminance will always have a positive lobe at
159 either $\pi/2$, or $3\pi/2$ ($1/4$ wavelength offset). Double peaks occur by virtue of local minima at
160 these locations. Thus we can identify double crested peaks by examining the extrema at
161 $\pi/2$, and $3\pi/2$; if both are minima then the peak is double crested. This in turn can be
162 determined from the second derivative of luminance:

163

164 $\frac{d^2L}{dx^2} \approx a\omega^3 \sin(\omega x) \cos\left(d + \frac{\pi}{2}\right) \sin\left(\frac{\pi}{2} - e\right) - \cos\left(\frac{\pi}{2} - e\right) a^2 \omega^4 \cos(2\omega x)$ eqn4

165

166 setting $x=\pi/2$ or $3\pi/2$. We can thus in principle find the lower limit of tilt giving single peaks
167 for each combination of slant, amplitude and frequency. For the amplitude : wavelength ratio
168 used in the current study (0.12) this lower limit is depicted by the border of the inner black
169 lozenge in supplementary Fig S2b. It is clear that, in the absence of double peaks,
170 luminance will always peak at an offset of $1/4$ wavelength from the physical surface peak and
171 that when double crested peaks do occur they will always lie either side of lobe centred on $1/4$
172 wavelength offset. When full frequency doubling occurs luminance peaks will always occur
173 at the peaks and troughs of the surface. Finally frequency doubling will always occur when
174 the elevation of the light is $\pi/2$ (frontal lighting) or when the lighting direction is either zero or
175 π (lighting from above or below)

176

177 Diffuse lighting where the surface is illuminated from all directions will produce approximately
178 sinusoidal shading regardless of the surface orientation (Figure 2). Comparing Figures 1 and
179 2 we note that point source illumination produces either a $1/4$ wavelength offset between

surfaces peaks and luminance peaks or results in frequency doubling² , whereas diffuse lighting produces neither an offset nor frequency doubling.

182

Figure 2 about here – Single column

184

It follows from the superposition rule that a mixture of diffuse and point source lighting will produce waveforms that are approximately sinusoidal but with luminance peaks that are offset from the physical surface peaks by some amount between 0 and $\frac{1}{4}$ wavelength. The addition of two sine waves with the same frequency but different amplitudes and phases being a sine wave with the same frequency but intermediate phase. This will hold so long as the point source term is not dominated by frequency doubling. There will also be a localised peak at the surface troughs due to the diffuse lighting component. Figure 3 plots the luminance profiles for oriented surfaces under mixed illumination in the format of Fig 1. The offset between the fundamental and the physical surface clearly changes with the physical orientation of the surface; as indicated in Fig 3. We note that even for vertically oriented surfaces lit from above the shading profile is dominated by the fundamental not the frequency doubled component.

197

Figure 3 about here double column

199

At least in terms of the generative models outlined above and in the supplementary file the relationship between surface profiles and shading is different for the three types of lighting even though sinusoidal shading is a reasonable approximation in all cases. Point source lighting produces $\frac{1}{4}$ wavelength offsets; diffuse lighting - zero offset; and mixed illumination offsets that vary with surface orientation. We thus propose sinusoidal shading patterns as a diagnostic stimulus for the default illumination model used in human shape-from-shading. If people were apply the inverse of the appropriate generative model (at least approximately) to estimate shape, we could determine which model had been adopted by observing the offsets between luminance and perceived surface peaks (inter-peak offset) and the tendency towards frequency halving at some orientations (undoing the doubling found in the point source case). This assumption is central to our method so we expand on it below.

211

212 Point source assumption: how might a $\frac{1}{4}$ wavelength inter-peak offset arise in human vision?
213 If people were to assume point source lighting then the peaks of the perceived surface
214 should (in general) be shifted away from the luminance peaks. The linear relationship noted

² Under conditions where approximately sinusoidal shading is achieved, see supplementary file.

215 by Pentland (1988) will be valid if the direction/elevation of the light source is oblique (or
216 assumed to be oblique). In Pentland's (1988) model for human shape-from-shading,
217 luminance components are subject to a 90° phase shift in the frequency domain. Thus for
218 sinusoidal shading under Pentland's model, perceived surface peaks will be offset by $\frac{1}{4}$
219 wavelength from luminance peaks: nicely undoing the offset produced by point-source
220 shading in the first place (see Figures 1 & supplementary file). Alternatively recovering
221 shape-from-shading is sometimes characterized as an integration process in which
222 perceived surface gradient is proportional to luminance. Integration would also produce a $\frac{1}{4}$
223 wavelength inter-peak offset for sinusoidal shading, although the presence of bounding
224 contours will alter the integration process by setting its boundary conditions. Thus the
225 generative model, Pentland's model, and integration models all support the notion that
226 sinusoidal luminance profiles should be seen as sinusoidal surfaces with a $\frac{1}{4}$ wavelength
227 offset between perceived surface peaks and luminance peaks *if* the observer assumes a
228 point light source.

229

230 In cases where a point source illuminant is aligned with the surface undulations (eg. upper
231 trace in Figure 1) the shading profile is dominated by a quadratic component (frequency
232 doubling). If people were to allow for such quadratic shading we would expect them to see
233 stimuli aligned with their assumed point source direction as surfaces undulating at half the
234 spatial frequency of the shading pattern. However, Pentland (1988) has shown that people
235 do not allow for quadratic shading when interpreting shape-from-shading although, as we
236 outline in Section 1.4, this alone does not mean that people do not assume point source
237 lighting in this special case.

238

239 Diffuse source assumption: how might a zero inter-peak offset arise in human vision? Langer
240 and Bülthoff (2000) found that when surfaces are rendered under diffuse illumination
241 humans adopt the dark-is-deep rule whereby peaks in the perceived surface align with
242 luminance peaks. The strict application of the dark-is-deep rule would predict a small
243 localized peak in the perceived surface in the bottom of valleys due to the local peak in
244 luminance at such points, but Langer & Bülthoff (2000) found that people do not perceive
245 peaks at these locations. Rather, their data were best characterized by a model in which the
246 luminance profile associated with the surface was first blurred, attenuating small local peaks,
247 and then interpreted according to the dark-is-deep rule. Note that such blurring could render
248 the luminance profile of Figure 2b identical to a sinusoidal profile and that if the stimulus is
249 itself sinusoidal, blurring, by say a Gaussian filter, will only alter the amplitude of the signal:
250 not its shape or position . If people were to assume diffuse lighting Langer and Bulthoff's

251 (2000) blur+dark-is-deep model would predict a sinusoidal surface interpretation but with no
252 offset between perceived luminance peaks and surface peaks.

253

254 Mixed source assumption: how might intermediate inter-peak offsets arise in human vision?
255 There is little support in the literature for this case. However, noting that shading profiles
256 under mixed illumination tend to be quite irregular, suitable offsets could be achieved in one
257 of two ways: first by reversing the image generation process under a mixed lighting
258 assumption; and second a combination of the blur+dark-is-deep rule and Pentlands (1988)
259 model with a stimulus dependent weighting between the two. We show later that for
260 sinusoidal shading patterns these two alternatives make identical predictions.

261

262 Thus we are confident that people are, in principle, capable of interpreting our stimuli
263 according to either a point or diffuse lighting assumption; and we can see a route by which a
264 mixed illumination assumption might be implemented. The question is which assumption
265 dominates.

266

267 1.3 Bas-Relief ambiguity

268 Point-source lighting of surfaces produces ambiguous luminance profiles due to the
269 generalised Bas-Relief (GBR) ambiguity and the related convex/concave ambiguity
270 (Belhumeur, et al., 1999; see also D'Zmura, 1991). Any shaded surface can be modified by
271 a GBR transformation which when coupled with a suitably transformed lighting and albedo
272 profiles will produce the same luminance profile as the original surface and lighting
273 conditions. Humans are thus unable to make good judgements about (for example) the
274 amount of relief applied to two sculptures. This ambiguity has some relevance to shape-from-
275 shading studies in general. Of more critical interest here however is the convex/concave
276 ambiguity in which a convex surface lit from below looks identical to a concave surface lit
277 from above. Prior assumptions for convexity and lighting-from above serve to stabilise this
278 percept and prevent perceptual flipping (Liu & Todd, 2004; Sun & Perona, 1998; and
279 Mamassian & Goucher, 2001). However, the convexity prior will not apply to sinewave
280 shading which appears corrugated (both convex and concave) and lighting from above will
281 only function for near horizontal stimuli. There is a strong chance then that the perceived
282 position of peaks in near vertical stimuli will flip between two possible positions. We explicitly
283 test for this.

284

285 1.4 *Experimental predictions*

286 Our main aim was to assess the default illumination model used by observers to interpret the
287 perceived shape of simple shading patterns in the absence of other cues to surface shape.

288 We also wanted to avoid explicit cues to the nature of the light source. Sinusoidal luminance
289 patterns can approximate the shading obtained under both point-source illuminants (Figure
290 1) and, to a lesser extent, diffuse illumination (Figure 2) while avoiding the above
291 confounding factors; we therefore used sinusoidal gratings as our shading stimuli (we did not
292 use rendered stimuli; see Section 2 for further justification). Observers were free to adopt
293 any lighting hypothesis in order to ‘make sense of’ the stimuli. We presume that observers
294 may have a preference for lighting-from-above (Adams, et al., 2004; Mamassian & Goutcher,
295 2001; Ramachandram, 1992; Sun & Perona, 1998) when adopting a point-source
296 hypothesis, but this is by no means fundamental to the experiment.

297
298 An important diagnostic case occurs when sinusoidal shading patterns align with the
299 observers preferred lighting direction for point source illumination. Given that most observers
300 prefer lighting from above (Mamassian and Goucher, 2001) this special case most often
301 corresponds to a vertical sinewave. We have no reason to suppose that this is any less
302 common a visual experience than any other sinewave. We test five predictions for this
303 critical case. (1) If people perceive such surfaces to be lit from their preferred direction by a
304 point source, and have at least an implicit model of the physics of shading under such
305 conditions, then they would perceive such a surface to have half the frequency of the
306 luminance profile (undoing the quadratic shading or frequency doubling seen for vertical
307 surfaces in Figure 1). (2) Alternatively, people might perceive such surfaces as lit by a point-
308 source but alter their estimate of the direction of this source consistently to one side or the
309 other. If this were so they would perceive a surface with the same frequency as the shading
310 and would retain the inter-peak offset expected at other orientations. (3) We might, however,
311 expect such an interpretation to be bi-stable owing to the convex-concave ambiguity (section
312 1.3) which is most intrusive when shading gradients are orthogonal to the observer’s
313 preferred light source (Sun and Perona, 1998). Such bi-stability would result in the inter-peak
314 offset flipping between two locations either side of zero. Here the average offset would fall to
315 zero but the distribution of offsets would become bi-modal. (4) If (as suggested by Langer &
316 Bülthoff, 2000) people can switch between point- and diffuse lighting interpretations
317 depending on stimulus cues, they might prefer a diffuse model for (close to) vertical stimuli. If
318 so they should shift from using the luminance=gradient ‘rule’ to the ‘dark-is-deep’ rule and
319 inter-peak offsets will vary accordingly. (5) Finally, shape-from-shading may fail for some
320 stimuli. Specifically, sinusoidal stimuli may fail to elicit a depth percept at some orientations,
321 causing people to perceive surfaces as flat at these orientations, thus degrading estimates of
322 inter-peak offset.

323

324 **2. Experiment 1. Inter-peak offsets**

325 The purpose of this experiment was to measure the spatial offset between the luminance
326 peaks of a shading pattern and the associated peaks of the perceived surface. In particular
327 we asked how this inter-peak offset varies with stimulus orientation in the frontal plane. We
328 used a haptic matching task in which observers adjusted the position of a haptically defined
329 sinusoidal surface to match that of a visually perceived surface. In contrast to most studies
330 of shape-from-shading, but in common with Pentland (1988) and Kingdom (2003), our stimuli
331 (see Figure 3) were sinusoidal gratings imposed on iso-tropic textures. They were not
332 rendered surfaces. The textures help to articulate the shading cue but introduce no depth
333 cues in themselfes. We used these stimuli because they give the observer freedom to
334 interpret the shading cue in the absence of other cues to shape or overt cues to the nature of
335 the light-source. We justify this as follows: (1) Our stimuli contain no geometric cues, either
336 in the form of distortions in the texture or bounding contours, that might otherwise bias the
337 shape-from-shading process. (2) They do not include any sharp luminance edges that could
338 be associated with shadows nor do they contain double-crested peaks, and so they do not
339 promote a point-source lighting interpretation. (3) They do not contain mini-peaks between
340 each luminance peak and therefore do not promote a diffuse lighting interpretation either. (4)
341 As anisotropic stimuli they are largely uninformative about the direction of the light source
342 (Koenderink, et al., 2007). Despite the fact that the visual stimuli were not produced by a
343 graphical rendering of a model surface, observers readily perceived the stimuli as corrugated
344 surfaces as we show in the supplementary file (Section S2).

345

Figure 4 about here (2 column)

347

348 By using highly under-constrained stimuli we hope to reveal internal observer biases. In
349 particular our stimuli are mostly free from cues that might promote either a point or diffuse
350 lighting interpretation (although they are a better approximation for point-source lighting: *cf*
351 Figures 1 and 2). In this sense we differ from Langer & Bülthoff (2000) who used realistic
352 rendering to bias observers to one or other light source type and Tyler (1998) whose radial
353 sine waves could not be interpreted as lit by a single point-source.

354

355 Based on the results of Schofield et al. (2006; see also Pentland, 1988, and the control
356 experiment in the supplementary file) we suppose that humans naturally perceive sinusoidal
357 luminance profiles as sinusoidally corrugated surfaces. However, it is possible that humans
358 adopt a very flexible approach to shape-from-shading, balancing a number of *a priori*
359 constraints so as to perceive the combination of surface shape and illumination profile that is
360 most likely to occur in real world situations. Therefore, given our overall aim of assessing the
361 lighting model used by observers, we felt it important – at least in the first instance – to fix

362 the surface interpretation. Asking observers to match haptically defined sine waves to visual
363 stimuli should enhance the impression that the visible surfaces were sinusoidal (Wijntjes,
364 Volcic, Pont, Koenderink, & Kappers, 2009) thus leaving their internal lighting model as the
365 only thing free to vary in order to ‘make sense of’ the stimuli presented.

366

367 *2.1 Method*

368 2.1.1 Procedure and stimulus details.

369 Observers adjusted the position of the undulations of a virtual haptic surface to match the
370 perceived undulations in a visually presented stimulus. Visual stimuli (see Figure 3)
371 consisted of an isotropic, binary visual noise texture (mean contrast=0.1) whose luminance
372 values were multiplied by a sinusoidal profile ($1+c.\sin(2\pi fx)$; spatial frequency $f=0.4$ c/deg,
373 luminance contrast $c= 0.2$) so as to emulate multiplicative shading in which the local mean
374 luminance of the surface texture is modulated but not its local contrast. Such signals can be
375 produced by adding a sinewave luminance modulation while modulating the local amplitude
376 of the noise texture in phase with the luminance signal (LM+AM, see Schofield, et al., 2006
377 for details). The orientation of this shading pattern varied over the range 0-165° at 15°
378 intervals. Note that we measured stimulus orientation clockwise from vertical but later (and in
379 supplementary Fig S1) express positive increments in the direction of the illuminant as anti-
380 clockwise rotations. We use this convention because, in terms of the shading pattern
381 produced, a clockwise rotation of the stimulus is equivalent to an anti-clockwise shift in the
382 direction of the illuminant.

383

384 The wavelength of the sine wave modulation was 25mm and its phase was randomized on a
385 trial-by-trial basis. An orthogonal sinusoidal signal comprising both a luminance modulation
386 and an anti-phase amplitude modulation (LM-AM, see Schofield et al, 2006) was added to
387 each stimulus. This component was irrelevant to the current study but was included because
388 the experiment was part of a larger study where observers’ perception of the LM-AM
389 component was relevant. We have previously shown that the LM-AM combination is seen as
390 flat stripes in these plaid patterns and that the perception of the LM+AM component varies
391 little with the presence of this extra cue (Schofield et al., 2006; Schofield, Rock, Sun, Jiang &
392 Georgeson, 2010).

393

394 Haptic stimuli consisted of a virtual surface with sinusoidal undulations collocated with the
395 visual stimulus and presented at the same orientation and spatial frequency as the shading
396 signal. These stimuli were presented via a small force-feedback robot arm with a pen-like
397 stylus. The arm provided physical resistance whenever the observer tried to move the stylus
398 tip through the virtual surface. Observers held the stylus with their dominant hand and gently

399 stroked the virtual surface. The initial position of the surface relative to the visual stimulus
400 varied at random from trial to trial. Surface amplitude was fixed at ± 3 mm (amplitude = 0.12
401 of a wavelength). Three markers (not shown in Figure 3) were added to the visual stimulus:
402 one at fixation and two at opposite edges of the stimulus positioned such that the alignment
403 of the three markers indicated the direction in which observers should stroke the haptic
404 surface in order to feel the undulations. The position of the stylus tip was marked by a small
405 circle to provide visual feedback of the stylus location. Visual and haptic stimuli were
406 generated on a PC computer and observers adjusted the position of the haptic stimulus
407 using keys 4 and 6 on the computer's numeric keypad. Numbers placed next to the outer
408 markers in the visual stimulus indicated which key to press to move the haptic surface
409 towards each marker. The haptic surface could be moved in either 1.4 or 0.35 mm steps,
410 toggled as required by pressing key 5. Observers heard a long tone after each 1.4 mm
411 movement and a short tone after each 0.35 mm movement.

