Available online at www.sciencedirect.com
ScienceDirect
Vision Research 47 (2007) 919–924
Vision Research
www.elsevier.com/locate/visres

Competition between newly recruited and pre-existing visual cues during the construction of visual appearance

Benjamin T. Backus [a,*], Qi Haijiang [b]

[a] Department of Psychology, University of Pennsylvania, 3401 Walnut St., C-Wing, Room 302-C, Philadelphia, PA 19104-6228, USA
[b] Department of Bioengineering, University of Pennsylvania, 3401 Walnut St., C-Wing, Room 302-C, Philadelphia, PA 19104-6228, USA

Received 11 August 2006; received in revised form 10 November 2006

Abstract

During perception, conflicting visual cues often trade against each other. Recent cue recruitment experiments show that the visual system can be conditioned to use artificial visual cues during the perception of a bistable stimulus. Does the visual system treat the new cue as an independent source of information, separate from the long-trusted cues that were used to train it? If so, presence of the long-trusted cue should not be sufficient to block the new cue's effect. Here, we show that a newly recruited cue (stimulus location) and a long-trusted, pre-existing cue (binocular disparity) trade against each other: they contribute simultaneously to the direction of perceived 3D rotation of a Necker cube. We also show that the new position cue was based primarily on retinal position, so early visual areas may mediate the cue's effect.

Keywords: Cue recruitment; Cue combination; Cue learning; Binocular disparity; Necker cube

1. Introduction

The visual system extracts signals ("cues") from the retinal images in order to construct visual percepts, and in certain cases classical (Pavlovian) conditioning procedures can be used to teach the visual system to use new cues (Haijiang, Saunders, Stone, & Backus, 2006). This learning has been demonstrated by showing that the new cue is effective in test stimuli that do not contain the long-trusted cues that were used during training. There are several reasons to investigate whether a newly learned cue is effective in stimuli that *do* contain those long-trusted cues. First, it could rule out the possibility that the learning in previous experiments manifested itself *only* because long-trusted cues were absent from test stimuli. If such were the case, it would mean that a newly recruited cue is used in a manner that is qualitatively different from long-trusted cues, so it would be harder to make a general argument that cue recruitment experiments reveal how our perceptual systems come to use cues in the natural environment. Second, it is useful to know if the effectiveness of a new cue can be measured by putting it into competition with other cues. For example, this would allow the experimenter to measure (and track) additional learning that occurs for a new cue after it becomes 100% effective when used by itself. Finally, showing that the new cue and a pre-existing cue are effective simultaneously would support the claim of Haijiang et al. (2006) that the new cue affects the appearance of the stimulus rather than some post-perceptual decision about how to respond.

Under normal conditions, several natural cues are often simultaneously informative about a given aspect of a scene. For example, binocular disparity and perspective cues can both be informative about surface slant (e.g., Banks & Backus, 1998). These redundant cues do not always agree with each other: a given cue need not co-vary perfectly with the property of the world about which it is informative (Brunswik, 1956), and cue measurement also adds noise. As

* Corresponding author. Fax: +1 215 746 6848.
E-mail addresses: backus@psych.upenn.edu (B.T. Backus), haijiang@seas.upenn.edu (Q. Haijiang).

0042-6989/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2006.12.008

a consequence the visual system must adopt a strategy to combine or choose among discrepant cues. Existing models of cue combination suppose that these strategies are near-ideal. Such models describe the combined use of redundant cues within a general framework of probabilistic inference (Backus & Banks, 1999; Brunswik, 1956; Hebb, 1949; Helmholtz, 1910/1925) or, more concretely, Bayesian inference (Adams & Mamassian, 2004; Geisler & Kersten, 2002; Hochberg & Krantz, 2004; Kersten, Mamassian, & Yuille, 2004; Knill & Richards, 1996; Maloney, 2002; van Ee, Adams, & Mamassian, 2003; Weiss, Simoncelli, & Adelson, 2002). Within this framework, cue recruitment can be described as utilization of a new signal for the purpose of estimating some property of the world, as reflected during perception by that property's *appearance*, to improve the system's estimate of the property.

Unfortunately, adopting such a framework is not sufficient for predicting how the system will combine a newly recruited cue with pre-existing cues. Is it more optimal to give weight or not to give weight to a new cue, when long-trusted cues are also present? In the case of natural cues, an experimenter might in principle measure, estimate, or at least sample the multivariate likelihood function that describes the joint probability between the cues' values and states of the world to be estimated, and from this determine the best estimate for any configuration of cue values. In the case of training stimuli that contain just two values for each cue (as in the present study), it is impossible to predict the system's cue combination strategy merely from an assumption that learning behavior is ideal.

