The human visual system's assumption that light comes from above is weak

Yaniv Morgenstern, Richard F. Murray[1], and Laurence R. Harris

Department of Psychology and Centre for Vision Research, York University, Toronto, ON, Canada M3J 1P3

Edited by Dale Purves, Duke University Medical Center, Durham, NC, and approved June 14, 2011 (received for review January 19, 2011)

Every biological or artificial visual system faces the problem that images are highly ambiguous, in the sense that every image depicts an infinite number of possible 3D arrangements of shapes, surface colors, and light sources. When estimating 3D shape from shading, the human visual system partly resolves this ambiguity by relying on the light-from-above prior, an assumption that light comes from overhead. However, light comes from overhead only on average, and most images contain visual information that contradicts the light-from-above prior, such as shadows indicating oblique lighting. How does the human visual system perceive 3D shape when there are contradictions between what it assumes and what it sees? Here we show that the visual system combines the light-from-above prior with visual lighting cues using an efficient statistical strategy that assigns a weight to the prior and to the cues and finds a maximum-likelihood lighting direction estimate that is a compromise between the two. The prior receives surprisingly little weight and can be overridden by lighting cues that are barely perceptible. Thus, the light-from-above prior plays a much more limited role in shape perception than previously thought, and instead human vision relies heavily on lighting cues to recover 3D shape. These findings also support the notion that the visual system efficiently integrates priors with cues to solve the difficult problem of recovering 3D shape from 2D images.

3D shape perception | visual psychophysics | Bayesian modelling

Most people see Fig. 1 as a bas-relief footprint illuminated from the top of the page, even though it depicts a concave footprint illuminated from the bottom. This percept illustrates the light-from-above prior, the human visual system's implicit assumption that light shines from overhead (1–5). In most environments light originates above the horizon, so the light-from-above prior is a reasonable assumption that helps us choose the most probable interpretation of ambiguous images.

How does the human visual system resolve contradictions between the light-from-above prior and lighting cues in the many scenes where light does not shine from directly overhead? The evidence is inconclusive: some researchers have argued that visual lighting cues completely override the prior (6, 7), whereas others have maintained that lighting cues have either no effect at all on the perceived lighting direction that guides shape-from-shading (8–10), or less influence than nonvisual factors such as head orientation and the direction of gravity (11, 12). This question is pivotal, however, for understanding 3D shape perception. If the visual system relies heavily on the light-from-above prior instead of estimating lighting direction from visual cues, then human shape-from-shading mechanisms must not require accurate estimates of lighting direction (13) and hence differ profoundly from classic computer vision approaches to shape-from-shading (14).

To determine how the visual system resolves contradictions between the light-from-above prior and lighting cues, we probed shape-from-shading percepts under a range of lighting conditions. We showed observers ambiguously shaded disks embedded at random orientations in scenes where shading and shadow cues indicated the true direction of lighting (Fig. 2). The ambiguous disks could be interpreted as bumps illuminated from one direction or as dents illuminated from the opposite direction. Six observers judged whether a target disk in each scene looked like a bump or a dent. The lighting cues were sometimes strong (Fig. 2*A*) and sometimes weak (Fig. 2*B*). In separate blocks, light came from one of six evenly spaced directions (12 o'clock, 2 o'clock, 4 o'clock, and so on). In each condition (two lighting cue strengths × six lighting directions) we found the orientation at which disks looked most like bumps, and we took this to be the lighting direction that guided shape-from-shading processes in that condition (3, 15). We call this the "effective lighting direction." For example, if disks looked most bump-like when their brighter half was 30° clockwise of vertical, then we took the effective lighting direction to be 30° clockwise of vertical. To find the direction of each observer's light-from-above prior [which previous studies have found is not always exactly overhead (3)], we also measured the effective lighting direction in a block of trials in which the ambiguous disks appeared on flat circular surfaces that provided no lighting direction cues (Fig. 2*C*).

Results and Discussion

When the lighting cue direction θ_{cue} was the same as the light-from-above prior direction θ_{prior}, the effective lighting direction θ_{eff} was naturally the same as well, because the lighting cues simply reinforced the prior. However, we found that when the lighting cue direction shifted away from the prior direction, the effective lighting direction also shifted away. Fig. 3 plots the shift of the effective lighting direction away from the prior, $\theta_{eff} - \theta_{prior}$, as a function of the shift of the lighting cue direction away from the prior, $\theta_{cue} - \theta_{prior}$. Under strong lighting cues the effective lighting direction closely tracked the lighting cue direction (Fig. 3*A*), indicating that lighting cues almost completely overrode the prior. Intriguingly, under weak cues the effective lighting direction was neither the prior direction nor the cued direction, but instead was approximately halfway between the two (Fig. 3*B*). (Fig. S1 shows detailed data from a typical observer, and Fig. S2 shows the directions of individual observers' light-from-above priors.)

Evidently observers dealt with contradictions between the light-from-above prior and lighting cues by using an effective lighting direction that was a compromise between the two, and the compromise depended on how strong the lighting cues were. To examine this strategy more closely, we used a vector sum model of how observers combine information from two or more cues to estimate a direction (11, 16, 17). In this model the prior direction and the lighting cue direction are represented by unit vectors $\mathbf{v_{prior}}$ and $\mathbf{v_{cue}}$, respectively, and the effective lighting direction is a weighted sum of the two vectors, $\mathbf{v_{eff}} = w_{prior}\mathbf{v_{prior}} +$

Author contributions: Y.M., R.F.M., and L.R.H. designed research; Y.M. and R.F.M. performed research; Y.M. and R.F.M. analyzed data; and Y.M., R.F.M., and L.R.H. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. E-mail: rfm@yorku.ca.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1100794108/-/DCSupplemental.