412

413 2.1.2 Equipment and calibration

414 Stimuli were presented in a modified ReachIN haptic workstation (ReachIN Technologies
415 AB, Stockholm, Sweden). The visual stimuli were presented on a 17" Sony Trinitron CPD
416 G200 CRT monitor mounted at an angle of 45° above a horizontal half-silvered mirror.
417 Haptic stimuli were presented via a Phantom-Desktop (SensAble Technologies Inc, Woburn,
418 MA) force feedback arm located beneath the mirror. Observers looked into the mirror at a
419 downward angle and thus perceived the visual stimulus to be beneath the mirror and
420 approximately perpendicular to the line of sight. The effective viewing distance was about 57
421 cm. Visual stimuli were calibrated against the monitor's gamma characteristic using look up
422 tables in a BITS++ graphics interface (CRS Ltd, Rochester, UK) which also served to
423 enhance the available grey level resolution to the equivalent of 14 bits. Values in the look up
424 tables were determined by fitting a four-parameter monitor model (Brainard, Pelli, & Robson,
425 2002) to luminance readings recorded with a CRS ColourCal photometer.

426

427 2.1.3 Observers

428 The 15 observers had normal or corrected to normal vision and, with the exception of
429 authors AJS & PR, were paid for their time and unaware of the purposes of the experiment.
430 They each undertook at least five observations at each orientation. Six of these observers
431 (from a pilot study) were not tested at orientations 15, 75, 105 & 165°. Observations were
432 made in a darkened room so that even though the mirror was half-silvered the observers
433 could not see their own hand through it. A hood was fitted to the monitor such that observers
434 could not view the screen directly. Head position was not physically constrained, but the
435 arrangement of the hood and the need to keep the haptic stimuli at a comfortable distance

436 for one's arm served to limit head position. Head orientation was not constrained but
437 observers were told to keep their heads upright. Viewing was binocular and so the visual
438 stimulus provided stereoscopic cues to flatness. However, we have previously shown
439 (Schofield, et al., 2006) that a robust percept of shape-from-shading can be derived from
440 such stimuli, and binocular presentation avoids the rivalry associated with monocular
441 presentation for some observers.

442

443 2.2 Analysis

444 The final position of the haptic surface was recorded at the end of each trial as was its offset
445 relative to the nearest luminance peak in the shading profile (see Figure 4). Positive offsets
446 (expressed as proportions of a wavelength) indicated that the perceived surface peak was
447 below the luminance peak (or to the right at 0°) for orientations from 0-165°. Data were
448 analysed by first taking medians (not all distributions were normal) then extrapolating the
449 recorded data to cover the full range of haptic directions from 0 to 360°. To do this we
450 exploited the fact that the orientation of the visual stimuli repeated every 180° whereas offset
451 direction repeats only every 360°. Hence, positive offsets in the range 180-345° would
452 correspond to a perceived surface peak above (to the left at 180°) of the luminance peak.
453 Thus the extrapolated data in the range 180-345° were set to the negative of those recorded
454 over the range 0-165°. This extrapolation is relevant for the modelling in Section 5.

455

456 2.3 Results and discussion

457 Figure 5 plots offsets between luminance peaks and perceived surface peaks (inter-peak
458 offsets) as a function of stimulus orientation for the nine observers who provided
459 observations at all orientations. There were considerable individual differences in behaviour
460 but strong common themes emerged. Inter-peak offsets varied with stimulus orientation and
461 typically ranged between 0 and ¼ wavelengths. The majority of observers (10 out of 15)
462 produced their largest offset at orientations close to horizontal (90 & 270°) and their smallest
463 offset close to vertical (0°) as exemplified by observers HS, HW, AS, AJS, AO & RCL. Five
464 observers produced their largest and smallest offsets at some other orientations (e.g. PJ &
465 AT). Based on the models described later we define the orientation orthogonal to each
466 person's maximum offset as their illuminant aligned orientation. This orientation generally
467 corresponds to a zero-crossing in the model offset traces of Figure 5 and is the point at
468 which the perceived ridges 'ran' towards the observer's preferred light source as estimated
469 by the model. Most observers perceived surface peaks to be below and to the right of the
470 luminance peaks (consistent with lighting from above the line of sight), but three placed their
471 surface peaks above the luminance peaks (e.g. AT, consistent with lighting-from-below).
472 Seven observers showed a smooth transition between their maximum and minimum offsets

473 (e.g. HW, HS, AT, AJS, AO) whereas the remainder had more abrupt transitions. For
474 example, all of PJ's offsets were close to $\frac{1}{4}$ wavelength; none were close to zero. The
475 maximum absolute offset for some observers was noticeably less than $\frac{1}{4}$ wavelength (eg,
476 HS & AO).

477

478 **Figure 5 about here (single column)**

479

480 We were worried about possible contamination from the convex/concave ambiguity. The
481 perceived surface may be more ambiguous at some orientations than others and flips in the
482 positions of perceived peaks could reduce offsets. If this were the case we would expect
483 standard deviations to increase with decreasing offsets and for orientation with low offsets to
484 have bi-modal distributions. We calculated the coefficient of bimodality

485 $[b = (1 + \text{skewness}(x)^2)/(\text{kurtosis}(x) + 3((n - 1)^2/((n - 2)(n - 3))))]$ for each observer
486 at each orientation where n is the number of observations and where kurtosis is defined as
487 being zero for a normal distribution. We then correlated this metric with offset magnitudes. If
488 the concave/convex ambiguity were a problem we would expect a *negative* correlation
489 between b and offset magnitudes. The mean correlation was significantly different from zero
490 on a one sample t-test but was positive ($\bar{r}=0.2$, $t=3.6$, $df=14$, $p=.003$) implying that offset
491 distributions tend to be bimodal when median offsets are large not small. Thus the
492 concave/convex ambiguity cannot have resulted in the reduced offsets recorded.

493

494 In the introduction we proposed inter-peak offsets as a means to assess the nature of
495 people's assumed light source. Point source interpretations should lead to $\frac{1}{4}$ wavelength
496 offsets, a diffuse lighting assumption will produce no inter-peak offset and a mixed lighting
497 assumption predicts offsets that depend on orientation. While some participants perceive
498 surface peak to be offset from luminance peaks by $\frac{1}{4}$ wavelength at some orientations no
499 offset was found at other orientations and some participants never perceived an offset as
500 large as $\frac{1}{4}$ wavelength. The similarity between perceived inter-peak offsets and the pattern
501 of physical inter-peak offsets observed for mixed illumination (Fig 3 & Model A, Section 5.1)
502 suggests that many people assume a mixed lighting model .These results support prediction
503 4 (Section 1.4).

504

505 **3. Experiment 2: Perceived depth magnitude does not vary with stimulus orientation.**

506 We were concerned to ensure that the magnitude of the perceived undulations did not vary
507 systematically with orientation, and that there was no association between inter-peak offset
508 and perceived depth. In particular, we wanted to verify that participants did not see illuminant

509 aligned stimuli as flat, as such a result might imply a failure to perceive shape-from-shading
510 at the given orientation.

511

512 *3.1 Method*

513 Seven observers (all naïve to the purpose of the experiment; six of whom had previously
514 taken part in Experiment 1) were presented with visual stimuli identical to those of
515 Experiment 1 and additional single oblique LM+AM stimuli (left side of Figure 3). They were
516 asked to adjust the amplitude of a collocated haptic surface to match that of the visually
517 perceived undulations. Haptic stimuli were aligned with the LM+AM components of plaid
518 stimuli and the offset between the haptic and visual stimuli was set to each observer's
519 preferred offset at the given orientation, as determined in Experiment 1. Surface depth was
520 adjusted in 2 or 0.5 mm steps by pressing keys on the keypad (8 for deeper, 2 for shallower,
521 and 5 to toggle between step sizes). Observers heard a long tone after each 2 mm
522 adjustment and a short tone after each 0.5 mm adjustment. Observers could not set
523 amplitude negative and were warned with a tone of any attempt to do so. The initial
524 amplitude was set to a random value in the range 0-8mm (mean to peak). Three visual
525 markers indicated the orientation along which to feel but the outer markers appeared without
526 numbers. All other experimental details were as Experiment 1.

527

528 *3.2 Results and discussion*

529 There was no systematic variation in perceived surface amplitude with orientation for either
530 plaid or single oblique stimuli (Figure 6). Importantly perceived depth amplitude did not
531 approach zero (flat) for any participant at any orientation. With the exception of AT (min
532 offset at 45°) and VC (did not participate in experiment 1), observers produced their smallest
533 *inter-peak offsets* (see Fig 5) for stimuli oriented close to 0°, but there is no sign of a
534 corresponding dip in *perceived depth amplitude* at 0° (or 45° for AT; Fig 6). To test for a
535 systematic relationship between perceived depth amplitude and absolute inter-peak offset
536 we correlated these judgments for the six participants who took part in both studies. A
537 positive correlation would indicate that when observers aligned surface peaks with
538 luminance peaks they also saw the stimulus as flat. With the exception of AS correlations
539 were either very weak or negative and the mean correlation across the six observers was
540 very weak and non significant ($\bar{r} = 0.009$ for the plaids and -0.03 for the single oblique
541 stimuli). Thus we conclude that our inter-peak offsets are valid at all orientations and
542 prediction 5 of Section 1.4 is rejected.

543

544 **Figure 6 about here (single column)**

545

546 **4. Experiment 3: Perceived frequency is constant with stimulus orientation.**

547 Experiment 1 had the limitation that observers could not adjust the frequency or shape of the
548 haptic surface to match that of the visually perceived surface. The frequency of the haptic
549 surface was always equal to that of the luminance signal and it was sinusoidal to match the
550 luminance variations. This was done so as to reinforce a sinusoidal surface interpretation
551 thus leaving the observers' internal illumination model as the only 'adjustable' parameter
552 available to them in making their interpretations. Although there is evidence that humans
553 readily perceive sinusoidal shading profiles as sinusoidal surfaces (Pentland, 1988;
554 Schofield et al., 2006; Supplementary data) it is possible that our use of a haptic match
555 stimulus forced observers into perceiving our stimuli in an unrealistic fashion. In particular,
556 they may have wanted to report some stimuli as having a frequency half that of the shading
557 as would be consistent with (say) vertical undulations lit from above (see Figure 1). In this
558 experiment we asked observers simply to mark the locations of perceived surface peaks and
559 troughs in the absence of any haptic cue to surface shape or frequency. Thus observers
560 were free to perceive the surface as non-sinusoidal and as having a frequency different from
561 that of the luminance signal.

562

563 *4.1 Method*

564 Six participants from Experiment 2 (excluding VC) and four new observers (SH, TP, LA, & IH
565 all with normal or corrected vision) viewed single, multiplicative (LM+AM), sinusoidal
566 luminance modulations of the textured surface (see left hand side of Figure 3). Two variants
567 of the experiment were conducted. The four new participants viewed stimuli in the haptic
568 workstation although the Phantom device was not used and the stimuli were displayed and
569 calibrated via a CRS-VSG2/5 graphics card. A modified hood which extended down to the
570 edge of the mirror was used. Observers looked through a slit in this hood and as a result the
571 viewing distance was reduced to 40cm. The seven observers from experiment 2 viewed
572 stimuli outside of the haptic workstation on a vertically oriented 21" Sony GDM F520 monitor.
573 The viewing distance was extended to maintain spatial frequency of the sinewave stimuli on
574 the larger monitor. Otherwise, the experimental setup was identical to Experiment 1.

575

576 Sinusoidal shading profiles (LM+AM alone, sf=0.4 c/deg, see Figure 3) were presented at 12
577 orientations in the range 0-165°. Observers were instructed to mark the positions of peaks
578 and troughs of the perceived surface by moving a red marker along a track defined by two
579 blue markers (lower panel of Figure 4; markers shown as white and black respectively in
580 print version). The position of the blue markers was chosen at random from trial to trial but
581 their spacing was fixed (2.04 cycles of modulation) and the track was always orthogonal to
582 the shading pattern. The red marker started 0.04 cycles away from one blue marker and

583 observers were told to mark features in order starting from this end of the range. Observers
584 were, however, allowed to track back and forth to home-in on features. The position of the
585 red marker was controlled using two bi-directional keys on a CRS CB3 button box (one each
586 for coarse and fine adjustments). The third key was pressed up to mark a peak and down for
587 a trough respectively. Thus, the direction of the marker key should have alternated on all
588 trials.

589

590 The position of each marked location was recorded relative to the luminance profile of the
591 visual stimulus. The distance between the marked features was also recorded. The data
592 were screened to remove trials where the direction of the marker key did not alternate (e.g.
593 where observers claimed to see two peaks without an intervening trough). The number of
594 trials that were ignored due to this screening process was very small.

595

596 *4.2 Results and discussion*

597 The point- and diffuse-source assumptions lead to two predictions for illuminant aligned
598 stimuli. A diffuse lighting interpretation would result in observers seeing a surface at the
599 same frequency as the shading signal. Point-source model would promote frequency halving
600 (undoing the frequency-doubling found for quadratic shading). Any perceptual flipping
601 between these interpretations would alter the mean peak-to-trough spacing and increase
602 standard deviations. The perceived distance between neighbouring peaks and troughs was
603 close to $\frac{1}{2}$ wavelength of the luminance signal at all orientations (Figure 7). Observers
604 always perceived the surface undulations to have the same spatial frequency as the
605 luminance signal regardless of stimulus orientation; there was no evidence for frequency
606 halving at any orientation. There is no evidence that standard deviation varied systematically
607 with orientation either suggesting that our observers saw a stable percept at all orientations.
608 These results confirm those of Experiment 1 and allow us to reject the prediction that
609 observers would perceive frequency halving at some orientations (prediction 1, Section 1.4).

610

611 Figure 8 shows the inter-peak offsets recorded for the four new observers in Experiment 3.
612 Offset profiles are similar to those of Experiment 1 confirming that the previous result was
613 unlikely to be due to the haptic matching method used. Comparing the results of the four
614 new participants with those of the seven participants from Experiment 2 we see that neither
615 past experience with the haptic task nor the exocentric inclination of the stimulus (backward
616 slant of 45° in the haptic workstation) affect either the peak-to-trough spacing or the pattern
617 of inter-peak offsets.

618

619 **Figures 7 and 8 about here (single column)**

620

621 **5. Modelling.**

622 We now propose two philosophically distinct models to explain our data. Noting that humans
623 do not solve shape-from-shading veridically (Pentland, 1988; Zhang et al., 1999) we do not
624 attempt to construct a machine vision algorithm to solve shape-from-shading to such
625 precision. Many machine vision algorithms exist and interested readers are directed to
626 Zhang et al. (1999) for an early review of such methods. It should be noted, however, that
627 most of these methods assume a collimated (point-like) light source of known direction and
628 require either iterative optimization of a cost function seeded with information such as
629 occluding boundaries or the iterative propagation of information from seed points in the
630 image such as intensity peaks. We avoided such methods because: (i) we are interested in
631 human performance (not veridical shape recovery), (ii) our stimuli lacked many of the
632 features that are required to make machine algorithms work, and (iii) we wanted to avoid
633 methods that assume point-source lighting.

634

635 5.1 Model A: Assumed mixed illuminant.

636 Model A starts from the assumption that the human visual system has developed to process
637 natural scenes and is thus optimised to a world that is mostly illuminated by a mixture of
638 point and diffuse lighting or at least upwardly baised diffuse lighting (Dror et al., 2004; Mury,
639 et al., 2009; Teller et al., 2001). The model assumes that human vision can, at least
640 approximately, invert the generative processes that produce shading on a surface given
641 some knowledge of the light source, and that when the stimulus provides few clues to the
642 light source composition a default illumination model is adopted in order that the inverse
643 generative process can function. We draw a parallel here with Langer & Bülthoff (2000) who
644 found that when a stimulus was rendered under point lighting observers mapped shape-
645 from-shading as if under point lighting whereas surfaces lit diffusely were mapped according
646 to a blur+dark-is-deep rule which is more appropriate for diffuse lighting.

647

648 In order to predict the default lighting adopted by each individual we generated the
649 luminance profiles of physical sinusoidal surfaces lit by mixtures of point and diffuse lighting
650 and then estimated the offset between the physical- and luminance-peaks by taking the
651 Fourier transform of the luminance profile and extracting the phase of the component equal
652 to the frequency of the original surface. This is equivalent to the blur imposed by Langer &
653 Bülthoff. We then used MatLab's *fmincon* function, which finds the optimal constrained
654 parameters for an arbitrary, user defined model with a user defined cost function (we used
655 sum of squared errors as our cost function), to find the direction of the assumed point light
656 source and balance between point and diffuse lighting that produced offset profiles that best

657 matched those for each observer (see Fig 5). These lighting models being fixed for all
658 stimulus orientations.

659

660 The luminance profiles used to derive these fits were generated from Eqn 2 with two
661 additional terms to describe the contribution from the diffuse source.

662

663

$$L \approx (1 - \gamma) \left(\cos\left(\frac{\pi}{2} - e\right) - 0.12 \sin(x) \cos\left(\frac{\pi}{2} - \varphi + \lambda\right) \sin\left(\frac{\pi}{2} - e\right) \right. \\ \left. - \cos\left(\frac{\pi}{2} - e\right) (0.12^2 \sin^2(x))/2 \right) + \gamma(0.065 \cos(x) + 0.045 \cos(2x))$$

664 Eqn5

665 Where γ described the balance between point and diffuse lighting (high γ = diffuse), φ is the
666 orientation of the surface corrugations (positive = clockwise), λ is the direction of the default
667 point source (positive = anti-clockwise) and the constants were appropriate to our stimuli (ie
668 surface depth was 0.12 wavelengths and 0.065 and 0.045 provide the appropriate weighting
669 for the first and second harmonics of the diffuse source modelled as described in the legend
670 of Fig 2), ω is omitted as we assume it equal to 1. We further assumed elevation (e) = 30°.

