

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.

© 2006 Elsevier Ltd. All rights reserved.

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).

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 (Doshier, 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 $\sim 80\%$ 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.)

did not contain the occluder, and the disparity signal was reduced in magnitude to be 4/6, 1/6, -1/6, or -4/6 of the natural value ($\pm 6/6$). Positive values specified that rightward moving elements of the cube were stereoscopically nearer than leftward moving elements, which evoked “rightward” rotation of the cube.

Stimuli were red/green anaglyphs, presented by rear-projection onto a large display screen 200 cm from the trainee. They depicted a wire-frame cube with random dots on its faces, rotating in depth about a vertical axis at 0.75 rad/s. The stimulus on a given trial was visible for 2s, and after that until the trainee responded. The separate horizontally-moving probe dot was shown at the middle of the display and it specified the response mapping on that trial. The cube appeared 33 cm above or below the center of the screen (Top or Bottom position, respectively), and subtended 14° of visual angle on a screen 1.2 m × 1.8 m in size. Trainees were instructed to fixate a small square at the center of the screen; eye position was not monitored (the logic of the experimental design does not require accurate fixation, but see Experiment 4). Other details are described in Haijiang et al. (2006).

2.3. Task and procedure

Each session contained 240 training trials and 240 test trials presented in alternation. Experiments 1, 3, and 4 were conducted in single sessions lasting about 45 min; Experiment 2 was conducted in two sessions. Because the task was to judge whether the probe dot moved the same direction as the front or the back of the cube, responses were uncorrelated with the cues. No feedback was given. Spontaneous reversals were reported as well but were very rare because display time was short (Long, Toppino, & Kostenbauder, 1983). For each type of trial (training or test) we allowed at most three consecutive trials at the same location. The experiment was conducted according to human subject protocols approved by the Institutional Review Board at the University of Pennsylvania.

3. Experiment 1: tradeoff between new and long-trusted cue

The test group saw training stimuli in which the position cue (“POSN”) was 100% correlated with stereo and occlusion. Thus, the cube always appeared to rotate one way at the Top position and the other way at the Bottom position. By definition, position and rotation were 100% “positively” correlated on training trials if the cube always rotated rightward when it was above the fixation mark and leftward when it was below the fixation mark; they were 100% “negatively” correlated if rotation was always rightward below the mark and leftwards above. The control group saw training stimuli in which position was randomly paired with the direction specified by stereo and occlusion (“no correlation”). Two groups of eight trainees ran in the experiment: four trainees ran in each of the two correlated conditions, and eight trainees ran in the uncorrelated condition.

Each panel in Fig. 2 shows the test-trial data for one trainee. “Percent judged as rightward rotation” is plotted as a function of disparity for Top and Bottom test trials, respectively. Two cumulative normal curves are fitted to the data for each trainee. For all trainees in the test condition (left side of Fig. 2), POSN caused a horizontal shift of the psychometric function in the expected direction. The absence of similarly large effects in the control condition provides further evidence that effects in the test condition did not result from pre-existing biases nor from viewing the experimental stimuli *per se*.

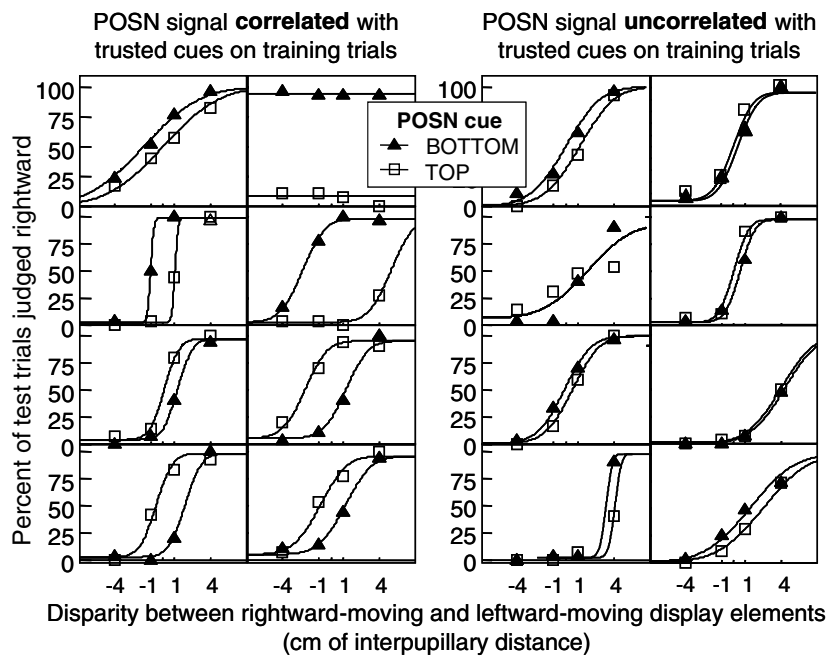


Fig. 2. Results from Experiment 1: tradeoff between newly recruited cue and long-trusted cue. POSN was correlated with long-trusted cues on training trials for eight trainees in the test condition (left graphs), and uncorrelated for eight trainees in the control condition (right graphs). The abscissa is the disparity added to the rotating cube on test trials. The ordinate is the percentage of test trials seen as having right-hand rotation. Each data point was computed from 30 test trials. Curves are maximum likelihood fits for a pair of cumulative Gaussians sharing a common slope (so there were three free parameters in the fit: μ_1 , μ_2 , and σ). In the test condition, half of the trainees (upper four panels on the left) viewed training trials with positive correlation and half (lower four panels on the left) viewed training trials with negative correlation; hence the Top and Bottom data are shifted in different directions in the upper four panels as compared to the lower four panels. This separation is not observed for the control group.