Fig. 1. This photograph is usually seen as a raised footprint illuminated from above, even though it is actually an indented footprint illuminated from below. (Photograph courtesy of Manuel Cazzaniga.)

$w_{cue}\mathbf{v_{cue}}$. The weight ratio w_{prior}/w_{cue} determines whether the prior or the cues have a greater influence on the effective lighting direction. This model is well established in the literature on the subjective vertical (11, 17). We have recently shown that it is largely equivalent to a Bayesian cue combination model that assigns reliability weights to noisy directional cues [here the prior, which we treat as just another cue (18, 19), and the lighting cues] and combines the cues by making a maximum-likelihood direction estimate (16). Thus, observers who obey the vector sum model are following an efficient statistical cue combination strategy. In *SI Materials and Methods* we describe the vector sum model and our fitting methods in detail.

Fig. 3 shows the fit of the vector sum model. Under strong lighting cues the prior-to-cue weight ratio was $w_{prior}/w_{cue} = 0.13 \pm 0.04$ (maximum-likelihood fit and bootstrapped 95% confidence interval), confirming that strong cues had a much greater effect than the prior on the effective lighting direction ($w_{cue} > w_{prior}$), but also showing that the prior had a measurable residual influence ($w_{prior} > 0$). Under weak cues the weight ratio was $w_{prior}/w_{cue} = 1.10 \pm 0.14$, showing that even our weak cues had approximately as much influence as the prior ($w_{cue} \approx w_{prior}$). Because the prior and the weak cues had approximately equal influence, the weak-cue stimulus (Fig. 2*B*) can be seen as a visual representation of how much directional information observers received from the prior: clearly very little. The large spread in data points at lighting directions near ±180° in the weak condition is also accounted for by the vector sum model, which predicts highly variable direction estimates when the prior and the lighting cues have approximately equal weights but indicate opposite directions (16). In particular, although the fitted curve predicts that the effective lighting direction matches the prior direction when the prior and lighting cue directions are 180° apart, for some observers the effective lighting direction was actually much closer to the lighting cue direction. Fig. S3 and Table S1 report fits to individual observers' data, and we discuss individual differences further in *SI Discussion*.

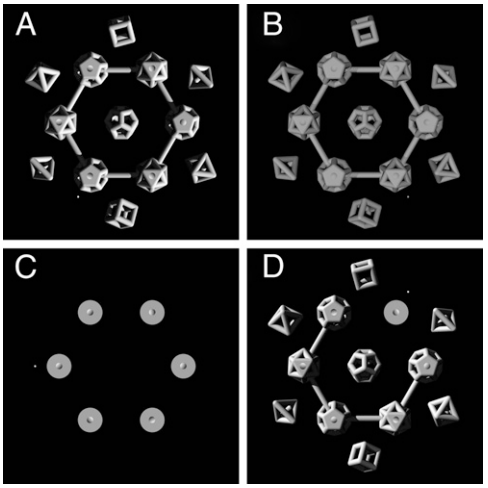

Fig. 2. Typical stimuli. Stimuli in (*A*) strong cue, (*B*) weak cue, (*C*) no cue, and (*D*) no local cue conditions. Observers judged whether the shaded disk next to the small white dot looked like a bump or a dent. Here the lighting directions are (*A*) 4 o'clock, (*B*) 8 o'clock, and (*D*) 10 o'clock.

As an independent measure of lighting cue strength, we measured how precisely observers could rotate an on-screen arrow to indicate the lighting direction in the strong and weak cue conditions. We found that observers' angular errors had circular SD 33° under strong lighting cues and 66° under weak cues (20). For comparison, the circular SD of responses randomly distributed within ±90° of the correct direction is 54°. The high error for judging lighting direction using the weak cues shows that these cues were almost unusable, making it all the more remarkable that they were even partially able to override the light-from-above prior. [Explicit estimates of lighting direction depend on the prior and lighting cues in a manner consistent with our model of the effective lighting direction that guides shape-from-shading (21), but it is nevertheless possible that our observers' explicit estimates of lighting direction differed from the implicit estimates that guided their responses in the main experiment (22).]

We found that observers assigned approximately as much weight to the weak cue stimulus as to the light-from-above prior, and we also found that observers' explicit lighting direction estimates based on the weak cue stimulus had a circular SD of 66°. This suggests that 66° can be taken as a rough estimate of the width of the light-from-above prior. This value is similar to the spread of lighting directions that an observer may encounter over the course of the day: the circular SD of directions uniformly distributed over the upper semicircle is 54°, as is the circular SD of lighting directions uniformly distributed over the upper hemisphere and projected into the frontoparallel plane.

Do these findings mean that the light-from-above prior plays little role in shape-from-shading in real-world scenes that are rich with lighting cues? One caveat is that the lighting cues in our experiment were directly adjacent to the ambiguous disks. In real scenes lighting can vary from place to place, so perhaps the visual system only allows lighting cues to affect the perceived shape of immediately adjacent objects (23) and relies on the light-from-above prior for interpreting large regions that have no local lighting cues. To test this possibility, we eliminated local lighting cues in the strong cue stimulus by showing the target disk on a flat surface and removing the two rods adjacent to the target disk (Fig. 2*D*). Eliminating local lighting direction cues had little effect: the effective lighting direction at the target disk still closely tracked the direction of the lighting cues, even though the nearest cues were on separate objects at least 3.3° away ($w_{prior}/w_{cue} = 0.00 \pm 0.03$; Fig. 3*C*).