671

672 The left hand side of Table 1 shows the SSE and parameter values for each observer. The
673 model produces a relatively good fits to the data although data from those observers having
674 a smooth offset profile with a high peak were not fit well. Resulting fits are shown by the solid
675 lines in Figs 5 and 8.

676

677 [Table 1 about here](#)

678

679 5.2 Model B: Mixed processing model.

680 Model A relies on the observer being able to at least approximately invert the generative
681 process in order to estimate shape-from-shading; it does not articulate a means by which
682 this is achieved. Given that shape-from-shading estimates are often not veridical this
683 inversion seems unlikely. We now present an alternative, mechanism driven, account of our
684 data.

685

686 5.2.1 Outline

687 Model B starts with the assumption that humans process all stimuli with two shape-from-
688 shading modules whose output is then combined in a stimulus specific way. This
689 combination could be the result of flipping between two hypothesised surface shapes but,
690 given our data, we think a linear combination of the two hypothesised surface is more likely.

691
692 We implement a cut-down version of Pentland's (1988) model for human shape-from-
693 shading which produces a linear mapping between luminance and perceived surface shape
694 with a $\frac{1}{4}$ wavelength offset. We augment this model with a version of Langer & Bulthoff's
695 (2000) blur+dark-is-deep model. These two models give a reasonable account of human
696 shape-from-shading under point- and diffuse-lighting assumptions respectively. Critically
697 when presented with sinusoidal shading patterns they will both produce sinusoidal depth
698 profiles but Pentland's (1988) model will shift the perceived surface peak by $\frac{1}{4}$ wavelength
699 relative to the luminance peak³ whereas the two peaks will align in the output of Langer and
700 Bulthoff's (2000) model. The principal innovation of this model is to combine the two
701 approaches above such that, when fit to the data, the balance between the point- and
702 diffuse-lighting interpretations can be inferred. The use of sinusoidal stimuli greatly simplifies
703 the model. Because each sub-module will produce a sine wave output we need not
704 implement the models in full but can simulate their action with appropriately phase shifted
705 sine waves.

706

Figure 9 about here (double column)

708

709 The upper arm of Model B also provides an estimate of the lighting direction..Note that
710 although sinusoidal shading is highly ambiguous with relation to the direction of the light
711 source Koenderink et al. (2007) have shown that anisotropic shading patterns give rise to
712 very stable estimates of illumination direction up to 180° flips. In this case people estimate
713 the light direction to be orthogonal to the dominant orientation in the shading pattern. In
714 practice there are two directions orthogonal to the dominant orientation in each stimulus;
715 both are equally valid estimates and we deal with this ambiguity in section 5.2.2.

716 The two surface interpretations are combined in a weighted sum. Each arm has a weight (β
717 and $1-\beta$ for the diffuse and point interpretations respectively) and β is fixed for each
718 observer. The point interpretation has an additional variable weight which depends on the
719 observer's estimate of the likelihood that illumination will come from the direction specified
720 by the stimulus. If an observer had a preference for lighting from above (say) this would be
721 expressed as a strong weight for vertical lighting and a weak weight for horizontal lighting.

722

723 When presented with a sinusoidal input the two arms of this model will produce sinusoidal
724 surface profiles: one offset by $\frac{1}{4}$ wavelength from the luminance profile, the other having no
725 offset. The weighted sum of two sine waves with the same frequency but different phases is

³ A model based on integration (surface gradients proportional to luminance) would also produce a $\frac{1}{4}$ wavelength offset.

726 a sine wave with intermediate phase. For the purposes of our fits the model is described by
727 the following simple equation $s = \cos(x + \arctan(\cos(\theta - \lambda) \cdot (1 - \beta) / \beta))$, where θ is
728 orthogonal to the dominant orientation in the stimulus - and λ is the preferred lighting
729 direction which, unlike θ , indexes anticlockwise. Thus the inter-peak offset predicted by the
730 model will be closer to whichever of the two interpretations carries the stronger weight. As
731 the relative weighting of the two components varies with stimulus orientation so does the
732 inter-peak offset.

733

734 *5.2.2 Direction-dependent weighting function*

735 We made the variable weighting function sinusoidal such that a negative weight would be
736 assigned if the illumination was predicted to come from the direction opposite to the
737 observer's preferred direction. Recall that there are two directions orthogonal to a given
738 stimulus orientation and both are candidates for the perceived lighting direction. In our
739 framework lighting from one direction would produce a positive inter-peak offset relative to its
740 own direction whereas lighting from the opposite direction would produce a negative inter-
741 peak offset relative to its own direction – lighting direction repeats every 360°. However, in
742 terms of the stimulus both the predicted offsets will be in the same direction because
743 stimulus orientation repeats every 180°. The weighting function can thus produce negative
744 weights and hence negative offsets allowing us to model the data of participants like AT. The
745 orientation with the most positive weight is deemed to be the observer's preferred lighting
746 direction and this is adjusted to fit the data best.

747

748 *5.2.3 Analysis and Results*

749 We fit (using *fmincon*) the model to the individual data from Experiments 1 & 3 dashed-lines
750 in Figures 5 and 8). Note that the dashed lines exactly overlay the predictions from Model A.
751 Model parameters and SSE's for all observers are shown on the right hand side of Table 1.
752 Note that the two models produce nearly identical λ 's and SSE value; although β and γ are
753 not identical they are perfectly correlated. These results strongly suggest that the two
754 models are mathematically equivalent, a fact we prove in Supplementary section S3.

755

756 Model B (see Fig 9) implies that perceived depth amplitude will vary with orientation because
757 the amplitude of s depends on w . We did not observe any such variation which is a
758 challenge to the specific form of the model although this anomaly can be reconciled by
759 supposing that the amplitude of surface s is normalised to the stimulus contrast.

760

761 *5.3 Interpretation*

762 The parameters for Model A should be interpreted as follows. The direction of the observers
763 preferred light source (but not its elevation) is given by λ (Note that lighting direction is
764 indexed anti-clockwise whereas stimulus orientation indexes clockwise) and the balance
765 between the diffuse and point components is determined by γ such that high γ suggests a
766 mostly diffuse default illuminant. For Model B λ again the observers preferred light source
767 and β is the weighting term. A high γ (β) means that the observer prefers a diffuse source
768 interpretation. A low γ (β) implies that a point source is preferred when viable.

769

770 Low γ or β also results in a flat-topped model offset profile with an abrupt transition between
771 extreme offsets (eg. observer PDJ in Figure 5). High γ or β results in smoother transitions
772 and a lower maximum offset. Note that where an observer's maximum inter-peak offset is
773 large (close to $\frac{1}{4}$ wavelength) the models will tend to prefer a low γ or β , resulting in abrupt
774 transitions. Thus observers with a large maximum offset but smooth transitions present a
775 challenge to the models. For Model B at least data from these observers might be better fit
776 by assuming a weighting function of a different shape. There was considerable variation in γ
777 and β across participants, suggesting that some preferred the diffuse lighting interpretation
778 more than others.

779

780 *5.4 Dealing with plaid stimuli*

781 The above models consider only stimuli comprising single sine wave luminance profiles. The
782 stimuli used in Experiment 1 were more complex plaid stimuli in which one orientation
783 faithfully represented multiplicative shading (LM+AM) while the other did not (LM-AM). We
784 have shown elsewhere that observers treat LM-AM as if it were a flat reflectance change
785 (Schofield et al, 2006; Schofield et al., 2010). Layer segmentation – the separation of
786 components into shading vs reflectance changes – is a complex issue in itself but humans
787 seem to be able to perform such a separation (Kingdom, 2008). We assume that layer
788 segmentation takes place before shape-from-shading such that our plaid stimuli present
789 themselves as single sine waves as far as shape-from-shading is concerned. Elsewhere we
790 propose a model for how layer segmentation is achieved in LM/AM plaids (Schofield, et al.,
791 2010). Layer segmentation will also separate the noise textures in our stimuli from the
792 shading patterns.

793

794 *5.5 Convex/concave ambiguities and high γ (β)*

795 In the introduction we outlined five predictions. One (prediction 3) concerned the
796 convex/concave ambiguity and the possibility of perceptual flipping between two equally
797 likely surface profiles. We noted that this would predict small inter-peak offsets for some
798 stimuli but that it was also make the data these orientations bimodal. We discounted this

799 hypothesis in section 2.3 because we found a positive relationship between absolute offsets
800 and bimodality. That is, bimodality was associated with large offsets not small offsets.
801 However, the analysis of Experiment 1 merged the results from observers such as PDJ with
802 abrupt transitions and those of observers such as HS with smooth transitions. Is it possible
803 that only those with smooth transitions suffer perceptual flipping and that this explains the
804 smoothness of their offset data? We reasoned that if people with smooth transitions (high γ)
805 suffered perceptual flipping more than those with abrupt transitions (low γ) then the individual
806 offset-bimodality correlations measured in Experiments 1 and 3 should themselves correlate
807 negatively with γ . Recall that offset-bimodality correlations will be negative if perceptual
808 flipping occurs at orientations with small offsets. Although this relationship was negative it
809 was relatively weak and not statistically significant ($r=-.26$, $df=17$, $p=0.283$). We also tested
810 the correlation between people's mean coefficient of bimodality and γ which should be
811 positive if flipping/bimodality is the cause of smooth offset profiles (high γ). This correlation
812 was very weak negative and not significant ($r=-.0051$, $df=17$, $p=.84$). Finally we measured
813 the correlation between γ and coefficients of bimodality associated with individual's smallest
814 offsets. Again this should be positive if perceptual flipping is causing smooth offset profiles; it
815 was not ($r=-.11$, $df=17$, $p=.64$). We conclude, as in section 2.3, that that the concave/convex
816 ambiguity was not responsible for producing the smooth offset profiles and high γ values
817 noted in our data.

818

819 **6. Discussion.**

820 People perceive sinusoidal luminance shading as a sinusoidal surface undulating at the
821 same spatial frequency as the luminance profile (see Schofield, et al., 2006, Pentland 1988,
822 Experiment 3, and supplementary file); dismissing prediction (1), see section 1.4.

823

824 Perceived inter-peak offsets varied systematically with orientation. This result is not
825 consistent with the assumption of a single, pure point source (even one with variable
826 direction; prediction 2), since that would predict no change in inter-peak offset with stimulus
827 orientation. This finding is not consistent with a fully diffuse light source either, since that
828 would predict no offset at any orientation. The variation in inter-peak offsets was not
829 accompanied by a reduction in perceived depth amplitude, nor was it due to perceptual
830 flipping between multiple, equally likely, surface interpretations; so predictions 3 and 5 are
831 also dismissed.

832

833 Our data can be modelled by assuming the observer is able to, at least approximately,
834 reverse the image generation process using a mixed, but fixed, internal lighting assumption
835 (Model A) or that they generate two surface interpretations which are linearly combined with

836 weights determined by the stimulus (Model B). In either case model fits suggest that
837 observers adopt a mixture of point and diffuse lighting. The two models are mathematically
838 equivalent for sine wave shading patterns so our data cannot discriminate between them.

839

840 Our finding that observers seem to adopt a mixed point and diffuse lighting model is
841 consistent with the results of lightness judgements found by Bloj, Ripamonti, Mitha, Hauck,
842 Greenwald & Brainard (2004). A mixed lighting model is also consistent with the data on
843 natural illumination which show a largely diffuse illumination with an upward bias – that is a
844 combination of diffuse and directional components (Dror, et al., 2004; Mury et al, 2009;
845 Teller, et al., 2001). It would make sense if humans adopted an illumination model which
846 was close to the naturally occurring illumination profile. The inclusion of an explicit point light
847 source (rather than a more general upward bias as might be more common in natural
848 settings; Mury et al, 2009) facilitated matches to individual data. It seems likely that
849 individuals have a preferred lighting direction that is generally from above but which varies
850 between observers and can be modified by experience (Adams et al., 2004; Mamassian &
851 Goutcher, 2001; Sun & Perona, 1998). This suggests to us that a discrete point component
852 rather than a general upward bias is appropriate.

853

854 Model B is limited to the understanding of sinusoidal shading. It could be expanded to deal
855 with (that is, ignore) reflectance changes by the inclusion of a preceding layer-segmentation
856 stage (see for example that proposed by Schofield et al., 2010). It might also be extended to
857 more complex natural patterns by implementing the linear shading (Pentland, 1988) and
858 blurred dark-is-deep (Langer and Bülthoff, 2000) models in full. A method based on
859 Pentland's (1982) model for finding the illumination direction could serve to expand the
860 illuminant direction estimation process to more natural images (see also Gerhard & Maloney,
861 2010). It would be interesting (but beyond the scope of the current paper) to test such a
862 model against human performance for more complex scenes. Pentland's (1988) model alone
863 does reasonably well in such situations.

864

865 Model A is similarly limited to surfaces with uniform albedo and can also be augmented by a
866 preceding layer-segmentation stage. In theory this model can deal with any type of surface
867 however in practice any implementation would require that the image generation process be
868 inverted. This amounts to solving the shape-from-shading problem given an assumed light
869 source which may prove pragmatically difficult for the general case.

870

871 Our use of sinusoidal shading patterns may over-emphasise the diffuse lighting component.
872 Our stimuli contain no sharp edges that might indicate hard shadows and thus the presence

873 of a point light source. Further, our stimuli may promote the perception of a Lambertian
874 surface with little or no specular component. Images with more obvious specular highlights
875 may require a different interpretation from the one outlined here. However, our models are
876 more generally applicable if we allow the γ (or β) to vary with stimulus content such as hard
877 edges and specular highlights. Schemes such as Freeman's (1994) generic view framework
878 might serve to adjust γ in more complex scenes if diffuse lighting were included as a
879 candidate lighting model. Non-sinusoidal shading profiles, especially those with occlusions,
880 might indicate harsher – more point-like – lighting, giving the point source component of the
881 model a greater weight. We note that people are rather good at estimating the true light field
882 in well articulated, object rich, scenes (Koenderink et al., 2007) and that in such cases
883 internal lighting biases may not apply at all. However, while some stimulus types might
884 provide little evidence that there is a diffuse component within the human default lighting
885 model, we argue that the most general model must contain such a component.

886

887 A potential criticism of our method is that people may not actually perceive our stimuli as
888 conveying realistic depth percepts but might rather associate luminance variations with
889 depth in order to follow the instructions given; an experimenter effect. We reject this for three
890 reasons. First, it is unlikely that all of the naive participants would adopt the same ‘false’
891 association between luminance and depth to please the experimenter and that *none* would
892 set their depth / gradient estimates to zero if they in fact saw our stimuli as flat. Second, if
893 observers had adopted a simple association between luminance and depth we think it
894 unlikely that this association would vary systematically with stimulus orientation. Third, we
895 have shown elsewhere (Schofield et al., 2010) that observers see the LM-AM components of
896 our plaid stimuli as flat. The LM-AM and LM+AM components contain the same luminance
897 variation and the AM sub-components create relatively subtle variations in pixel-wise
898 luminance values. Therefore it seems unlikely that a ‘false’ association between luminance
899 and depth would be applied to LM+AM stimuli alone and much more likely that observers
900 genuinely perceive LM+AM as conveying depth.

901

902 7. Conclusion

903 People perceive sinusoidally corrugated luminance patterns as sinusoidal surfaces of the
904 same spatial frequency as the luminance waveform. In general perceived surface peaks are
905 offset from the luminance peaks and these inter-peak offsets vary with stimulus orientation.
906 This result is not consistent with an internal lighting model that is either a pure point source
907 or fully diffuse illumination, but it is consistent with a weighted mixture of the two lighting
908 types. Such as mixed illumination model is consistent with everyday experience of biased

909 diffuse illumination as found on a cloudy day, in the illumination field of a typical room with
910 an overhead light and light-coloured walls, or from the sun embedded in a diffusing sky.

911

912 Acknowledgments

913 This work was funded by EPSRC grants GR/S07254/01 & EP/F026269/1 to AJS &
914 GR/S07261/01 to MAG. EPSRC took no part in the study or manuscript preparation. We
915 thank Harriet Allen and Martin Banks for commenting on early drafts of the manuscript, and
916 Mike Landy and an anonymous reviewer for their helpful suggestions.

917

918 *Author contributions.* AJS devised the study, collected some of the data for experiments (1
919 &2), devised and implemented experiment 3 and the control experiment, conducted much of
920 the analysis and was the principle author of the manuscript. PBR designed and implemented
921 experiments 1 and 2 and collected data for them. MAG advised on study design and made
922 considerable revisions to the final manuscript.

