Cognition 127 (2013) 99–104

Contents lists available at SciVerse ScienceDirect

Cognition

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

Brief article

Learning different light prior distributions for different contexts

Iona S. Kerrigan, Wendy J. Adams *

Psychology, University of Southampton, Southampton SO17 1BJ, UK

A R T I C L E I N F O

Article history:
Received 13 June 2012
Revised 26 November 2012
Accepted 17 December 2012
Available online 30 January 2013

Keywords:
Shape-from-shading
Human visual perception
Light priors
Context specific learning

A B S T R A C T

The pattern of shading across an image can provide a rich sense of object shape. Our ability to use shading information is remarkable given the infinite possible combinations of illumination, shape and reflectance that could have produced any given image. Illumination can change dramatically across environments (e.g. indoor vs. outdoor) and times of day (e.g. mid-day vs. sunset). Here we show that people can learn to associate particular illumination conditions with particular contexts, to aid shape-from-shading. Following a few hours of visual–haptic training, observers modified their shape estimates according to the illumination expected in the prevailing context. Our observers learned that red lighting was roughly overhead (consistent with their previous assumption of lighting direction), whereas green lighting was shifted by 10°. Greater learning occurred when training for the two contexts (red or green light) was intermingled rather than when it was sequentially blocked.

1. Introduction

Humans cope with reddish illumination at sunset or flickering coloured lights at the disco – managing to decompose shading patterns into reflectance and shape variations – but how? Our impressively robust ability to estimate our surroundings, given complex and ambiguous retinal input relies heavily on prior knowledge – we bias perceptual estimates toward the most likely scenes. For example, we bias estimates of illumination direction toward overhead (e.g. Adams, 2007; Kleffner & Ramachandran, 1992) and estimates of surface shape toward convexity (Adams & Mamassian, 2004; Langer & Bülthoff, 2001) in alignment with the statistics of our environment (Potetz & Lee, 2003). Such assumptions, or 'priors', facilitate the notoriously underconstrained problem of recovering shape-from-shading. Here we investigate whether observers can further refine this process by learning that particular illumination conditions are more likely in particular contexts.

For optimal performance, humans should (i) respond to long-term changes in scene statistics by updating their priors and (ii) select the correct prior for a given context. We know that humans do the former: in contrast with chickens (Hershberger, 1970), human observers change their light prior in response to appropriate haptic (Adams, Graf, & Ernst, 2004) or visual feedback (Adams, Kerrigan, & Graf, 2010). Here we ask whether humans also do the latter: can we learn different prior assumptions for different contexts? There is no clear consensus: although Adams et al. (2004) found that a modified light-prior generalised to novel stimuli, Adams et al. (2010) noted that modified light-priors were retained for several weeks beyond training, after observers had returned to their normal environment, in which lighting was presumably, on average, overhead. This latter finding suggests that observers learnt separate, context-dependent light priors, with the experimental set-up acting as a contextual cue.

Here we ask whether humans can learn two light priors, each invoked by a different illumination colour. To induce colour-dependent learning, visual–haptic feedback was modulated by the simulated illumination colour: when scenes were illuminated by red light, feedback was consistent with the observer's baseline light prior distribution. In contrast, under green illumination, feedback was consistent with a new lighting distribution.

* Corresponding author. Tel.: +44 02380 593629.
E-mail addresses: i.s.kerrigan@soton.ac.uk (I.S. Kerrigan), w.adams@soton.ac.uk (W.J. Adams).

0010-0277/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cognition.2012.12.011