Conclusion

The finding that the role of lighting direction cues in shape-from-shading depends heavily on lighting cue strength may explain the inconsistent conclusions of previous studies. Studies that used complex illuminated objects with strong lighting direction cues concluded that lighting cues guide shape-from-shading (6, 7), whereas studies that used weaker lighting cues concluded that they have little effect (8–12). We discuss the lighting cues in previous studies in further detail in *SI Discussion*.

Given the attention paid to the light-from-above prior in previous literature (1–12, 15, 24, 25), one would think that it played a crucial role in shape perception. In fact, the light-from-above prior has a surprisingly weak influence and is easily overridden. Using a weak prior is a rational strategy for the visual system to follow in a world where knowing the current lighting direction is important and where, on average, light comes from overhead, but

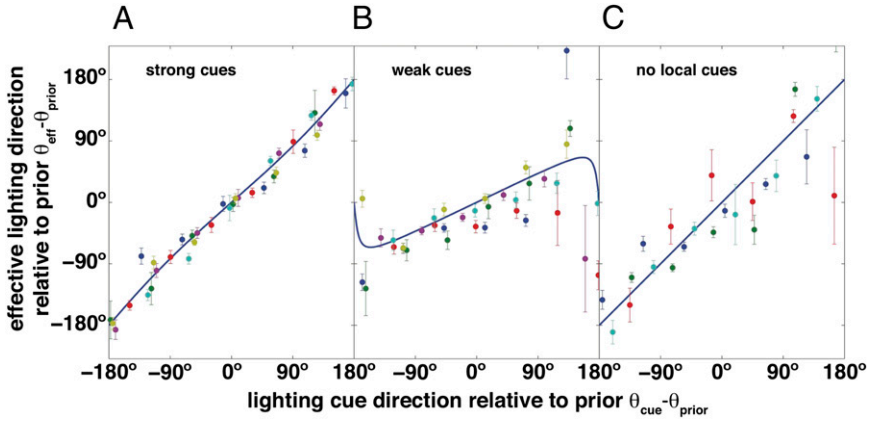

Fig. 3. Effective lighting direction as a function of lighting cue direction. Panels correspond to (*A*) strong cue (six observers), (*B*) weak cue (six observers), and (*C*) no local cue conditions (four observers). Angles are measured relative to the direction of each observer's individually determined light-from-above prior. Different colors correspond to different observers. Error bars indicate SE. Solid lines are maximum-likelihood fits of the vector sum model.

where there are large variations in lighting direction that are reliably cued by shading and shadows.

Materials and Methods

Stimuli and Apparatus. The stimuli were computer-generated images of matte objects rendered in RADIANCE (26). The simulated lighting consisted of two distant point sources. One source was in one of six evenly spaced directions at 12 o'clock, 2 o'clock, 4 o'clock, and so on, 30° toward the viewer from the frontoparallel plane (like clock hands bent 30° forward from the clock face). The second source was in the direction of the viewer. In the strong cue condition the first source was brighter, and in the weak cue condition the second was brighter. In both cases the brighter source contributed 85% of the illuminance of the frontoparallel planes that the ambiguous disks appeared on, which was always 120 cd/m^2. The no cue stimulus (Fig. 2*C*) showed ambiguous disks on six circles of diameter 2.3° and luminance 120 cd/m^2. The ambiguous disks had diameter 0.63°, and the scene subtended 17.4° horizontally.

Procedure. In each condition (strong cue, weak cue, no local cue), observers participated in seven blocks. Six blocks showed scenes with six lighting directions, and the seventh showed scenes without lighting cues (Fig. 2*C*). On each trial, six ambiguous disks appeared at various orientations, with a white dot next to the target disk (e.g., lower left in Fig. 2*A*). The observer judged whether the target disk looked like a bump or a dent. We showed six disks because observers found the task easier if they could compare the target disk with other disks.

Analysis. We made a maximum-likelihood fit of a periodic function to the probability of a "bump" response as a function of the orientation of the target disk, for each observer in each condition. The effective lighting direction was the orientation where the fitted curve peaked. Fig. 3 shows fits of the following equation, which follows from the vector sum model:

$$\theta_{eff} - \theta_{prior} = \arctan2\left(\sin\left(\theta_{cue} - \theta_{prior}\right), \left(w_{prior}/w_{cue}\right) + \cos\left(\theta_{cue} - \theta_{prior}\right)\right). \quad [1]$$

Here *arctan2* is the four-quadrant inverse tangent. See *SI Materials and Methods* for further details.

ACKNOWLEDGMENTS. We thank Minjung Kim and two anonymous reviewers for comments on the manuscript. This work was funded by grants from the Natural Sciences and Engineering Research Council and the Canada Foundation for Innovation.