923

924 References

- 925 Adams, W.J., Graf, E.W., & Ernst, M.O. (2004). Experience can change the 'light-from-
926 above' prior. *Nature Neuroscience*, 7, 1057-1058.
- 927
- 928 Belhumeur, P.N., Kriegman, D.J., Yuille, A.L., (1999) The Bas-Relief Ambiguity, *International
929 Journal of Computer Vision*, 35, 33-44.
- 930
- 931 Bloj, M., Ripamonti, C., Mitha, K., Hauck, R., Greenwald, S., & Brainard, D. H. (2004). An
932 equivalent illuminant model for the effect of surface slant on perceived lightness. *Journal of
933 Vision*, 4(9):6, 735-746, <http://journalofvision.org/4/9/6/>, doi:10.1167/4.9.6.
- 934
- 935 Brainard, D. H., Pelli, D.G., & Robson, T. (2002). Display characterization. In *Encyclopedia
936 of Imaging Science and Technology*. J. Hornak (ed.), pp.172-188, Wiley.
- 937
- 938 Brewster, D. (1826). On the optical illusion of the conversion of cameos into intaglios, and
939 intaglios into cameos, with an account of other analogous phenomena. *Edinburgh Journal
940 of Science*, 4, 99-108.
- 941
- 942 Cavanagh, P., & Leclerc, Y.G. (1989) Shape from shadows. *Journal of Experimental
943 Psychology: Human Perception and Performance*, 15, 3-27.
- 944

- 945 Christou, C.G., & Koenderink, J.J. (1997). Light source dependence in shape from shading.
946 *Vision Research*, 37, 1441-1449.
- 947
- 948 Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world
949 illumination. *Journal of Vision*, 4(9):11, 821-837, <http://journalofvision.org/4/9/11/>,
950 doi:10.1167/4.9.11.
- 951
- 952 D'Zmura, M., (1991) Shading ambiguity: Reflectance and Illumination, in *Computational*
953 *models of visual processing*, M.S. Landy, & J.A. Movshon (eds), pp 187-207, MIT Press,
954 Cambridge, MA.
- 955
- 956 Erens, R.G.F., Kappers, A.K.L., & Koenderink, J.J. (1993). Perception of local shape from
957 shading. *Perception & Psychophysics*, 54, 145-156.
- 958
- 959 Freeman, W.T. (1994). The generic viewpoint assumption in a framework for visual
960 perception. *Nature*, 368, 542-545.
- 961
- 962 Gerhard, H.E., Maloney, L.T. (2010). Estimating changes in lighting direction in binocularly
963 viewed three-dimentional scenes. *Journal of Vision*, 10(9):14,
964 <http://www.journalofvision.org/content/10/9/14>
- 965
- 966 Hess, E.H. (1950). Development of chicks' responses to light and shade cues to depth.
967 *Journal of Comparative and Physiological Psychology*, 43, 112-122.
- 968
- 969 Kingdom, F.A.A. (2003) Colour brings relief to human vision. *Nature Neuroscience*, 6, 641-
970 644.
- 971
- 972 Kingdom, F.A.A. (2008) Perceiving light versus material. *Vision Research* 48, 2090-2105.
- 973
- 974 Kleffner, D.A. & Ramachandran, V.S. (1992). On the perception of shape from shading.
975 *Perception & Psychophysics*, 52, 18-36.
- 976
- 977 Koenderink, J.J., & Pont, S.C. (2003). Irradiation direction from texture. *Journal of the*
978 *Optical Society of America A – Optics Image Science and Vision*, 20, 1875-1882.
- 979
- 980 Koenderink, J.J., Pont, S.C., van Doorn, A.J., Kappers, A.M.L., & Todd, J.T. (2007). The
981 visual light field. *Perception*, 36, 1595-1610.

- 982
983 Koenderink, J.J., van Doorn, A.J., Kappers, A.M.L., Pas, S.F.T., & Pont S.C. (2003).
984 Illumination direction from texture shading. *Journal of the Optical Society of America A –*
985 *Optics Image Science and Vision*, 20, 987-995.
986
987 Koenderink, J.J., van Doorn A.J., & Pont, S.C. (2004). Light direction from shad(ow)ed
988 random Gaussian surfaces. *Perception*, 33, 1405-1420.
989
990 Koenderink, J.J., van Doorn A.J., & Pont, S.C. (2007). Perception of illuminance flow in the
991 case of anisotropic rough surfaces. *Perception & Psychophysics*, 69, 895-903.
992
993 Langer, M.S., & Bülthoff, H.H. (2000). Depth discrimination from shading under diffuse
994 lighting. *Perception*, 29, 649-660.
995
996 Langer M.S. & Zucker, S.W. (1997). Shape-from-shading on a cloudy day. *Journal of the*
997 *Optical Society of America A – Optics Image Science and Vision*, 11, 467-478.
998
999 Liu, B., & Todd, J.T., (2004). Perceptual biases in the interpretation of 3D shape from
1000 shading. *Vision Research*, 44, 2135-2145.
1001
1002 Mamassian, P., & Goutcher, R. (2001). Prior knowledge on the illumination position.
1003 *Cognition*, 81, B1-B9.
1004
1005 Mury, A.A., Pont, S.C., Koenderink, J.J. (2009). Structure of light fields in natural scenes.
1006 *Applied Optics*, 48, 5386-5395.
1007
1008 Norman, J.F., Todd, J.T., & Orban, G.A. (2004). Perception of three-dimensional shape from
1009 specular highlights, deformations of shading, and other types of visual information.
1010 *Psychological Science*, 15, 565-570.
1011
1012 Pentland, A. (1988) Shape Information From Shading: A Theory About Human Perception.
1013 *Second International Conference on Computer Vision*, 404-413.
1014
1015 Pentland, A.P. (1982) Finding the illuminant direction. *Journal of the Optical Society of*
1016 *America*, 72, 448-455.
1017
1018 Ramachandran, V.S. (1988). Perception of shape-from-shading. *Nature*, 331, 163-165.

- 1019
- 1020 Rittenhouse, D. (1786). Explanation of an optical deception. *Transactions of the American*
1021 *Philosophical Society*, 2, 37-42.
- 1022
- 1023 Schofield, A.J., Hesse, G., Rock P.B., & Georgeson, M.A. (2006). Local luminance amplitude
1024 modulates the interpretation of shape-from-shading in textured surfaces. *Vision Research*,
1025 46, 3462-3482.
- 1026
- 1027 Schofield, A.J., Rock, P.B., Sun, P., Jiang, X., & Georgeson. (2010). What is second-order
1028 vision for? Discriminating illumination versus material changes. *Journal of Vision*, 10(9): 2,
1029 <http://www.journalofvision.org/content/10/9/2>, doi: 10.1167/10.9.2.
- 1030
- 1031 Stewart, A.J., & Langer, M.S. (1997). Towards accurate recovery of shape from shading
1032 under diffuse lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19,
1033 1020-1025.
- 1034
- 1035 Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience*, 1, 183-184.
- 1036
- 1037 Teller, S., Antone, M., Bosse, M., Coorge, S., Jethwa, M., & Masters, N. (2001). Calibrated,
1038 registered images of an extended urban area. *Proceedings of the IEEE computer Society*
1039 *Conference on Computer Vision and Pattern Recognition*, Kauai, Hawaii.
- 1040
- 1041 Todd, J.T., & Mingolla, E. (1983). Perception of surface curvature and direction of
1042 illumination from patterns of shading. *Journal of Experimental Psychology: Human*
1043 *Perception and performance*, 9, 583-595.
- 1044
- 1045 Tyler, C.W. (1998). Diffuse illumination as a default assumption for shape-from-shading in
1046 graded images. *Journal of Image Science and Technology*, 42, 319-325.
- 1047
- 1048 von Fieandt, K. (1949). The phenomenological problem of light and shadow. *Acta*
1049 *Psychologica*, 6, 337-357.
- 1050
- 1051 Wijntjes, M.W.A., Volcic, R., Pont, S.C., Koenderink, J.J., & Kappers, A.M.L. (2009). Haptic
1052 perception disambiguates visual perception of 3D shape. *Experimental Brain Research*, 193,
1053 639-644.
- 1054

1055 Zhang, R., Tsai, P-S., Cryer, J.E., Shah, M. (1999) Shape from shading: A survey. *IEEE*
1056 *Transactions on Pattern Analysis and Machine Intelligence*, 21, 690-706.

1057

1058

1059 Figure and Table Legends

1060

1061 Figure 1. A: Luminance profiles for sinusoidal surfaces under a point-source illuminant. Here
1062 and throughout the paper we take point-source to mean a highly concentrated but distant
1063 light source. Outer ring: example rendered stimuli. Inner ring: surface profiles (thin lines) and
1064 shading (luminance) profiles (thick lines) for each rendering. Sub-plot the x-axes represent
1065 position along the surface; y-axes represent luminance (thick lines) or height (thin lines).
1066 Sub-plots trace surface depth and luminance from left to right working along a lines from A to
1067 B in the images of the outer ring. This configuration is counter intuitive for some plot pairs
1068 but maintains a common reference frame. Note that as surface orientation repeats every
1069 180° , plots on the left of the figure mirror those on the right. Values of ψ indicate the offset
1070 between physical- and luminance peaks in wavelengths. The polar location of each inner
1071 plot represents the orientation of the surface (polar angle, see panel C) and elevation of the
1072 light source relative to the centre of the surface (radial distance from centre) where frontal
1073 lighting would be represented by a plot at the centre of the figure and oblique lighting by a
1074 plot on the outer circle. For example, the top most image and associated plot represent a
1075 vertical surface lit from above, the images and plots at 90 and 270° represent horizontal
1076 stimuli also lit from above. The inner plots are located at a distance from the centre of the
1077 figure appropriate for the elevation of the light source used in each rendering. These were
1078 45° for surfaces oriented at 90 and 270° , 41° for surfaces at $60, 120, 240$ and 300° , 27° for
1079 surfaces at $30, 150, 210$ and 330° , and 30° for surfaces at 0 and 180° . With the exception of
1080 0 and 180° surfaces, light elevations were chosen to avoid occlusions and double-crested
1081 peaks (see supplementary file). The depth amplitude (mean-to-peak) of the surfaces was
1082 0.12 of the undulation wavelength, matching that used in experiment 1. Rendered images
1083 are from *PoVRay* (Persistence of Vision Raytracer Pty. Ltd), and traces from *MatLab*, both
1084 assuming Lambertian shading. B: Lighting diagram depicting variable lighting direction d. C:
1085 Lighting diagram showing variable orientation for the physical surface.

1086

1087 Figure 2. Sinusoidal surface under diffuse illumination. a) surface as described in Figure 1
1088 rendered (using *PovRay*) under a spherical diffuse illumination model consisting of a
1089 spherical array of 400, randomly but evenly spaced light sources. Minor fluctuations in gray
1090 level are due to the sampling process. b) surface (dashed line, left axis) and luminance (solid
1091 line, right axis) traces for the central 2 cycles of a similar surface rendered in *MatLab* under a
1092 diffuse source sampled at 1568 random positions in front of the surface. The strength of
1093 each light in the latter diffuse model was $1/1568$ th of that for the source of Figure 1.

1094

1095 Figure 3. As figure 1 except rendered images are sub-plot luminance traces now show the
1096 case of mixed diffuse and point source lighting (weighted 0.75 diffuse, 0.25 point).
1097 Figure 4. Extracts from example stimuli arranged as Figure 1. To save space only the
1098 cardinal and 45° oblique orientations are shown. Stimuli labelled with dashed lines (top and
1099 right hand column) are from Experiments 1 and 2 where plaid stimuli were used those with
1100 solid lines (bottom and left hand column) from Experiments 2 and 3 where single gratings
1101 were used. Radial lines have been labelled to reflect stimulus orientation in the range 0–
1102 180°. These stimuli have been cropped for presentation, un-cropped versions are shown in
1103 supplementary Figure S3. No gamma correction has been applied to these stimuli but the
1104 noise contrast has been exaggerated to aid visualisation. However, despite this
1105 manipulation, we note that the example stimuli do not provide an especially good
1106 representation of the appearance of our stimuli within the lab setting. Specifically we are
1107 aware that people find it harder to perceive depth in our stimuli when presented in paper
1108 form than is the case during experiments. See Section 4.1 for a description of the markers
1109 on bottom panel (coloured white and blank in print but red and blue in the experiment and in
1110 the online version). The stimuli are best viewed online at 200% magnification.

1111

1112 Figure 5. Example inter-peak offset data from nine observers as a function of stimulus
1113 orientation (measured clockwise from vertical). Circles represent recorded data; triangles are
1114 extrapolated data (see Analysis). Lines represent model fits (see section 5). Error bars are
1115 standard deviations. AJS was an author.

1116

1117 Figure 6. Perceived surface amplitude measured as the distance between the zero-
1118 crossings of the haptic sine wave (dc position of the surface) and the haptic surface peaks in
1119 the direction normal to the surface plane. Traces show results for individual observers as a
1120 function of orientation. AS was not an author.

1121

1122 Figure 7. Peak-to-trough spacing versus stimulus orientation. Data points show the mean
1123 spacing between neighbouring peaks and troughs in the perceived surface at different
1124 orientations averaged across repeated trials. Data for observers TP to HS have been shifted
1125 vertically in integer steps for clarity. Error bars represent 95% confidence intervals.

1126

1127 Figure 8. Inter-peak offset data from Experiment 3. Details as Figure 5.

1128

1129 Figure 9. A) Lighting diagram showing orientation of stimulus ϕ and orientation of default
1130 point light source λ for Model A. B. Schematic diagram of Model B simplified for the case of
1131 sine-wave stimuli. Estimates of surface shape are predicted for a point source interpretation

1132 (upper arm) and a diffuse source interpretation (lower arm). These are combined with a
1133 variable weight determined by an estimate of the lighting direction.

1134
1135

Person	Model A			Model B		
	SSE _A	γ	Preferred light source (λ_A)	SSE _B	β	Preferred light source (λ_B)
AJS	0.005	0.45	6	0.005	0.34	6
PDJ	0.048	0.21	-25	0.048	0.14	-25
RCL	0.013	0.41	5	0.013	0.3	5
AO	0.033	0.34	6	0.033	0.25	6
HW	0.071	0.23	5	0.071	0.16	5
AS	0.081	0.64	-12	0.081	0.53	-12
HS	0.005	0.54	3	0.005	0.42	3
AT	0.061	0.33	132	0.061	0.23	132
KU	0.056	0.78	151	0.053	0.67	129
AC	0.004	0.69	31	0.004	0.59	31
JG	0.019	0.38	-13	0.019	0.28	-13
PS	0.027	0.5	6	0.027	0.38	6
PR	0.006	0.43	-4	0.006	0.32	-4
MH	0.049	0.36	-153	0.049	0.26	-153
SW	0.056	0.64	-46	0.056	0.52	-46
SH	0.012	0.09	13	0.012	0.06	13
TP	0.019	0.53	-5	0.019	0.41	-5
LA	0.009	0.54	-3	0.009	0.43	-3
IH	0.015	0.28	-7	0.015	0.2	-7

1136

1137

1138 Table 1.

1139 Model fit parameters. Model A, left-hand side: columns show the sum of squared errors
1140 between the modelled and observed inter-peak offsets, the weight (γ) applied to the diffuse
1141 interpretation, and the observers' preferred lighting direction (λ). Model B, right-hand side:
1142 SSE, weight β and λ . Lighting direction is given as positive = anti-clockwise shift from
1143 vertical. AJS and PR are authors.