It thus becomes interesting to know why, and to what extent, the system generalizes from the limited training sample to novel configurations of cues. Experimentally, we can infer the visual system's default strategy, and this can reveal the implicit assumption made by the system about how new cues ought normally to be combined with long-trusted cues. Long-trusted cues often do trade against each other, and more specifically, a perceptual attribute can often be modeled (for a moderately large range of cue values) as a weighted average of the values specified by each cue separately (Backus, Banks, van Ee, & Crowell, 1999; Backus & Matza-Brown, 2003; Buell & Hafter, 1991; Clark & Yuille, 1990; Johnston, Cumming, & Parker, 1993; Landy, Maloney, Johnston, & Young, 1995; Young, Landy, & Maloney, 1993). To anticipate, we found that a new cue and a long-trusted cue are likewise both given weight during the construction of appearance when both are present.

Our experiments used a perceptually bistable rotating Necker cube stimulus. We are not interested here to explain perceptual bistability *per se*, either why it occurs or the time course of alternation during prolonged viewing (e.g., Carter & Pettigrew, 2003; Mamassian & Goutcher, 2005). Bistable stimuli are useful in cue recruitment experiments because small amounts of learning can result in measurable perceptual biases (Haijiang et al., 2006; Wallach & Austin, 1954), and because it is easy for observers to reliably report the appearance of a binary perceptual attribute. Importantly, the direction of perceived 3D rotation of a Necker cube can be forced using binocular disparity cues (Dosher, Sperling, & Wurst, 1986) or by new cues (Haijiang et al., 2006). We trained observers' visual systems to use position as a cue, and presented test stimuli that contained this new cue but also various amounts of binocular disparity, to find out how the two cues would interact.

Experiment 1 tested the basic hypotheses. Experiments 2 and 3 ruled out the possibility that results in Experiment 1 were caused exclusively by short term position-dependent priming. Experiment 4 tested whether the bias caused by the position cue was a consequence of retinal position or position in the world.

2. General methods

2.1. Participants

Participants ("trainees") were undergraduates from the University of Pennsylvania who passed a test of stereoacuity and gave correct responses at least 90% of the time on training trials. Stereoacuity was assessed using anaglyph displays that contained nine diamonds in a 3×3 array subtending 1.8°. In each display one diamond had different disparity and the trainee had to identify it; passing the test required correct identification of a 5 arcmin (and greater) disparity difference and ~80% of potential trainees passed. Trainees were paid to participate and were naive to the hypotheses of the experiment.

2.2. Apparatus and stimuli

Stimuli were stereo movies that depicted a wire-frame cube rotating about a vertical axis. On each 2s trial, the trainee indicated the cube's rotation direction by judging whether a random-direction probe dot moved in the same direction as the front or the back of the cube (Fig. 1). *Training trials* contained two long-trusted cues: binocular disparity and an opaque occluder that passed through the cube. These cues successfully disambiguated the rotation direction. *Test trials* were similar to training trials, but

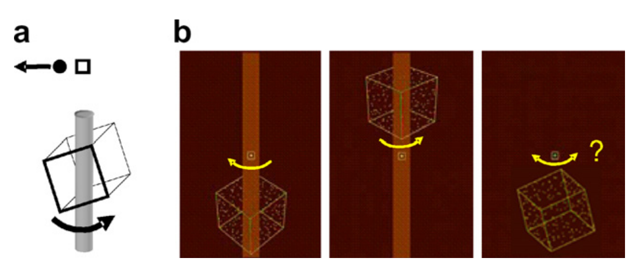

Fig. 1. Task (a) and stimulus (b). A dot was displayed near the fixation mark (a). On each trial it moved either left or right with equal probability, and the trainee pressed "2" if the dot appeared to move with the front of the cube, or "8" if it moved with the back. For the figure shown, the correct answer would be "8". The left and middle pictures in (b) show training stimuli, containing both the long-trusted depth cues (stereo and occlusion) and the new cue (stimulus position). In the configuration shown, the cube rotates leftward when it is below the fixation mark and rightward when it is above. By arbitrary definition, position and rotation direction are "positively" correlated in these panels; negative correlation for these cues is simply the opposite (rightward below fixation mark, leftwards above). The rightmost panel shows a test trial; it contains the position cue and a weak disparity cue that favors rightward rotation. Yellow arrows indicate object rotation direction. (For anaglyph images the reader is referred to Fig. 1 in the web version of this article.)