1. Metzger W (2006) *Laws of Seeing* (MIT Press, Cambridge, MA).
2. Ramachandran VS (1988) Perception of shape from shading. *Nature* 331:163–166.
3. Sun J, Perona P (1998) Where is the sun? *Nat Neurosci* 1:183–184.
4. Mamassian P, Goutcher R (2001) Prior knowledge on the illumination position. *Cognition* 81:B1–B9.
5. Adams WJ (2007) A common light-prior for visual search, shape, and reflectance judgments. *J Vis* 7(11):11, 1–7.
6. Berbaum K, Bever T, Chung CS (1984) Extending the perception of shape from known to unknown shading. *Perception* 13:479–488.
7. Ramachandran VS (1988) Perceiving shape from shading. *Sci Am* 259:76–83.
8. Erens RGF, Kappers AML, Koenderink JJ (1993) Estimating local shape from shading in the presence of global shading. *Percept Psychophys* 54:334–342.
9. Mingolla E, Todd JT (1986) Perception of solid shape from shading. *Biol Cybern* 53: 137–151.
10. Wagemans J, van Doorn AJ, Koenderink JJ (2010) The shading cue in context. *i-Perception* 1:159–178.
11. Jenkin HL, Jenkin MR, Dyde RT, Harris LR (2004) Shape-from-shading depends on visual, gravitational, and body-orientation cues. *Perception* 33:1453–1461.
12. Yonas A, Kuskowski M, Sternfels S (1979) The role of frames of reference in the development of responsiveness to shading information. *Child Dev* 50:495–500.
13. Koenderink JJ, van Doorn AJ (1980) Photometric invariants related to solid shape. *Opt Acta (Lond)* 27:981–996.
14. Zhang R, Tsai PS, Cryer JE, Shah M (1999) Shape from shading: A survey. *IEEE T Pattern Anal* 21:690–706.
15. Hess EH (1961) Shadows and depth perception. *Sci Am* 204:138–148.
16. Murray RF, Morgenstern Y (2010) Cue combination on the circle and the sphere. *J Vis* 10(11):15, 1–11.
17. Mittelstaedt H (1983) A new solution to the problem of the subjective vertical. *Naturwissenschaften* 70:272–281.
18. Mamassian P, Landy MS (2001) Interaction of visual prior constraints. *Vision Res* 41: 2653–2668.
19. Backus BT (2009) The mixture of Bernoulli experts: A theory to quantify reliance on cues in dichotomous perceptual decisions. *J Vis* 9(1):6, 1–19.
20. Fisher NI (1993) *Statistical Analysis of Circular Data* (Cambridge Univ Press, Cambridge, UK).
21. O'Shea JP, Agrawala M, Banks MS (2010) The influence of shape cues on the perception of lighting direction. *J Vis* 10(12):21, 1–21.
22. Hillis JM, Brainard DH (2007) Distinct mechanisms mediate visual detection and identification. *Curr Biol* 17:1714–1719.
23. Ostrovsky Y, Cavanagh P, Sinha P (2005) Perceiving illumination inconsistencies in scenes. *Perception* 34:1301–1314.
24. Rittenhouse D (1786) Explanation of an optical deception. *T Am Philos Soc* 2:37–42.
25. Brewster D (1826) On the optical illusion of the conversion of cameos into intaglios, and of intaglios into cameos, with an account of other analogous phenomena. *Edinb J Sci* 4:99–108.
26. Ward LG (1994) The RADIANCE lighting simulation and rendering system. *Proc SIGGRAPH '94*, ed, Glassner A (Association for Computing Machinery, New York), pp 459–472.

Supporting Information

Morgenstern et al. 10.1073/pnas.1100794108

SI Materials and Methods

Shape Judgment Experiment. ***Participants.*** The first author and five naïve observers participated in the strong lighting cue condition; the first author and five new naïve observers participated in the weak lighting cue condition; and four new naïve observers participated in the no local cue condition. All observers were students at York University, all except the author were paid $10/h, all reported normal or corrected-to-normal vision, and ages ranged from 20 to 34 y.

Stimuli. The stimuli were computer-generated scenes of Lambertian objects modeled after Platonic solids, rendered in RADIANCE. The simulated lighting consisted of two distant point sources. One source, which we call the frontal source, was located directly in front of the scene, in the direction of the virtual camera. The direction of the other source, which we call the directional source, varied from scene to scene, but it was always 30° toward the observer from the frontoparallel plane. When describing the direction of the directional source, we take 0° to mean 30° toward the observer from directly overhead, and positive angles to be clockwise of that direction. That is, a lighting direction θ corresponds to a lighting direction vector $(\sqrt{3}\sin\theta, \sqrt{3}\cos\theta, 1)$, with the observer on the +z axis.

We varied the strength of lighting direction cues by varying the proportion of light originating from the frontal and directional sources, subject to the constraint that the luminance of frontoparallel surfaces was always 120 cd/m^2. In scenes with strong lighting cues, the directional source contributed 85% of the illuminance of the frontoparallel surfaces. In scenes with weak lighting cues, the frontal source contributed 85% of the illuminance of the frontoparallel surfaces. The no cue stimulus used to measure the light-from-above prior consisted of six circles with diameter 2.3° of visual angle and luminance 120 cd/m^2. The no local cue stimulus was the same as the strong cue stimulus, except that on each trial the object at the randomly chosen location where the target disk appeared was replaced by the same type of circle used in the no cue stimulus, and the two rods attached to that object were removed. The scene subtended 17.4° horizontally, and the center-to-center distance between adjacent ambiguous disks was 5.7°. The luminance of the black background was 0.5 cd/m^2, and the peak luminance was 213 cd/m^2 in the strong and no local cue conditions and 127 cd/m^2 in the weak cue condition.