1144

1145

1146

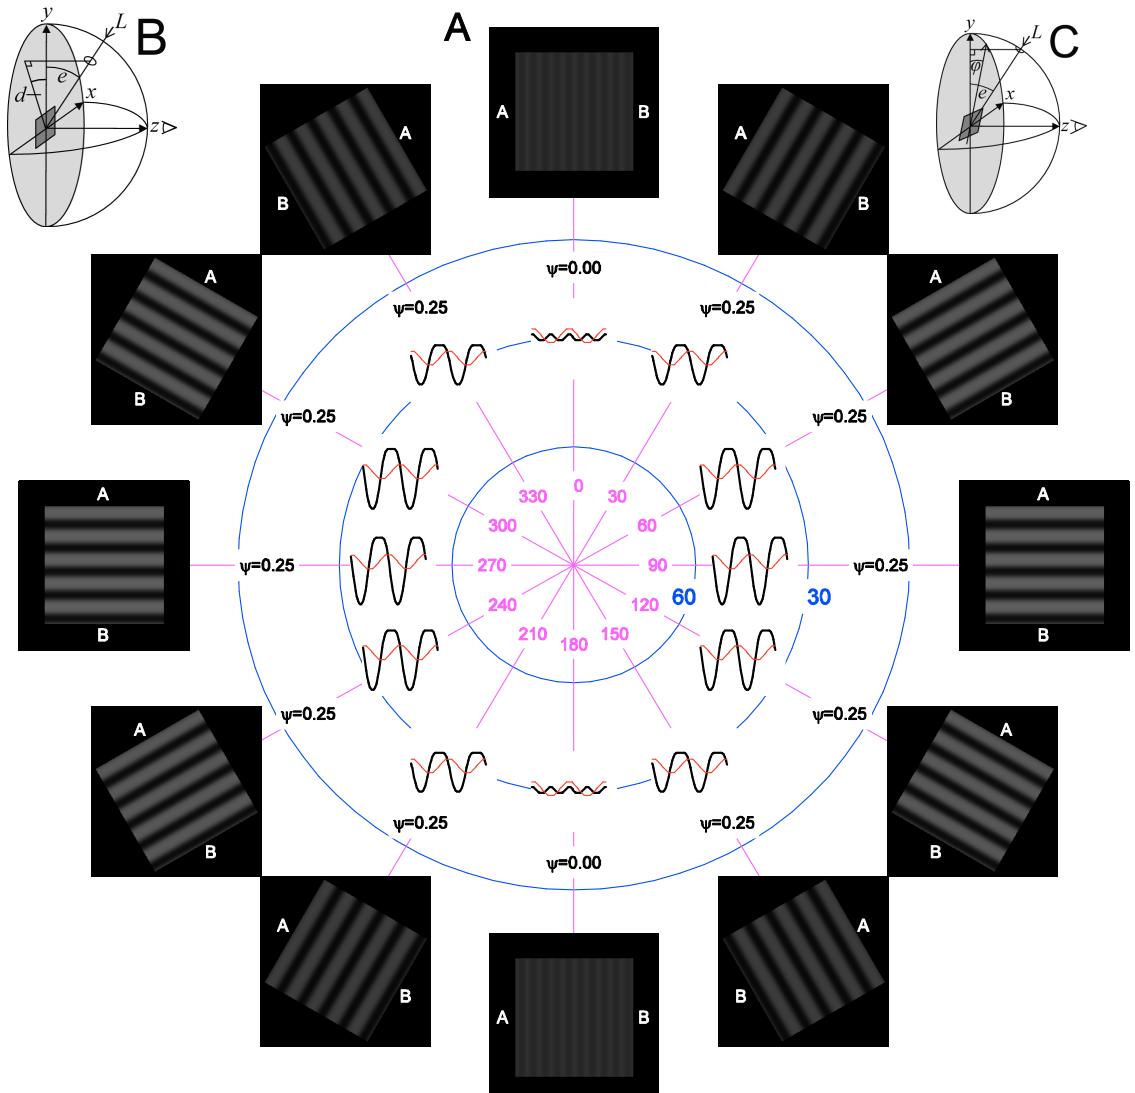
1147

1148

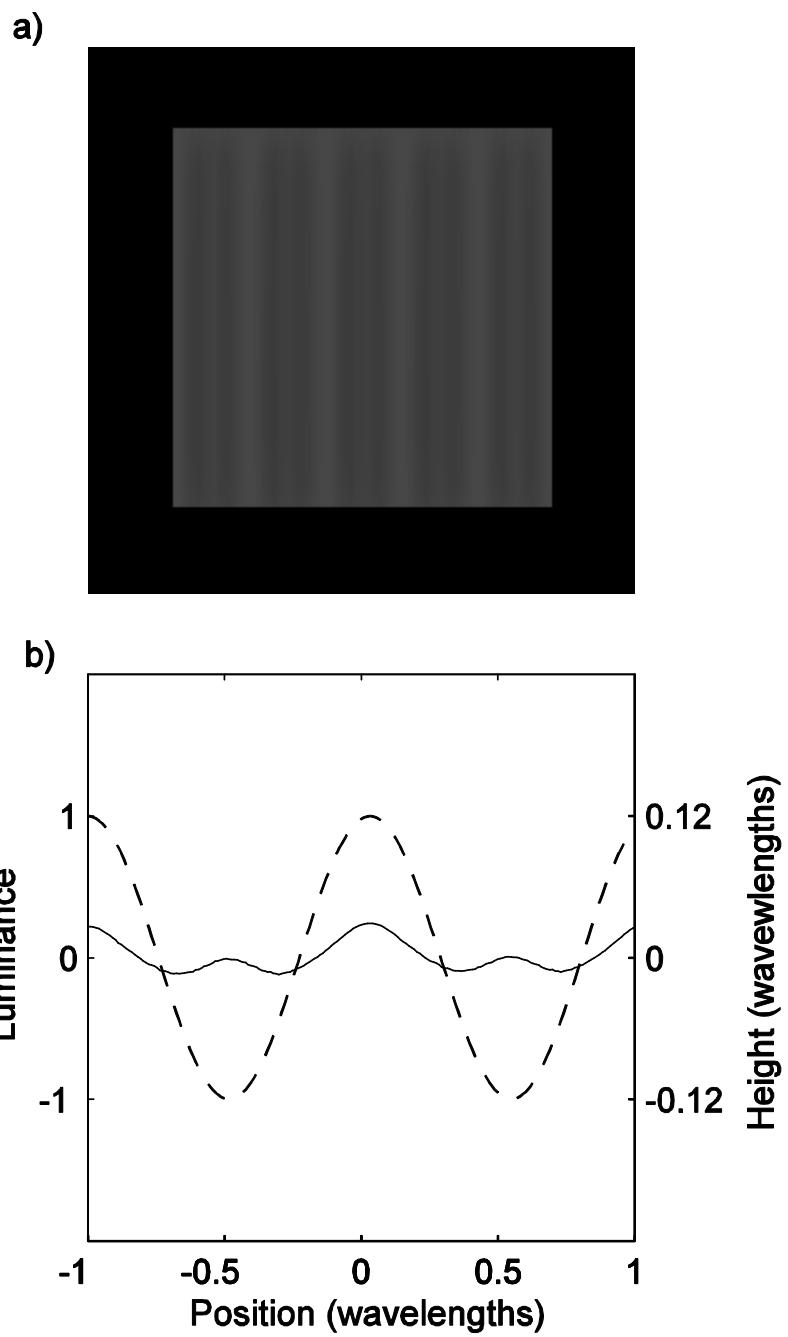
1149

1150

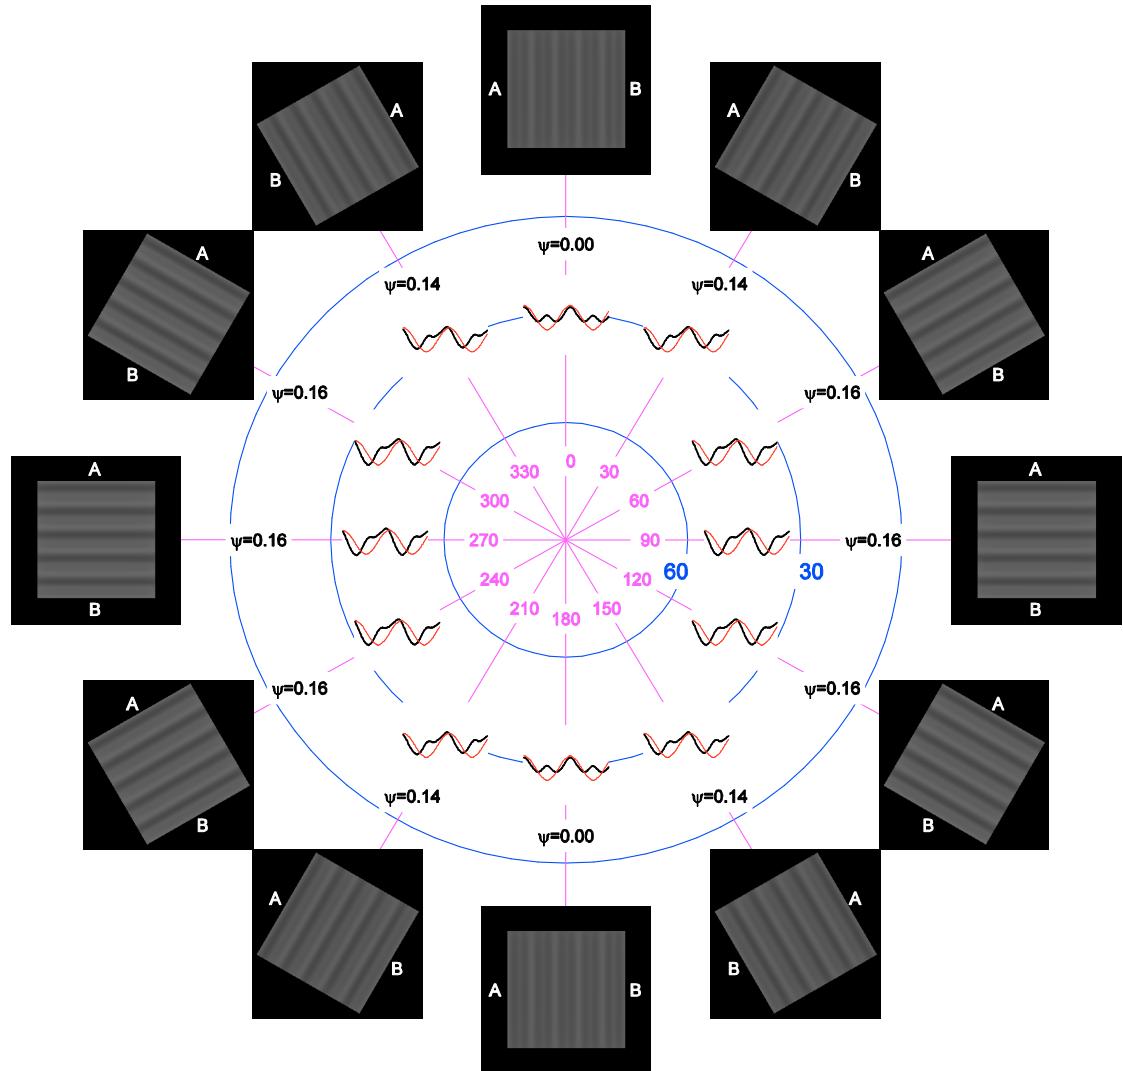
1151



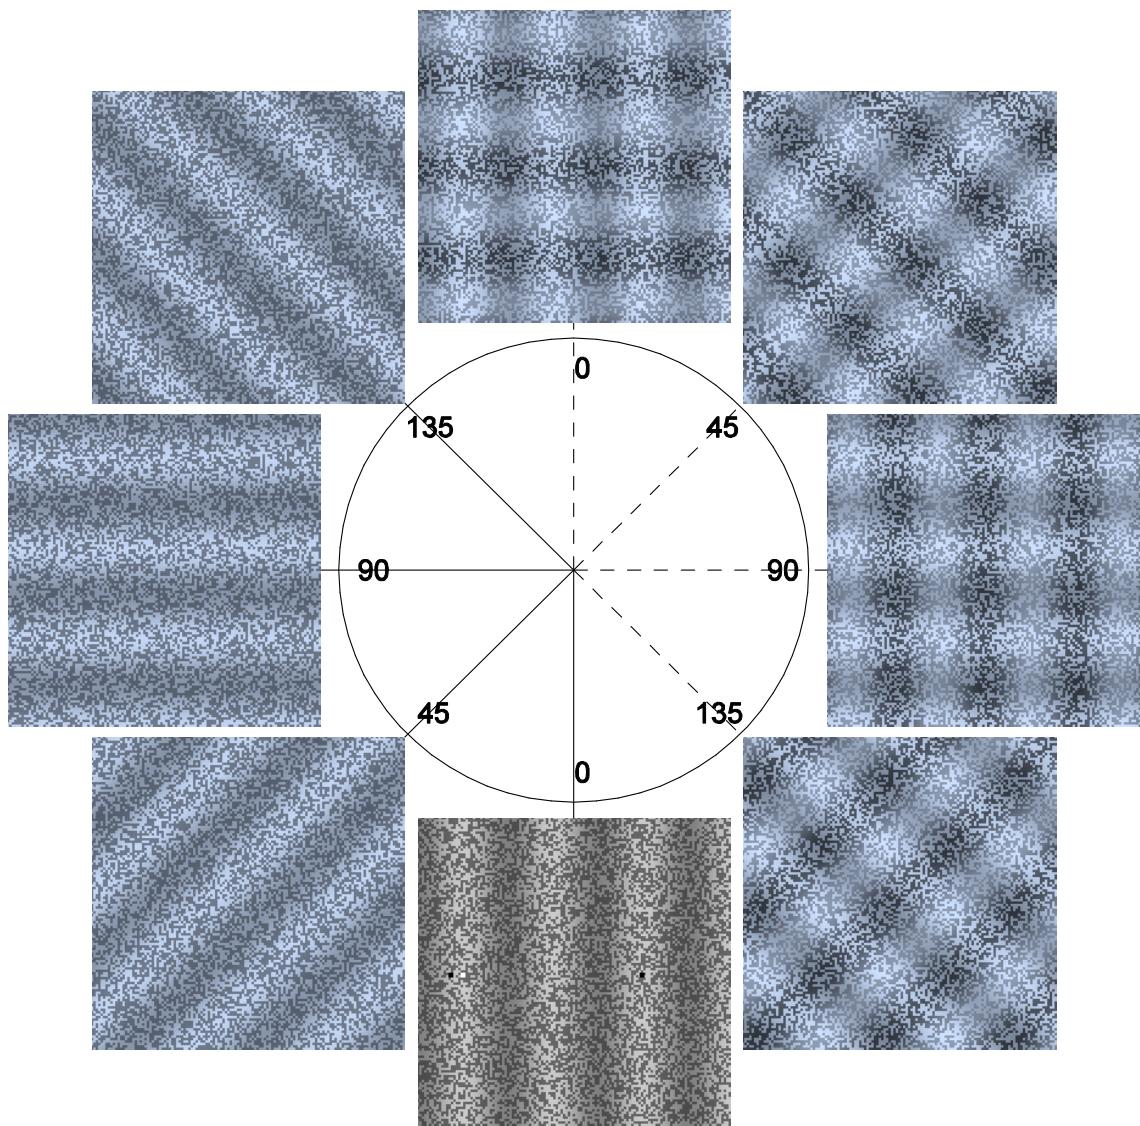
1152
1153 Fig 1



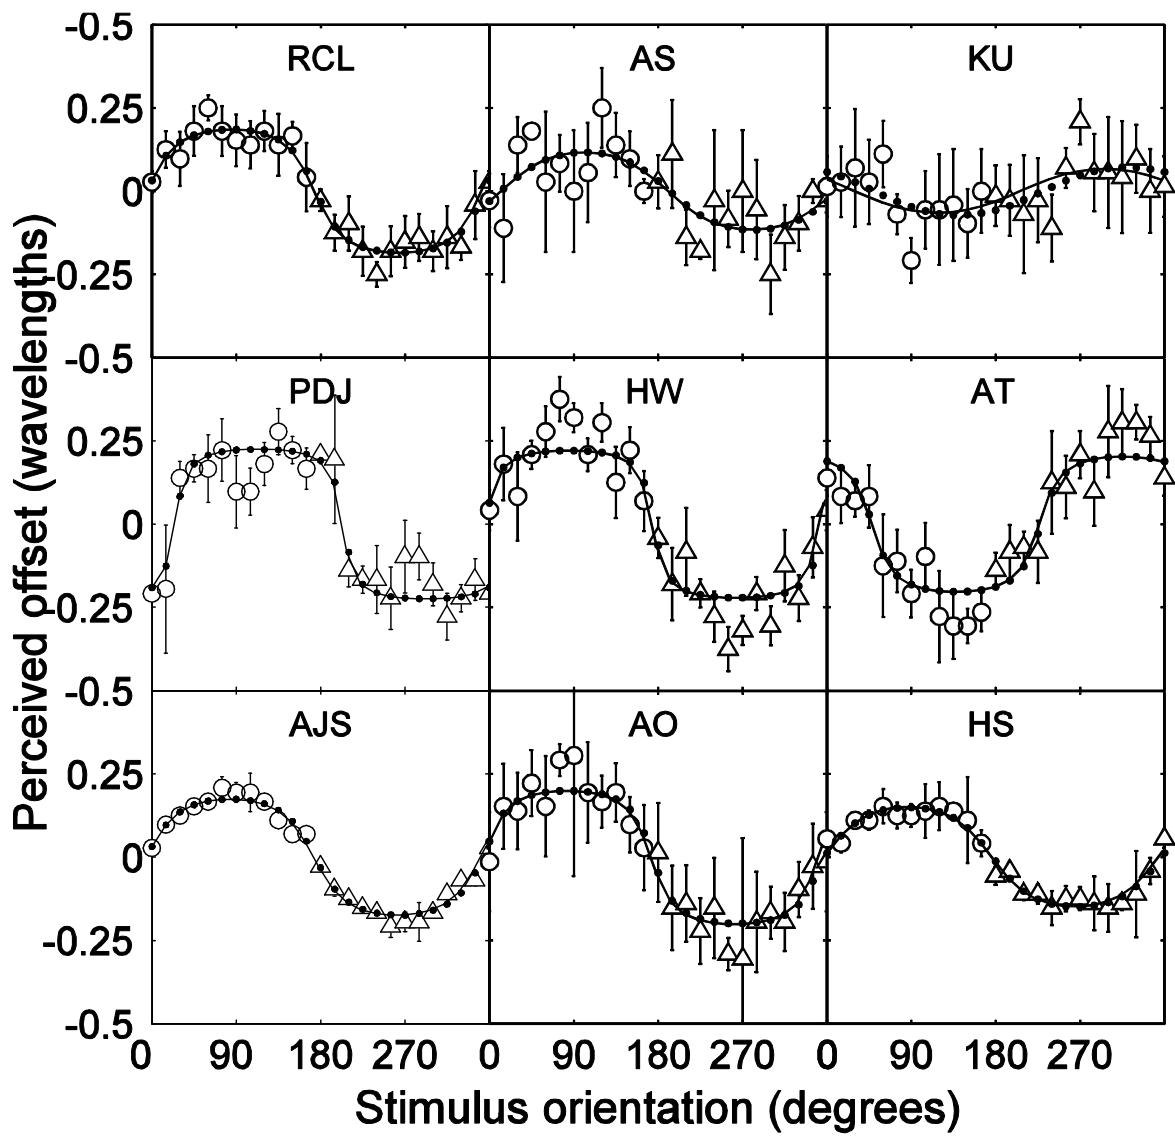
1154
1155 Fig 2



1156
1157 Fig 3

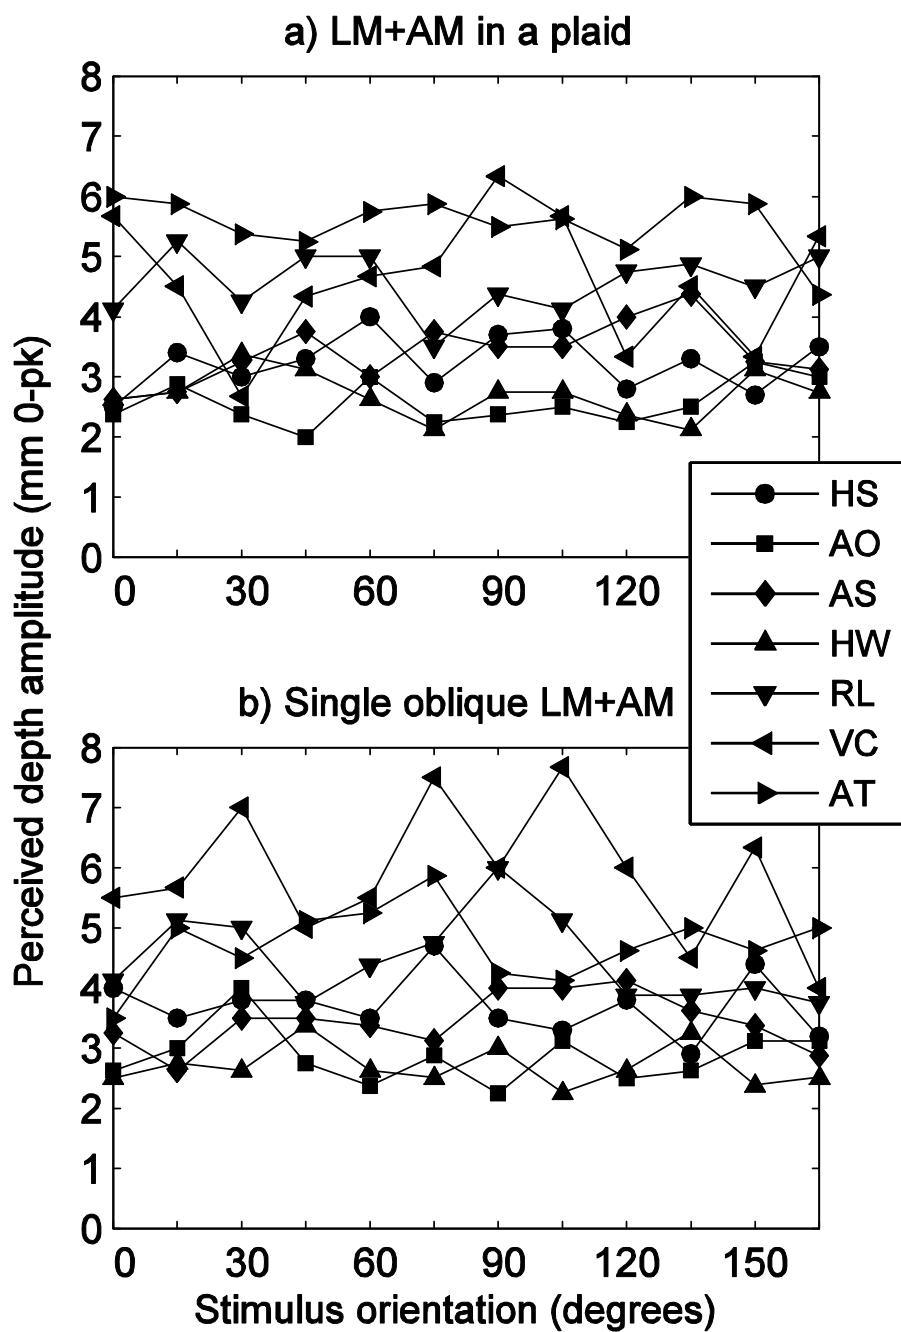


1158
1159 Fig 4



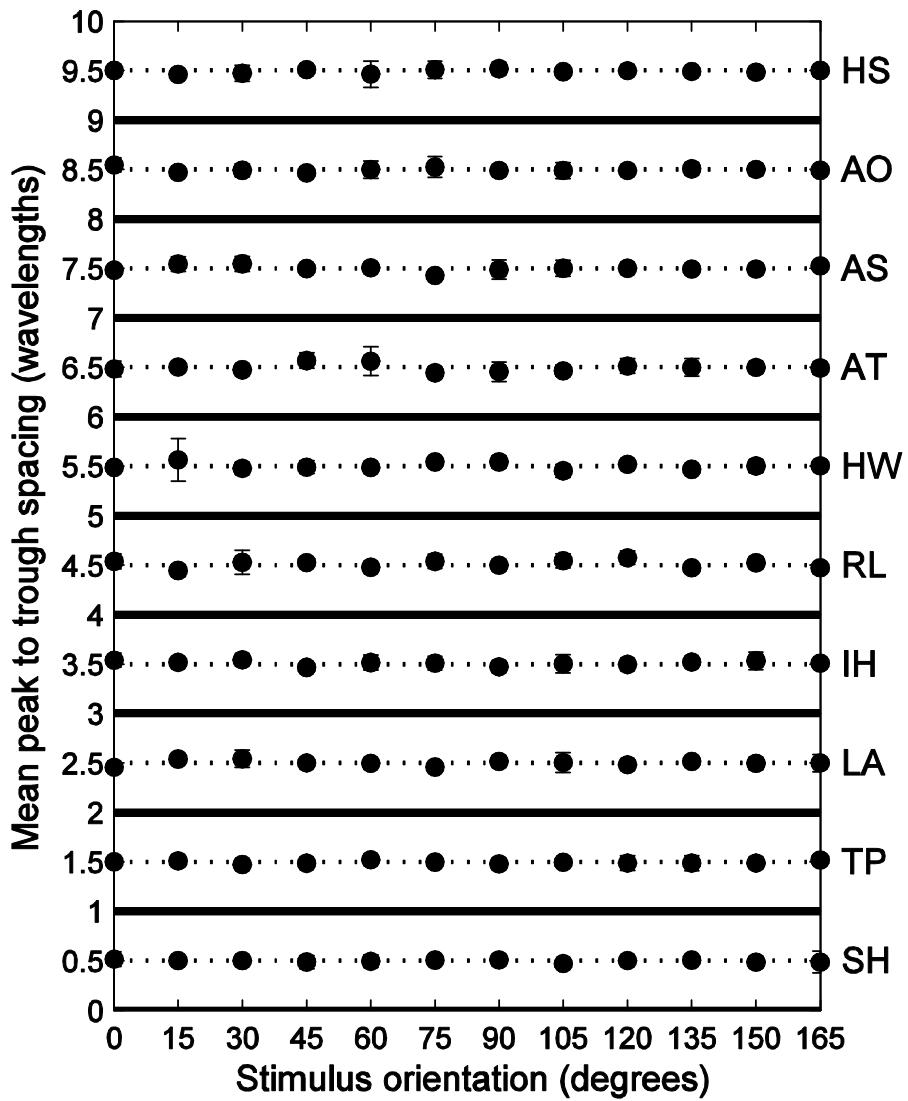
1160
1161

Fig 5

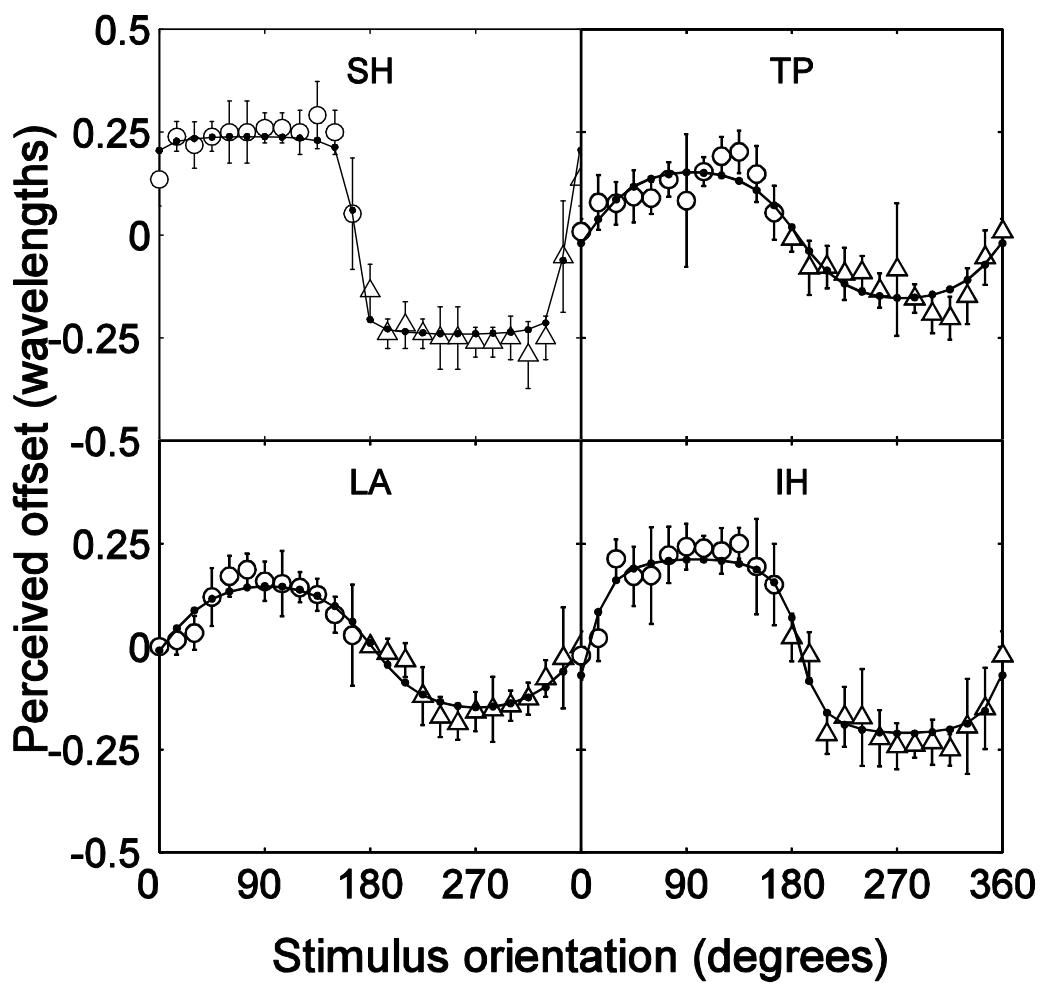


1162
1163

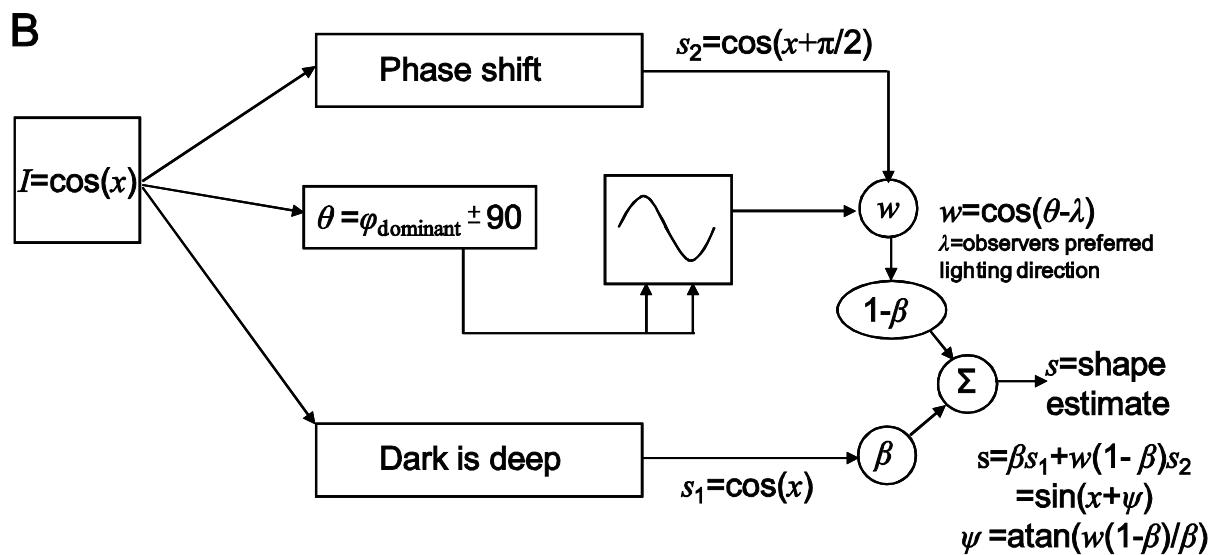
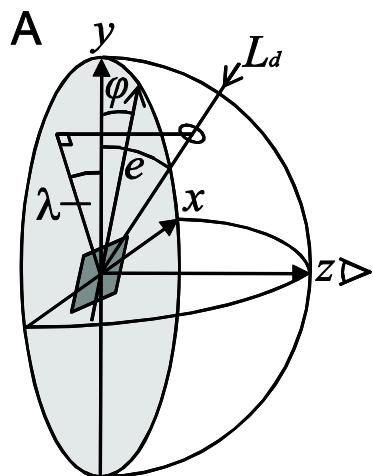
Fig 6



1164
1165 Fig 7



1166
1167 Fig 8



1168
1169 Fig9

Sun and sky: Does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading?

Andrew J Schofield, Paul B Rock, & Mark A Georgeson.

Supplementary data.

S1. The relationship between illumination direction and shading profile for sinusoidal surfaces.

One reason for using sinusoidal shading patterns (gratings) is that we know that humans interpret these as sinusoidal surfaces. We have previously shown, in line with Pentland (1988), that sinusoidal luminance signals (single gratings or plaids) produce a convincing percept of a sinusoidally undulating surface, of the same spatial frequency as the luminance profile (Schofield, Hesse, Rock, & Georgeson, 2006). Employing a generative model we now illustrate the point source lighting conditions under which sinusoidal undulations give rise to approximately sinusoidal luminance profiles. We presume that humans adopt one of these lighting interpretations when viewing sinusoidal gratings; if they adopt a point source interpretation at all.

The relationship between sinusoidal surfaces and their luminance profiles under point illumination is illustrated in Figure S1. This figure shows surface profiles (thin red lines) for a vertically oriented sinusoidal surface based on the fronto-parallel plane and undulating in the direction of the line of sight. This surface was rendered under a range of point light sources with the resulting luminance profiles shown by the thick black lines. The polar position of each trace represents the direction and elevation of the light source relative to the surface plane (see legend). The luminance profiles shown are seldom sinusoidal, but when the elevation of the light source is low and its direction oblique to that of the undulations luminance is dominated