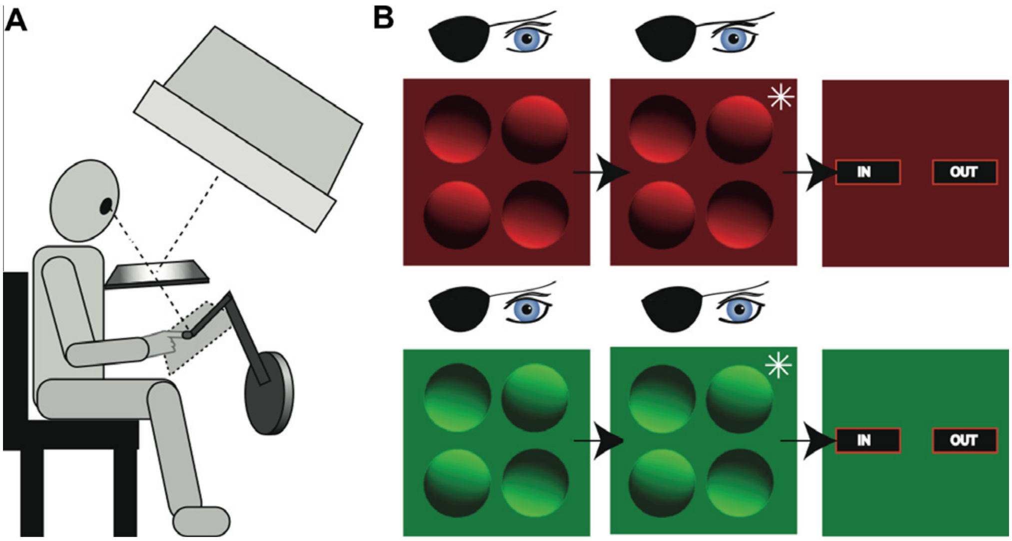

Fig. 1. Apparatus and visual test trials: (A) The visual–haptic experimental set-up. (B) Examples of visual-only test trials: the simulated lighting is either red (upper row) or green (lower row). Observers briefly viewed the four shaded discs (total presentation time 1.2 s, target object cued after 600 ms) before indicating whether the cued object was concave or convex (in or out). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2. Materials and methods

2.1. Apparatus and stimuli

Observers simultaneously viewed and felt virtual objects (see Fig. 1a). Haptic scenes were presented via a 'thimble gimbal' attached to a force-feedback device (Ghost libraries, PHANToM, SensAble Technologies). Visual stimuli (Figs. 1b and 2d–g), generated using OpenGL, were presented via a front-silvered mirror. Their perceived location (at a visual distance of 56 cm) matched the location of the haptic stimuli, giving the impression of a single visual–haptic scene. A headrest and bite bar maintained head position and an eye patch eliminated binocular depth cues. The room was completely dark, other than the light emitted by the visual display.

2.2. Visual test trials

Pre- and post-training trials contained solely visual (no haptic) information. Observers viewed four shaded discs, each subtending 5.6° and offset from the screen's centre by 5.3° (see Fig. 1). Each disc was consistent with a hemisphere squashed in depth by a factor of 2, illuminated by a distant light source. The slant of the light source (the angle between the lighting vector and the screen normal) was 68.2°. The light source tilt (the angle between the projected lighting vector and the vertical axis in the plane of the screen, θ) varied across trials. This illumination tilt, with object shape (convex vs. concave) determined the shading orientation of each disc. Within each trial, one, two or three discs had a shading gradient direction of θ and the remaining disc(s) had a shading gradient of $\theta + 180°$, such that observers generally perceived both convex and concave objects to be present. The simulated scene was white, with either a red or green simulated light source although stimuli were equally consistent with red and green scenes illuminated by white light.

Observers judged the shape (concave vs. convex) of one object (cued by a star). The observer's light prior was estimated from the set of 288 visual trials (24 equally spaced θ values × 2 colours × 6 repetitions), lasting approximately 10–15 min (see Fig. 2a).

2.3. Training trials

Visual–haptic training was similar to that used previously (e.g. Adams et al., 2004, see Fig. 2d–g). Observers viewed four shaded discs (as in test trials), but also explored the scene haptically by running a finger (in a thimble gimbal) over the simulated objects. This haptic information disambiguated each object's shape, and thus also the lighting direction. However, the relationship between shading orientation and haptic shape depended on colour (see Fig. 2b). On 'red' trials, stimuli were consistent with the observer's baseline light prior; haptic shape matched the observer's pre-training shape responses. On 'green' trials, however, the lighting direction was drawn from a range shifted by ±30° relative to the observer's baseline prior (13 observers were assigned a +30° shift, 13 a −30° shift). Thus, on 'green' trials, some objects previously perceived as convex now felt concave, and vice versa.

It is important to note that haptic feedback did not introduce an association between colour and shape; $p(\text{haptically convex}|\text{green}) = p(\text{haptically convex}|\text{red})$. Rather, for perception to become aligned with haptic feedback, the observer would have to learn a relationship between illumination direction and colour.

After haptically exploring the scene for a minimum of 7 s, including 'touching' all four objects, the observer pressed a button to continue. One of the objects then appeared visually (without haptics) in the centre of the