We showed shaded disks on flat frontoparallel surfaces attached to six of the objects. The disks were images of a convex section of a sphere, rendered under the same lighting as the rest of the scene. The spherical section was sufficiently shallow that it did not cast a shadow. Each disk subtended 0.63°. The disks were rendered in RADIANCE separately from the rest of the stimulus and were rotated about their midpoints and added to the stimulus using MATLAB image processing routines.

Stimuli were shown on an Apple iMac (liquid crystal display, pixel size 0.285 mm, resolution 1,680 × 1,050 pixels, software-linearized) in a dark room. Observers viewed stimuli binocularly from a distance of 0.57 m, and head position was stabilized by a chin rest.

Procedure. Each observer participated in seven 10- to 15-min blocks of ≈300 trials over 1 or 2 d. Six blocks showed scenes illuminated from angles −120° to 180° in 60° steps, and the seventh showed shaded disks in the no cue stimulus (Fig. 2*C*). Each block showed a single lighting direction. The seven blocks were run in random order, except that in the weak lighting cue and no local cue conditions, the no cue block was always run first.

At the beginning of each trial, the scene was shown without shaded disks for 1 s. Then the shaded disks appeared, and simultaneously a small white dot appeared next to a randomly chosen object, cueing the observer to press one of two keys to indicate whether the target disk on that object looked like a bump or a dent. The disks and the white dot remained on the screen until the observer responded, and then they disappeared and the next trial began. We showed six disks because pilot trials suggested that some observers found the task easier if they could compare the target disk with disks at other orientations. The target disk orientation was chosen using interleaved staircases in the strong cue condition and the method of constant stimuli in the weak and no local cue conditions. The five nontarget disks were set to the target disk orientation plus −120°, −60°, 60°, 120°, and 180° and randomly assigned to the five remaining surfaces.

Analysis. ***Effective lighting direction.*** To estimate the effective lighting direction for each observer in each lighting condition, we found the maximum-likelihood fit of the following unimodal, periodic curve to the probability of a "bump" response as a function of the orientation of the target disk:

$$p(\theta;\alpha,\tau,\sigma,a,b) = a + (b-a)\left(\Phi\left(\mathrm{mod}(\theta,360°);\alpha-\frac{\tau}{2},\sigma\right) - \Phi\left(\mathrm{mod}(\theta,360°);\alpha+\frac{\tau}{2},\sigma\right)\right). \quad [S1]$$

Here $\Phi(x;\mu,\sigma)$ is the normal cumulative distribution function, *mod* is the modulus function, θ is the orientation of the target disk, α is the orientation where the curve peaks, τ is the width at half-maximum, σ determines the steepness of falloff from the peak, a is the curve's minimum value, and b is the curve's maximum value. We took the peak orientation α, where disks were most bump-like, to be the effective lighting direction that guided shape from shading processes.

Cue combination model. We have shown elsewhere (1) that the vector sum model and the Bayesian cue combination model both predict that the effective lighting direction relative to the prior direction, $\theta_{eff} - \theta_{prior}$, is closely approximated by the following function of the lighting cue direction relative to the prior direction, $\theta_{cue} - \theta_{prior}$:

$$\theta_{eff} - \theta_{prior} = \mathrm{arctan2}\left(\sin\left(\theta_{cue}-\theta_{prior}\right), \left(w_{prior}/w_{cue}\right) + \cos\left(\theta_{cue}-\theta_{prior}\right)\right). \quad [S2]$$

Here *arctan2(y,x)* is the four-quadrant inverse tangent, equal to *arctan(y/x)* when $x > 0$ and *arctan (y/x)* + π when $x < 0$. w_{prior} and w_{cue} are the weights that the observer assigns to the prior and the lighting cues, respectively.

Eq. S2 has a straightforward interpretation. Suppose $\mathbf{v_{prior}} = (\theta_{prior}, 1)$ is a unit vector in polar coordinates that represents the prior direction, and $\mathbf{v_{cue}} = (\theta_{cue}, 1)$ is a unit vector in polar coordinates that represents the lighting cue direction. According to the vector sum model, the effective lighting direction is the direction of the vector $\mathbf{v_{eff}} = w_{prior}\,\mathbf{v_{prior}} + w_{cue}\,\mathbf{v_{cue}} = (\theta_{eff}, r_{eff})$. The equations for this vector sum in polar coordinates are

$$\theta_{eff} = \theta_{prior} + \mathrm{arctan2}\left(\sin\left(\theta_{cue}-\theta_{prior}\right), \left(w_{prior}/w_{cue}\right) + \cos\left(\theta_{cue}-\theta_{prior}\right)\right) \quad [S3]$$

$$r_{eff} = \sqrt{w_{prior}^2 + w_{cue}^2 + 2w_{prior}w_{cue}\cos\left(\theta_{prior}-\theta_{cue}\right)}. \quad [S4]$$

Eq. S2 is simply a minor rearrangement of Eq. S3.