by a component at the frequency of the undulations (i.e. linear shading; Pentland, 1988). Those parts of the surface most oriented towards the light source have the highest luminance, but do not, in general, correspond to surface peaks (that is, the points closest to the observer for fronto-parallel presentation or more generally points of maximum – convex – surface curvature). In such cases, luminance peaks are offset from surface peaks by $\frac{1}{4}$ wavelength. When the elevation of the light source is increased or its direction relative to the undulations is more acute, contrast is reduced and a component at twice the undulation frequency is introduced into the luminance profile (i.e. quadratic shading, Pentland 1988)¹. In such cases the average position of the two luminance peaks is offset by $\frac{1}{4}$ wavelength from the surface peaks. When the light source is positioned on the line of sight (90° elevation) or directed along the surface ridges (directions 0 and 180°) the frequency-doubled component dominates. Here surface and luminance peaks align but there are additional luminance peaks at the surface minima.

When the elevation of the light source is very low the surface partially occludes itself. Occlusions are worst for obliquely oriented light sources. In many cases occlusion does not shift the position of the luminance peak relative to the surface peak but shifts will occur in extreme cases.

¹ We note that sinusoidal undulations are actually degenerate with respect to the direction of the illumination when expressed in a Cartesian framework. The shape of the luminance profile depends only on the angle between the surface normal and the vector joining the light source to the surface. In our polar representation this single angle is determined by two components that we call elevation and direction. However this distinction is unimportant for two reasons. First, we make no strong claims about the direction of the preferred illuminant in this paper. Second, we did not use the rendered stimuli in this demonstration in our experiments. The purpose of this demonstration is merely to show that sinusoidal surfaces can give rise to approximately sinusoidal luminance profiles over a range of point source lighting conditions and that when this is the case the luminance peaks and physical peaks are offset by $\frac{1}{4}$ wavelength.

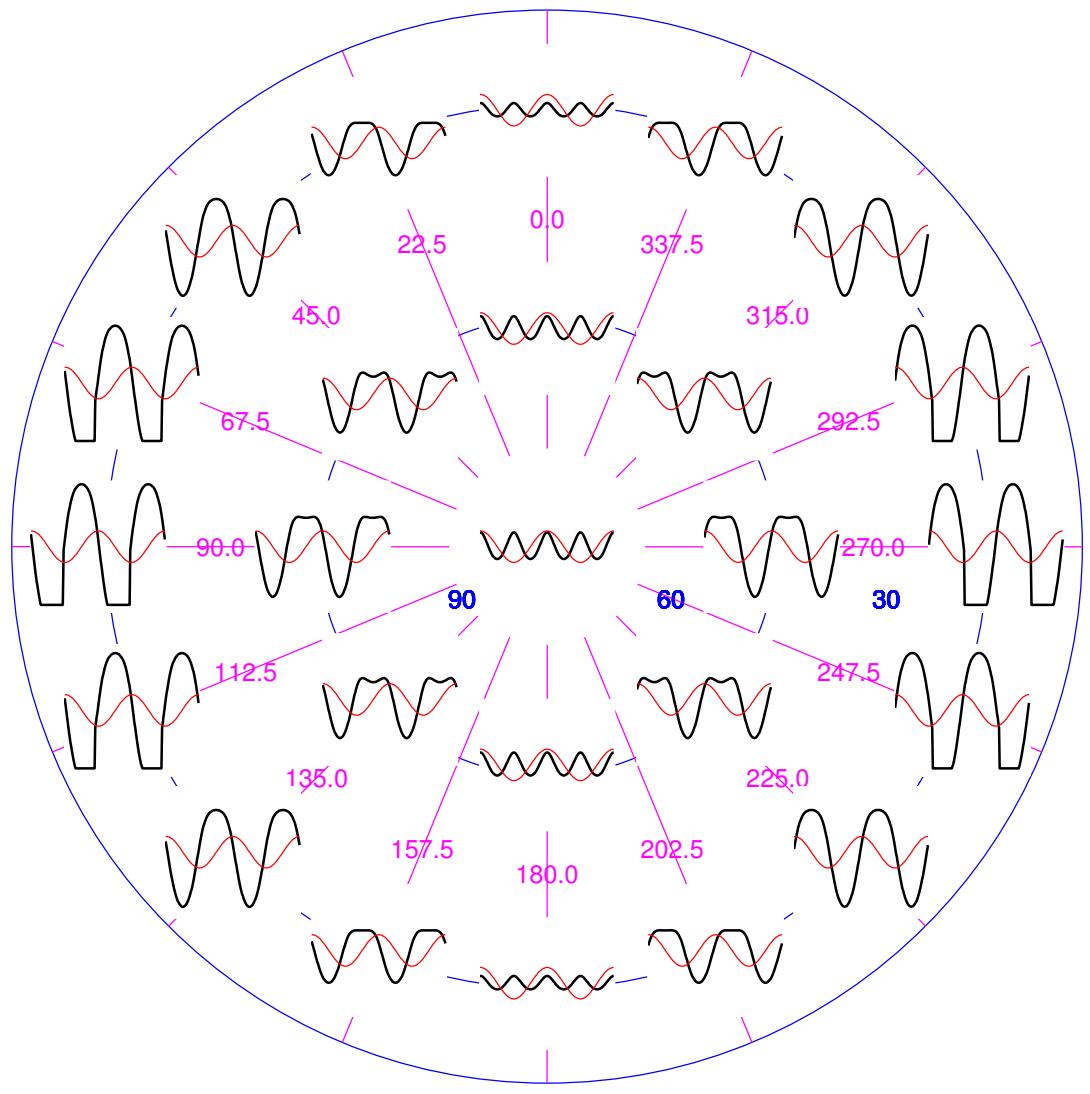


Figure 1. Luminance profiles for a sinusoidal surface under point-source illuminants. Thin red traces show the surface profile; Thick black traces show luminance after rendering the surface under point source illuminants from a range of directions. The surface was sinusoidally corrugated with ridges running top to bottom. Its depth amplitude (mean-to-peak) was 0.12 wavelengths (axes not to scale). The polar location of each pair of traces represents the direction and elevation of the light source relative to the centre of the surface. Polar orientation (magenta radial lines and labels) represents the direction of the light projected into the surface plane while the distance from the centre (blue circles and labels) represents the elevation of the light out of that plane. Note in main Fig 1 radial lines represent surface orientation labelled clockwise; here lighting direction is labelled working anti-clockwise such that relative direction is consistent across the figures. For surfaces based on a horizontal plane direction and elevation should be interpreted as the azimuth and elevation of the ‘sun’ respectively. However, if the surface is interpreted as based on a vertical plane direction and elevation should be interpreted as follows. When elevation = 90° the light is directly in front of the surface regardless of its direction (central traces) lights with elevations 60° and 30° are progressively closer to the surface plane; direction 0 is above the midpoint of the surface, 90 to its left, 180 below and 270 right. For each subplot the x-axis represents position along the surface while the y-axis represents luminance (thick lines) or height (thin lines). Traces were produced in *MatLab* (The Mathworks, Inc) using a Lambertian reflectance model (luminance is proportional to the cosine of the angle between the surface normal and the vector to the light source) with occlusion (luminance is zero when the surface occludes itself from the light source).

The renderings of Figure S1 and main Figure 1 show shading profiles for surfaces with relatively low relief (0.12 of a wavelength, mean to peak depth) illuminated by point light sources with varying direction and elevation. The plots suggest that for such surfaces occlusions are relatively rare for elevations greater than 30° and that while double-crested peaks occur there is a sizable range of elevations and directions for which the quadratic component is relatively small. Thus for low relief surfaces there is an elongated annulus in illumination space for which shading profiles are approximately sinusoidal. In this band – represented best in main Figure 1 – luminance peaks are offset from surface peaks by $\frac{1}{4}$ wavelength. Figure S1 also suggests that this offset persists even when some occlusions occur and that double-crested peaks are centred on this offset. This analysis suggests that $\frac{1}{4}$ wavelength offsets are common for such stimuli. We investigated this notion further by rendering surfaces under every possible illuminant in the space depicted in Figure S1 (sampled on a 181x181 position grid). We then measured a number of properties of the resulting shading images. We estimated the location of the dominant peak in the stimuli by measuring the phase of the frequency component with the same frequency as the surface (the fundamental frequency). This is represented by hue in Figure S2a with purple and green equal to $\frac{1}{4}$ wavelength offsets in the positive and negative directions respectively. We also measured the degree to which the shading profiles were sinusoidal by taking the ratio of the amplitude of the fundamental to all of the higher order components combined. This is represented by saturation in Fig S2a. Finally, we measured the degree of occlusion by simply measuring the proportion of ‘pixels’ on the surface that can ‘see’ the illuminant, 1= no occlusions, 0=surface totally occluded. This is represented by intensity in Fig S2a.

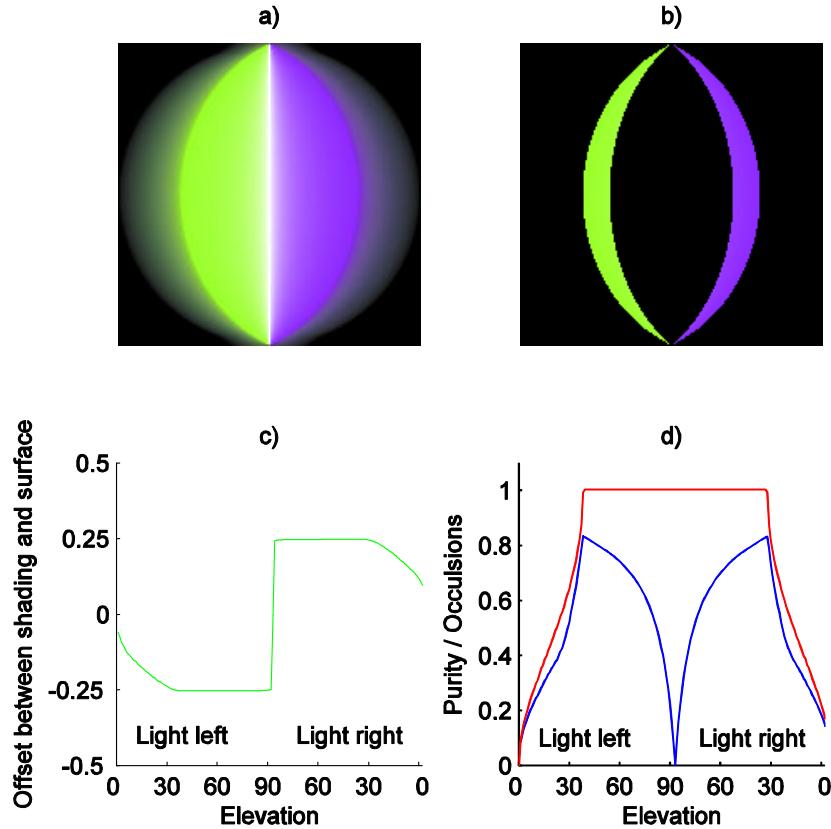


Figure S2. a) Representation of the illumination space. Hue corresponds to the offset between surface peaks and the fundamental component of the luminance waveform – normally the dominant peak. Saturation represents the ‘purity’ of the shading waveform given by the amplitude of the fundamental divided by the total amplitude of all other components. Intensity represents the proportion of surface points that can see the light source. b) as (a) but only those light locations that produce neither occlusion nor double peaks are coloured. c) cross section through (a) showing offsets. d) cross sections through (a) showing purity (blue) and ‘occlusions’ (red) where ‘occlusions’ records the proportion of surface positions that can see the light source.

Fig S2b shows offsets (hue), and purity (saturation), for only those lighting directions that produce no occlusions at all and which contain no evidence of a double peak.

Fig S2c shows offsets for a cross section through the midline of Fig 1a.

Fig S2d shows purity (blue) and occlusions (red) for the same cross section as Fig S2c.

It can be seen that the peak of the fundamental component of the shading profile is almost always $\frac{1}{4}$ wavelength offset from the surface peaks. This breaks down only when either the quadratic component in the shading profile begins to dominate or when the degree of occlusion is large. We conclude that for shading stimuli that are dominated by linear component shading - surface offsets will generally be $\frac{1}{4}$ wavelength. Importantly our sinusoidal stimuli can provide a reasonably good match to the shading profiles produced from surfaces whose half-height amplitude is 0.12 of their wavelength – that is, the haptic surfaces of Experiment 1.

The depth of the undulations represented in Figures S1 and S2 and main Figure 1 is not accidental but (as a proportion of wavelength) is the same that used for the haptic stimuli of Experiment 1. These figures clearly show that sinusoidal luminance profiles are a physically plausible representation of sinusoidal surfaces that are configured as in Experiment 1 and rendered under point light sources. Further, this illustration strongly suggests that in perceiving sinusoidal gratings as sinusoidally undulating surfaces humans should, if they adopt a point source illumination model, also perceive the surface peaks to be offset by $\frac{1}{4}$ wavelength from the luminance peaks. We assert this because, in the generative case, such offsets are ubiquitous whenever sinusoidal surfaces produce even approximately sinusoidal luminance profiles. We discuss models for how people might perceive luminance peaks to be offset from surface peaks by $\frac{1}{4}$ wavelength in Section 1.2 of the main paper.

We conclude that point source illumination of a low-relief sinusoidal surface produces approximately sinusoidal luminance profiles over a range lighting directions. Except for frontal lighting, lighting directed in the direction of surface ridges, and very oblique lighting, luminance profiles have the same fundamental frequency as the surface itself but their peaks are offset from the surface peaks by $\frac{1}{4}$ wavelength.

S2. Control Experiment 1: Do observers see sinusoidal luminance patterns as sinusoidal surfaces?

S2.1 Rationale

Our experiments rely on the observation that people perceive sinusoidal luminance gratings as sinusoidal surfaces. Our previous results (Schofield, et al., 2006) suggest that this is indeed the case, as do the results presented by Pentland (1988). However, neither of these studies used the same stimulus presentation as we do here. In particular they used vertically presented stimuli rather than stimuli inclined at 45°. Further, neither study formally tested the hypothesis that the shape of the perceived surface is sinusoidal versus some other profile. This control experiment addresses these issues using a gauge figure task to assess perceived gradient at points along each sinusoidal modulation and thence to estimate perceived surface shape.

S2.2 Method

Stimuli (see Fig S3) were similar to those of experiments 1 and 3 and consisted of either LM+AM / LM-AM plaids or single sinusoidal LM+AM gratings applied to samples of isotropic binary noise. We tested six different orientations (0, 30, 60, 90, 120, 150°) at eight different phases (0, 45, 90, 135, 180, 225, 270 and 315 degrees of phase). Each orientation / phase combination was tested 10 times. Stimuli were presented in the ReachIn haptic workstation described in Experiment 1 and thus on a plane tilted backwards at 45°. However, observers looked down at the stimuli at a similar angle so the stimulus plane was fronto-parallel. A gauge figure, comprising a 2D rendition of an unfilled ellipse with an orthogonal stick (see Fig S2) was placed at the centre of each image and observers were asked to adjust the apparent orientation of this figure until it appeared to lie on the surface. By changing the phase of the sinusoid we effectively moved it under the gauge figure allowing us to measure gradients at different positions on the surface. The gauge figure could only be adjusted in the direction orthogonal to the sinusoidal modulation under test. That is, for a vertical sinusoid the gauge figure could only be made to point to the left or right. The initial orientation of the gauge figure was randomised within the range ±85° degrees at the start of each trial. The phase of the untested LM-AM component of the plaid stimuli could be either 90 or 270 degrees such that the gauge figure was always at a zero crossing for this cue.

Six naive observers took part in the study, three each completing the plaid and single oblique conditions. None had participated in our previous experiments. All other experimental details were as described in Experiment 1 of the main study.

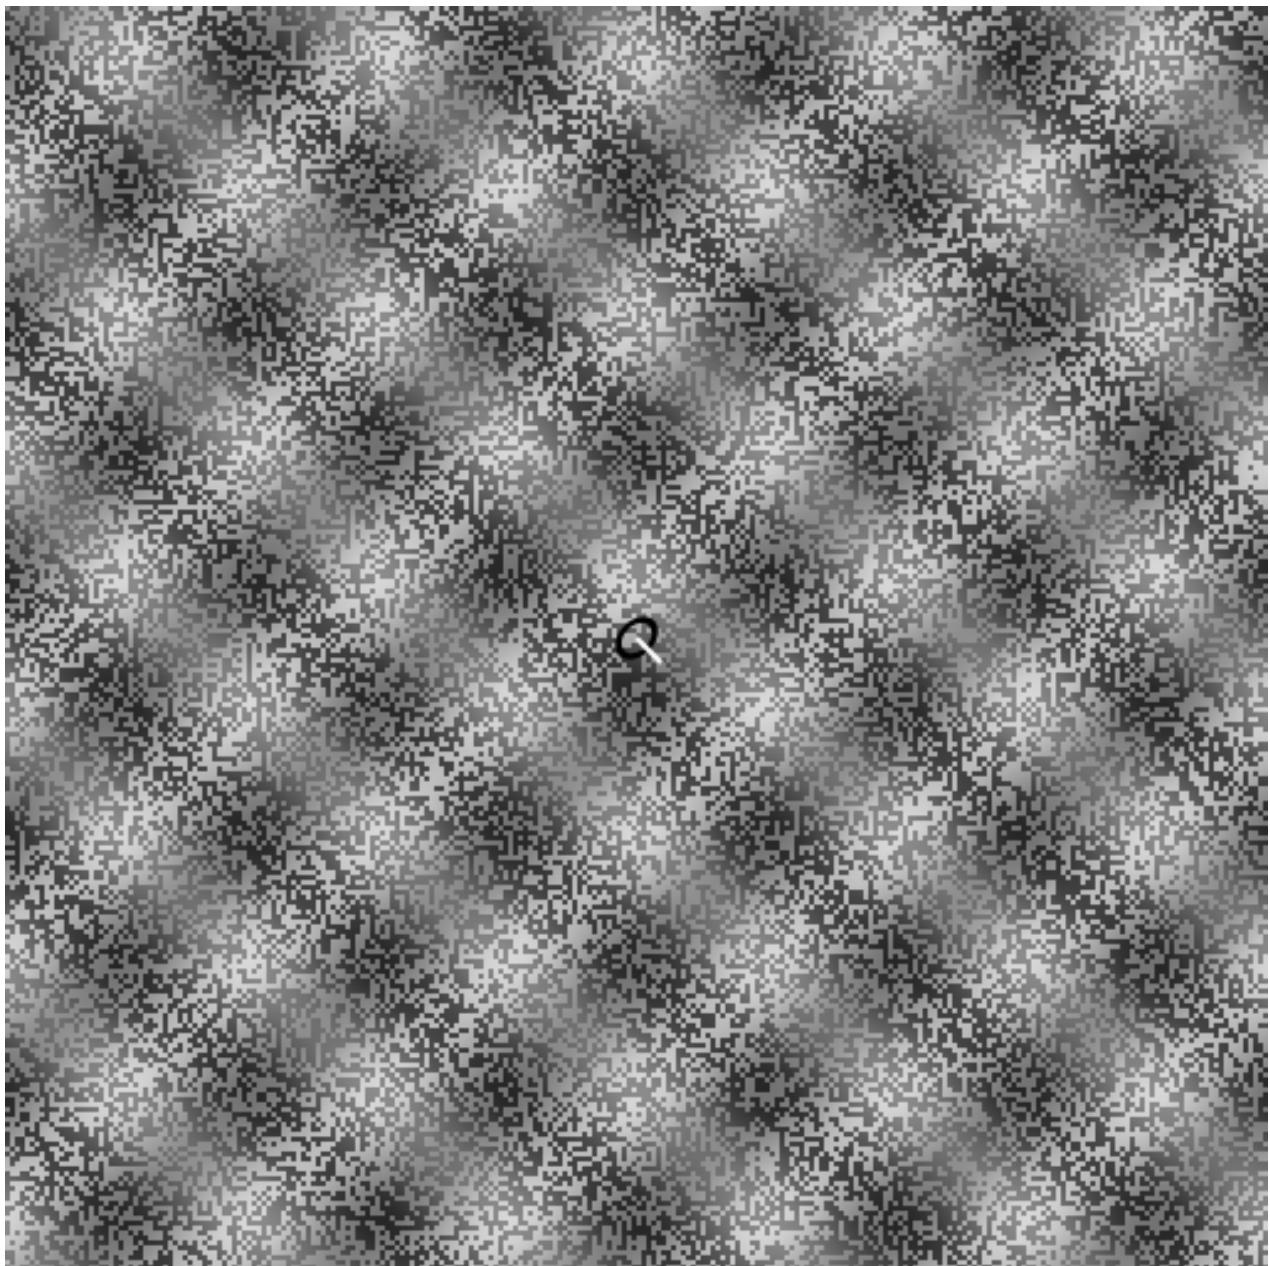


Figure S3a. Example plaid stimulus from control experiment 1 showing gauge figure.
Noise contrast has been exaggerated to aid visualisation.