We fitted Eq. S2 to the data in each panel of Fig. 3 as follows. First, we found the direction of each observer's prior by measuring the observer's effective lighting direction in the no lighting cue block (i.e., we measured $\theta_{prior,i}$ for each observer i). Second, we found the effective lighting direction for each observer under each lighting direction (i.e., we measured $\theta_{eff,i,j}$ for each observer i under each lighting direction $\theta_{cue,j}$). For each effective lighting direction estimate $\theta_{eff,i,j}$ we also bootstrapped a concentration parameter, $\kappa_{eff,i,j}$. (The concentration parameter κ of the von Mises distribution is higher for narrower distributions. For narrow distributions, κ is approximately the inverse of the variance, $\kappa \approx \sigma^{-2}$. Here, the bootstrapped value of $\kappa_{eff,i,j}$ indicates the certainty with which we were able to measure the effective lighting direction, not the concentration of the observer's posterior on lighting directions.) Finally, we found the weight ratio w_{prior}/w_{cue} that gave a maximum-likelihood fit of Eq. S2 to the measured prior and effective lighting directions. That is, we found the value of w_{prior}/w_{cue} that minimized this negative log likelihood:

$$-\sum_{i=1}^{n}\sum_{j=1}^{6}\log\left[f_{VM}\left(\theta_{eff,i,j}-\theta_{prior,i},\arctan2\left(\sin\left(\theta_{cue,j}-\theta_{prior,i}\right),\left(w_{prior}/w_{cue}\right)+\cos\left(\theta_{cue,j}-\theta_{prior,i}\right)\right),\kappa_{eff,i,j}\right)\right].$$

Here, n is the number of observers, and f_{VM} (θ, μ, κ) is the von Mises probability density function.

Lighting Direction Estimation Experiment.

The six observers who ran in the weak lighting cue condition of the shape judgment experiment also ran in the lighting direction estimation experiment. This experiment had two 36-trial blocks, one with strong and one with weak lighting cues. The blocks were run after the shape judgment experiment, in random order. On each trial the direction of the point light source was chosen from a uniform random variable ranging from 0° to 360°. Observers indicated the dominant lighting direction in the scene by rotating a knob to adjust an onscreen arrow (length 1.2°, located just below the scene). The stimuli were identical to those in the strong and weak conditions of the shape experiment, except that no ambiguous disks were shown. We found the circular SD of the direction errors in each condition, pooled across observers.

Discussion

Example of a Bump-vs.-Orientation Plot. Fig. S1 shows the proportion of "bump" responses as a function of target disk orientation, for a typical observer (initials XL in Table S1) in the strong lighting cue condition. For this observer, shaded disks were most likely to be seen as bumps when they were rotated so that their brighter halves were in the same direction as the scene lighting, indicating that the effective lighting direction closely tracked the true lighting direction.

Directions of Individual Observers' Light-from-Above Priors. Fig. S2 shows the directions of individual observers' light-from-above priors (i.e., the effective lighting direction in the no cue block). Some observers' priors were approximately overhead, but others were more than 90° left of vertical. Previous studies have also found a leftward bias in the light-from-above prior and have found large individual differences in the prior's direction. Furthermore, in our strong cue condition the blocks were ordered randomly, and recent experience in blocks with strong lighting cues may have temporarily shifted observers' priors in the no cue block. If this occurred, it is unlikely that it affected our results substantially, because we found that the light-from-above prior had almost no influence in the strong cue condition. Nevertheless, in the remaining conditions we measured the light-from-above prior in the first block, to avoid such learning effects. In the remaining conditions the prior directions spanned a similarly wide range, suggesting that the wide range of priors in the strong cue condition was not due to learning.

Our findings show that the light-from-above prior is extremely weak. This means that when measuring the direction of an observer's light-from-above prior, it is important to completely eliminate lighting direction cues. We ran our experiments in a dark room, but a very small amount of light entered the room under a door that was behind observers and to their left. It is conceivable (although we think unlikely) that this faint light source affected observers' responses. If it did, it would bias our estimates of observers' priors toward the bottom left. In any case, such a bias would not undermine our conclusions: it does not change the fact that the effective lighting direction closely tracked strong lighting cues and partly tracked even very weak lighting cues. Furthermore, any effect of such a faint light source would support our conclusion that the light-from-above prior is remarkably weak.

Individual Differences in Prior-to-Cue Weight Ratio. Fig. 3 in the main text shows that individual differences were small in the strong cue and no local cue conditions but larger in the weak cue condition. Fig. S3 shows the fit of the cue combination model (Eq. 1, main text) to individual observers' data, and Table S1 reports the fitted parameters. Just as previous studies have found large individual differences in the direction of the light-from-above prior, we found large individual differences in the strength of the prior relative to lighting cues. These differences were mostly obscured by strong lighting cues (which were stronger than all observers' priors) but were evident when weak cues approximately matched the average prior strength. When weak lighting cues were in the opposite direction to the prior, the effective lighting direction was in the direction of the prior for some observers (e.g., light blue triangles in Fig. S3*B*) and in the direction of the lighting cues for others (e.g., green squares in Fig. S3*B*).

Three confidence intervals in Table S1 are reported as 0.00 ± 0.00. These unrealistically precise estimates occurred when an observer's effective lighting directions were even further from the prior than the lighting cue direction was (e.g., prior at 0°, lighting cues at 30°, effective lighting direction at 40°). The vector sum model cannot accommodate this unusual pattern of results, and the closest it can come is to assign a prior-to-cue weight ratio of 0.00 (i.e., make the cues completely override the prior). In such cases a weight ratio of 0.00 was obtained on each bootstrap iteration, and the resulting confidence interval had zero width. Fig. S3 shows that such overshoots in effective lighting direction were not very large or frequent, and we believe that they occurred simply owing to sampling error in measuring the effective lighting direction.

Variability of Effective Lighting Direction. In the main text we pointed out that in the weak cue condition the effective lighting direction was highly variable when the lighting cue direction was opposite to the prior direction (Fig. 3*B*). This is consistent with the predictions of the vector sum model, for two reasons. First, the model predicts that when the prior and cue directions are directly opposed, the effective lighting direction is in the direction of whichever has the greater weight. Thus, when the prior and cue weights are approximately matched ($w_{prior} \approx w_{cue}$), as with our weak cue stimuli, different observers who assign slightly different weights to the prior and the cues may have very different effective lighting directions, depending on whether $w_{prior} > w_{cue}$ or $w_{prior} < w_{cue}$. Second, the model predicts that when the prior and the lighting cues are approximately equally weighted and in opposite directions, observers can estimate an effective lighting direction, but their estimate is uncertain. See Murray and Morgenstern (1) for further discussion of these properties of the vector sum model.

Previous Investigations of the Role of Lighting Direction Cues. Our findings account for the inconsistent conclusions of previous studies. We find that only strong lighting cues completely override the light-from-above prior. Previous studies that concluded that lighting cues affect shape judgments used stimuli much like ours, with complex illuminated objects providing strong lighting cues immediately adjacent to ambiguous shaded disks: Berbaum et al. (2) showed a hand in front of ambiguous disks, and Ramachandran (3) showed disks attached to a face and a corrugated cylinder. [Rittenhouse (4) and Brewster (5) reached similar conclusions, but they did not describe their stimuli in detail.] Erens et al. (6) and Wagemans et al. (7) concluded that illumination cues do not affect shape perception, but their stimuli were highly simplified and did not give a vivid impression of 3D shape or illumination. Yonas et al. (8) and Jenkin et al. (9) found that lighting cues had a small effect on shape judgments, but less than other factors like head orientation; Yonas et al. placed light sources next to flat displays of shaded disks with no other objects nearby, and Jenkin et al. showed shaded disks in a room lit diffusely by fluorescent lights, with no directly illuminated objects adjacent to the disks. Thus, consistent with our findings, the magnitude of lighting effects in previous studies seems have been determined by the strength of the lighting cues.

Previous Estimates of Prior Strength. Stone et al. (10) and Mamassian and Landy (11) have estimated the probability distribution of human observers' light-from-above priors. Although informative, neither study measured the strength of the light-from-above prior in the sense we are interested in here, for two reasons.

First, both studies were based on models that rely on assumptions about the degree of randomness in observers' responses. Stone et al. chose an arbitrary value for a decision noise parameter in their model. Mamassian and Landy assumed a probability-matching decision rule, whereby if an observer believes that the posterior probability of response R being correct is p, then the observer gives response R with probability p. In both cases, the strength of the estimated prior (e.g., its circular SD) depends on the assumed level of randomness. Because both studies' assumptions about randomness were not tested and were made largely for convenience, the probability distributions they estimated must be regarded as tentative and probably do not characterize the light-from-above prior in an absolute sense. Similarly, O'Shea et al. (12) arrived at an estimate of the von Mises concentration parameter of the light-from-above prior, but this estimate relied on the assumption that the concentration of the directional information provided by one of their lighting cues was 1.0.

Second, and more importantly, neither study examined conflicts between the light-from-above prior and visual lighting cues (nor was this their goal). Even if we knew the probability distribution of the light-from-above prior, to infer from this whether the prior or lighting cues dominate in illuminated scenes we would also have to know the probability distribution of lighting directions that the visual system infers from visual lighting cues. Instead, in the present study we evaluated the relative strength of the prior and lighting cues by measuring the effective lighting direction when they were in conflict.

1. Murray RF, Morgenstern Y (2010) Cue combination on the circle and the sphere. *J Vis* 10(11):15, 1–11.
2. Berbaum K, Bever T, Chung CS (1984) Extending the perception of shape from known to unknown shading. *Perception* 13:479–488.
3. Ramachandran VS (1988) Perceiving shape from shading. *Sci Am* 259:76–83.
4. Rittenhouse D (1786) Explanation of an optical deception. *T. Am. Philos. Soc.* 2:37–42.
5. Brewster D (1826) On the optical illusion of the conversion of cameos into intaglios, and of intaglios into cameos, with an account of other analogous phenomena. *Edinb J Sci* 4:99–108.
6. Erens RGF, Kappers AML, Koenderink JJ (1993) Estimating local shape from shading in the presence of global shading. *Percept Psychophys* 54:334–342.
7. Wagemans J, van Doorn AJ, Koenderink JJ (2010) The shading cue in context. *i-Perception* 1:159–178.
8. Yonas A, Kuskowski M, Sternfels S (1979) The role of frames of reference in the development of responsiveness to shading information. *Child Dev* 50:495–500.
9. Jenkin HL, Jenkin MR, Dyde RT, Harris LR (2004) Shape-from-shading depends on visual, gravitational, and body-orientation cues. *Perception* 33:1453–1461.
10. Stone JV, Kerrigan IS, Porrill J (2009) Where is the light? Bayesian perceptual priors for lighting direction. *Proceedings of the Royal Society B: Biological Sciences* 276: 1797–1804.
11. Mamassian P, Landy MS (2001) Interaction of visual prior constraints. *Vision Res* 41: 2653–2668.
12. O'Shea JP, Agrawala M, Banks MS (2010) The influence of shape cues on the perception of lighting direction. *J Vis* 10(12):21, 1–21.