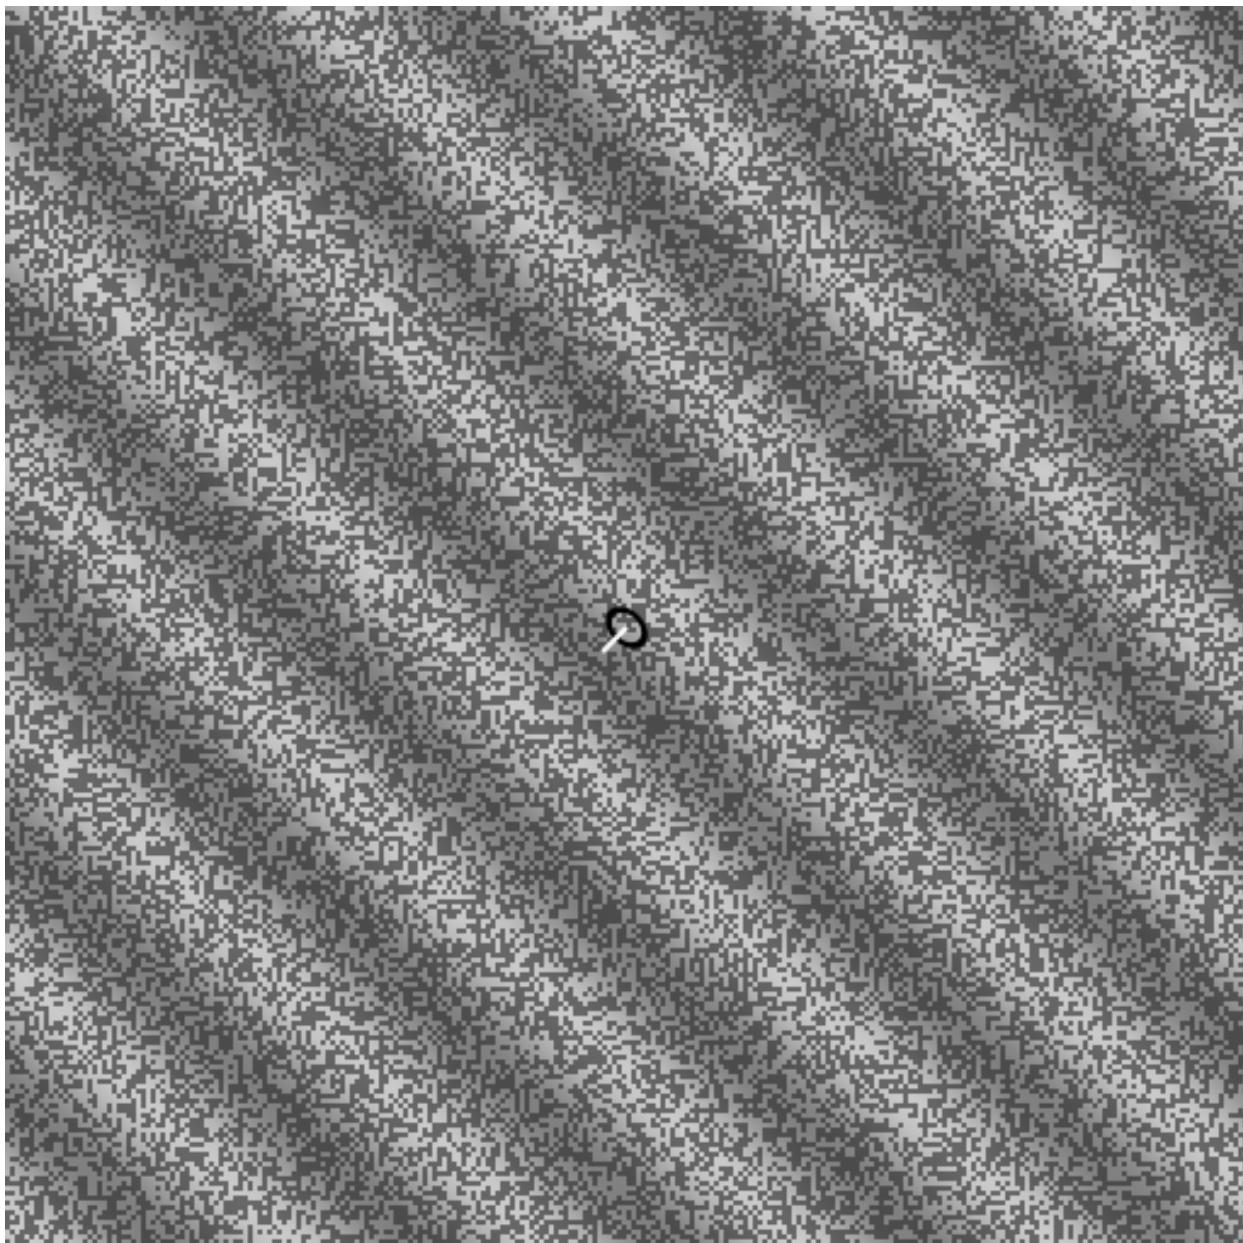


Figure S3b. Example single oblique stimulus from the control experiment showing gauge figure. Noise contrast has been exaggerated to aid visualisation.

S2.3 Analysis

We took the mean of the ten gradient estimates for each orientation / phase combination. We then collated the data for all phases at a given orientation and fit these data with three plausible gradient profiles representing 4 possible surface interpretations. The functions used were: sinusoid, square-wave, and saw-tooth. Once integrated to form perceived surfaces these gradient functions correspond to the following surface profiles: sinusoidal gradient profiles integrate to produce sinusoidal surface profiles with a phase shift of 90 degrees (1/4 wavelength), square-wave gradient profiles integrate to form triangle wave surfaces with a surface peak located at each falling edge in the gradient profile, negative going saw-tooth ramps integrate to produce a series of humps with sharp valleys (humps peak as the saw-tooth ramp crosses $y=zero$), positive saw-tooth ramps produce sharp ridges and shallow valleys (ridge peaks occur at the abrupt fall in the saw-tooth). We varied the phase (position), amplitude and dc level of each gradient profile in order to obtain the best (least-mean-square) fits. Fits were obtained using MatLab's *fminsearch* function and were compared by calculating the correlation between the predicted and actual gradients for each fit; R^2 values are reported. We were thus able to compare 4 candidate surface profiles: sinusoidal, triangular, smooth-humps and sharp-ridges. Although by no means exhaustive, we feel that these four profiles represent good candidates for the perceived surface in response to our sinusoidal shading waveforms given that they are universally, but informally, described as 'undulating corrugations' by our participants.

S2.4 Results

Figure S4 shows a typical dataset and fitted curves for one participant at one orientation. Table S1 shows the R^2 values for all six observers, six orientations and three fitted gradient profiles. R^2 values for the sine wave profiles are generally quite high (typically > 0.8). The sinusoidal model provided the best fit in 34 of 36 cases: at all orientations for four observers and at all but one orientation for the remaining two observers. We conclude that the sinusoidal gradient profile provides a better fit to the data than either the square-wave or saw-tooth profiles. We also conclude that observers see sinusoidal shading patterns as sinusoidal surfaces most of the time. We note that there was very little variation in the amplitude of the best fit gradient profiles with orientation. This result suggests that the surfaces are seen as having about equal depth amplitude at all orientations: confirming the results of Experiment 2.

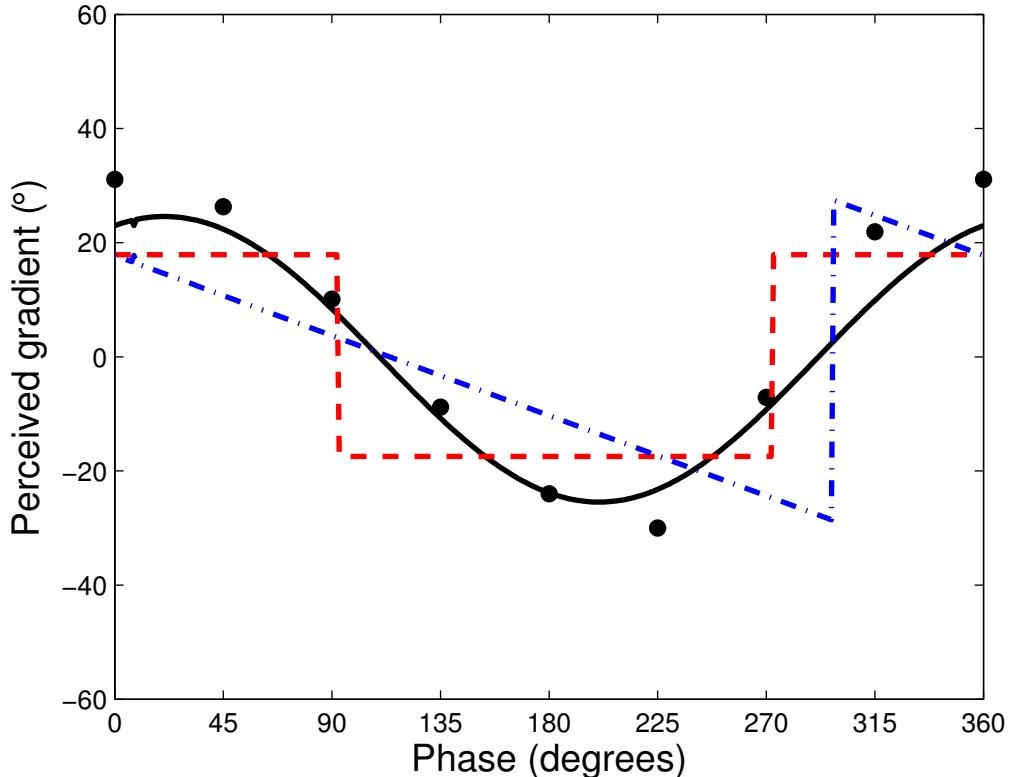


Figure S4. Gradient profile fits for an example dataset from one observer. Symbols, raw data; black solid line sinusoidal gradient profile; red dashed line, square-wave gradient profile; blue dot-dash line, saw-tooth gradient profile.

Orientation	CC			DD			HAW		
	Sine	Square	Saw tooth	Sine	Square	Saw tooth	Sine	Square	Saw tooth
0	.97	.76	.63	.99	.79	.66	.98	.72	.63
30	.7	.77	.76	.85	.69	.56	.82	.78	.73
60	.78	.63	.41	.93	.76	.56	.91	.75	.71
90	.81	.67	.6	.9	.69	.52	.95	.8	.74
120	.84	.65	.58	.91	.67	.5	.89	.84	.81
150	.82	.69	.46	.8	.67	.64	.86	.83	.68

Table S1a. R^2 values for the three gradient profiles for each observer at each orientation in the plaid condition.

Orientation	GM			KC			XJ		
	Sine	Square	Saw tooth	Sine	Square	Saw tooth	Sine	Square	Saw tooth
0	.87	.74	.83	.89	.14	.46	.84	.65	.62
30	.79	.73	.62	.66	.17	.51	.78	.7	.65
60	.88	.83	.67	.76	.75	.51	.81	.59	.44
90	.95	.75	.61	.58	.58	.27	.89	.69	.47
120	.95	.8	.63	.7	.71	.41	.88	.72	.56
150	.85	.77	.21	.71	.62	.35	.78	.64	.6

Table S1b. R^2 values for the three gradient profiles for each observer at each orientation in the single oblique condition.

As noted in Section S3.3 it is possible to estimate the position of the perceived surface peak relative to the luminance peak from the model fit parameters. We computed such estimates based on whichever of the three gradient profiles produced the better fit in each condition. The resulting inter-peak offsets are shown in Fig S5 which has the same format as Fig 3 in the main paper. Inter-peak offsets were fit using the mixed illumination model described in the main paper and the resulting fit parameters are shown in Table S2. These are consistent with those shown in the main paper.

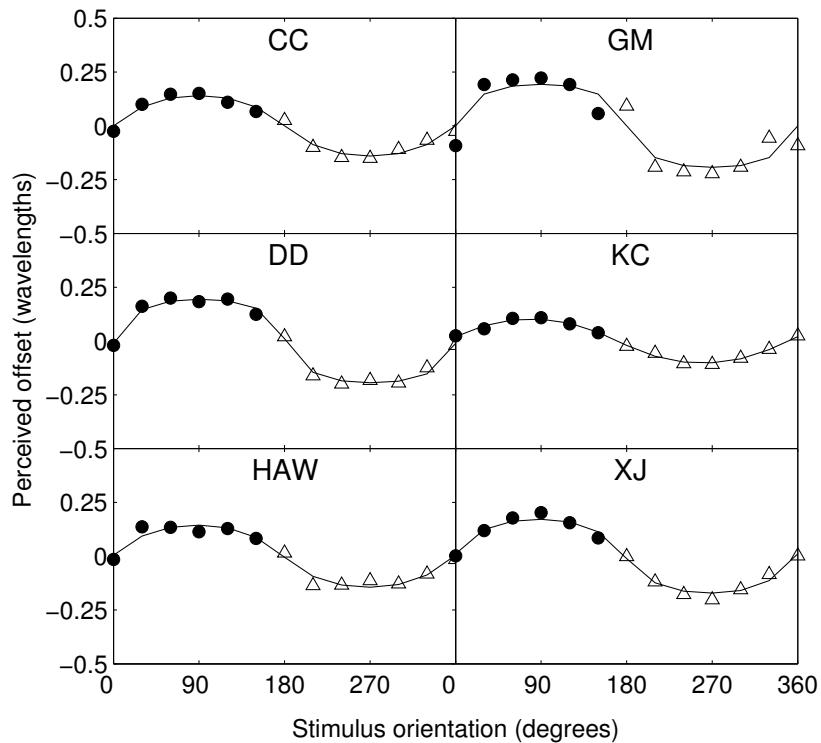


Figure S5. Perceived offsets as a function of stimulus orientation for the three participants in control experiment 1. Circles represent recorded data, triangles interpolated data. Lines represent model fits. For details of the data interpolation and modelling processes, see the main paper.

Stimulus	Person	SSE	β	Preferred light source (λ)
Plaid	CC	.002	0.45	0.7
Plaid	DD	.002	0.27	-1
Plaid	HAW	.003	0.44	2.8
Single	GM	.018	0.247	-5.8
Single	KC	.000	0.57	11.7
Single	XJ	.002	0.35	2.9

Table S2. Model fit parameters. Columns show the sum of squared errors between the modelled and observed inter-peak offsets, the weight (β) applied to the diffuse interpretation, and the observers' preferred lighting direction (the one giving the most positive weight for the point source interpretation). Lighting direction is given in degrees; positive = anti-clockwise shift from vertical.

S2.5 Discussion

The results of this control experiment confirm that observers tend to see sinusoidal shading patterns as sinusoidal surfaces. This finding adds further validation to our approach of asking observers to match haptically presented sine waves to our visual stimuli. The haptic method has the advantage that inter-peak offsets can be assessed more quickly and more directly than in the case with the gauge figure method. Nonetheless the derived inter-peak offsets and resulting model fits for this control are consistent with those obtained in the main study.

S.3 Model equivalence.

The similarity of the two model fits described in the main paper suggest that they two models are mathematically equivalent at least for the case of sinusoidal stimuli. We now demonstrate that this is indeed the case. Basing Model A offsets profiles only on the Fourier component at the frequency of the physical surface is equivalent to simplifying Eqn 5 keeping only those terms at that frequency:

$$L \approx (1 - \gamma) \left(-0.12 \sin(x) \cos\left(\frac{\pi}{2} + \varphi + \lambda\right) \sin\left(\frac{\pi}{2} - e\right) \right) + \gamma(0.065(\cos(x))) \quad \text{Eqn 6}$$

Recalling that e is constant, λ constant for a given fit and φ the direction of the surface corrugations the lumiance profile for a given φ becomes the sum of two orthogonal sine waves with the same frequency and different amplitudes: $\sin(x)$ with amplitude $-0.12(1 - \gamma) \cos(\pi/2 - \varphi + \lambda) \sin(\pi/2 - e)$ and $\cos(x)$ with amplitude 0.065γ . Thus the phase of Eqn 6 is given by:

$$\arctan\left(-\frac{0.12(1 - \gamma) \cos(\pi/2 + \varphi + \lambda) \sin(\pi/2 - e)}{0.065\gamma}\right) \quad \text{Eqn 7}$$

combining all the constants and given that θ in Model B is equal to $\varphi - \pi/2$ and that \cos is symmetric about zero we can rewrite eqn 7 this as:

$$\arctan\left(\frac{0.12 \sin(\pi/2 - e) (1 - \gamma) \cos(\theta - \lambda)}{0.065\gamma}\right)$$

Copyright Andrew John Schofield, University of Birmingham

Citation

Sun P., & **Schofield** A.J. (2012) Two operational modes in the perception of shape-from-shading revealed by the effects of edge information in slant settings. Journal of Vision, 12 (1): 12. doi:10.1167/12.1.12

This author post-print is provided under a Creative Commons: Attribution-NonCommercial-ShareAlike Licence Copyright Andrew J Schofield, University of Birmingham