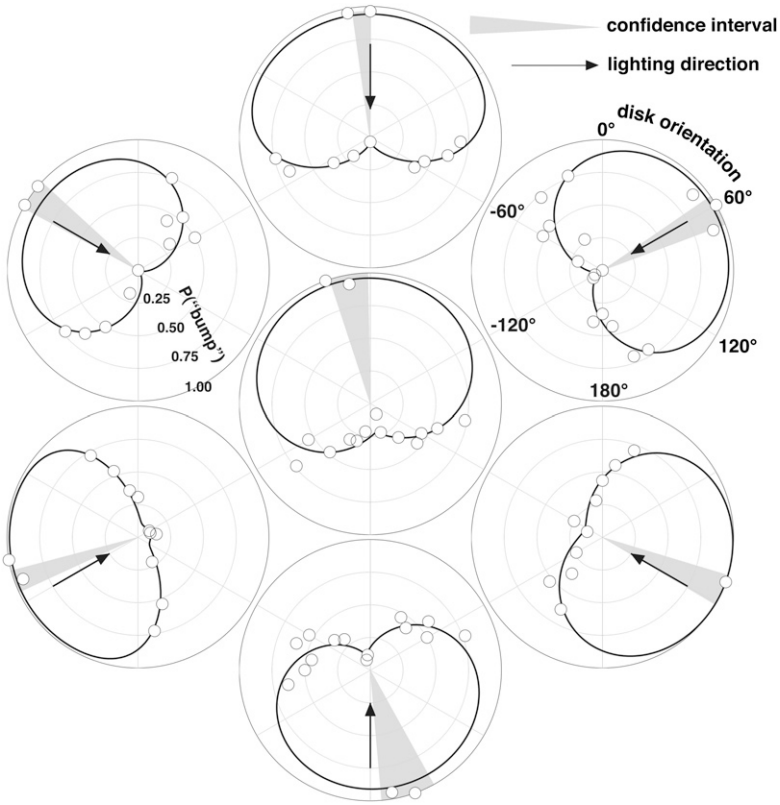

Fig. S1. Polar plots showing the proportion of "bump" responses as a function of disk orientation, for a typical observer (initials XL in Table S1) in the strong cue condition (Fig. 2*A*). Each polar plot corresponds to a single lighting condition. The angular position of each white data point represents a disk orientation, and the radial distance indicates the proportion of "bump" responses the observer gave to disks at that orientation. In each plot the arrow shows the true lighting direction, the solid curve is the fit of Eq. **S1** to the data points, and the gray wedge shows the bootstrapped 95% confidence interval for the orientation where the fitted curve peaks. The center plot shows results from the block without lighting cues (Fig. 2*C*).

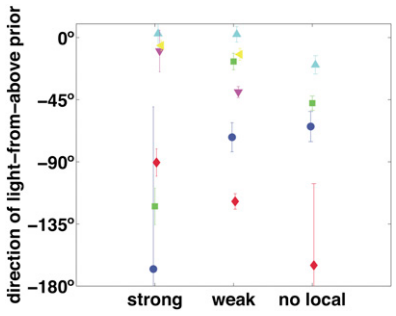

Fig. S2. Direction of individual observers' light-from-above priors in the strong, weak, and no local lighting cue conditions (related to Fig. 3). An angle of 0° is directly overhead, and negative angles are left of vertical. Error bars indicate bootstrapped SE. Data points have been jittered horizontally so that the error bars do not overlap. Each observer's data point has the same symbol and color as in Fig. S3 and Table S1. All data points are from different observers, except the two yellow triangles, which are from the first author (Y.M.).

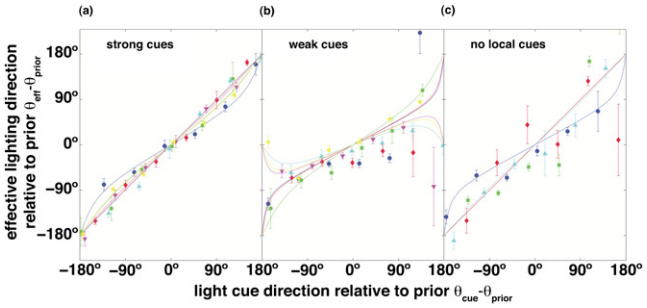

Fig. S3. Individual differences in the prior-to-cue weight ratio w_{prior}/w_{cue} (related to Fig. 3). Different colors correspond to different observers. Each solid line is fitted to the data points of the same color. Table S1 reports the fitted parameter values.

Table S1. Maximum-likelihood estimates and 95% confidence intervals for individual observers' prior-to-cue weight ratios, related to Fig. 3

Strong cues		Weak cues		No local cues		
Observer	w_{prior}/w_{cue}	Observer	w_{prior}/w_{cue}	Observer	w_{prior}/w_{cue}	Symbol in Fig. S3
DS	0.63 ± 0.10	CO	0.89 ± 0.16	JF	0.84 ± 0.41	●
II	0.27 ± 0.16	DC	0.52 ± 0.15	NS	0.00 ± 0.00	■
IV	0.03 ± 0.08	DS	0.86 ± 0.40	RD	0.00 ± 0.00	◆
VM	0.00 ± 0.00	KF	1.65 ± 0.40	XG	0.08 ± 0.10	▲
XL	0.12 ± 0.11	SF	1.12 ± 0.19			▼
YM	0.29 ± 0.09	YM	1.30 ± 0.13			◀

These are the parameters of the fitted curves in Fig. S3.