The *horizontal separation* of the curves shows how much disparity was needed to null the effect of POSN, which indicates the system's reliance on POSN relative to its reliance on disparity, i.e., how these two cues traded against one another. The *vertical separation* of the curves (at the point of greatest separation) shows the system's reliance on POSN, relative to binocular disparity and whatever unmodeled signals may have contributed to variability in the perceptual decision (such as signals from endogenous processes that resolve the competition between the two viable interpretations of the stimulus). The *steepness* of the curves indicates the system's sensitivity to disparity. These three statistics are not independent: if two trainees have the same horizontal separation, the trainee with steeper curves has greater reliance on both disparity and POSN (relative to sources of noise), while reliance on POSN relative to disparity is the same for both trainees.

It is hard to argue that disparity *per se* was directly visible in these stimuli. Instead, disparity became visible to the observer only through its effect on the apparent depth and rotation direction of the cube. Thus, the horizontal separation of the psychometric functions in Fig. 2 (left side, seven out of eight observers excepting top right panel) demonstrates a tradeoff between disparity and POSN during or prior to the construction of appearance. The results therefore reinforce the claim of Haijiang et al. (2006) that the appearance of a stimulus can be made to depend on a newly recruited cue.

From Experiment 1, we also conclude that a newly recruited cue and a long-trusted cue can influence appearance simultaneously. It is not necessary to remove long-trusted cues from the stimulus, as was done in Haijiang et al. (2006), to see the new cue's effect. We also conclude that the effectiveness of a new cue can be measured by the effect it has when put into conflict with a long-trusted cue.

4. Experiments 2 and 3: long-lasting nature of learning

Short term factors might cause position-dependent biases leading to separated psychometric functions such as observed for the test group in Fig. 2. Indeed, a Necker cube's perceived rotation direction can be influenced by recent trials at the same location (Long & Toppino, 1994). But normally by "cue recruitment" we mean a long-term modification of perception. Long-term modification was demonstrated by Haijiang et al. (2006) for monocular test stimuli, but not for stimuli containing both the POSN cue and binocular disparity. Experiments 2 and 3 test whether long-term learning contributes to the separation of the curves.

In Experiment 2, four trainees each ran in Experiment 1 two times, on successive days. The left panels of Fig. 3 show that trainees relied more on the POSN cue during their second session than during their first session, which demonstrates cumulative learning across days.

Experiment 3 was a counterconditioning experiment, in which the correlation between the POSN cue and the long-trusted cues was reversed half way through the session (after 240 total trials). This manipulation tests whether

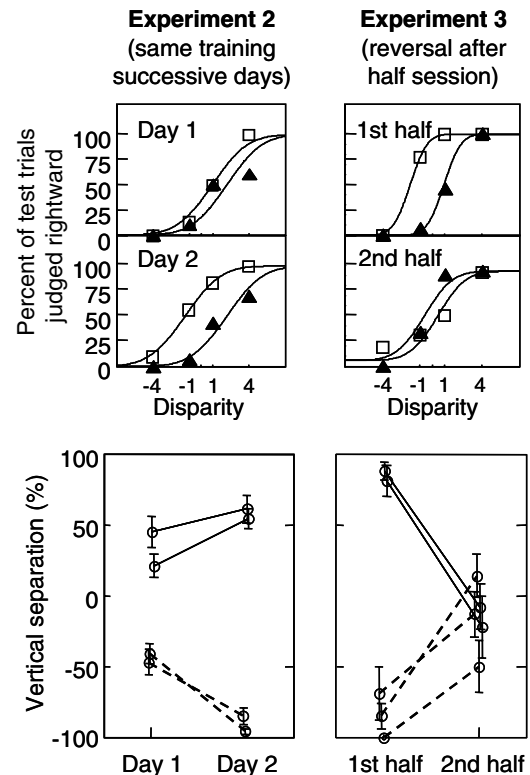


Fig. 3. Results from Experiments 2 and 3, showing that the new POSN cue had a long-lasting effect. Top panels show data from a single trainee in each experiment, respectively, in the format of Fig. 2, for Days 1 and 2 (Experiment 2, left) or for the first and second half of the session (Experiment 3, right). The bottom panels plot vertical separation between the two fitted psychometric functions for all trainees in each experiment, respectively. Initial training in each experiment was counterbalanced across trainees (half positive correlation, solid lines; half negative, dashed lines). In Experiment 2, all trainees showed additional reliance on POSN in the second session as compared to the first ($p < .01$ for three of the four trainees, normal test using estimated standard errors). In Experiment 3 the correlation between POSN and long-trusted cues was reversed half way through the experiment. The first 20% of each half session is excluded from the analysis. All trainees unlearned the initial correlation to some extent during the second half of the session, and some trainees achieved reversal, but no trainee learned the reversed correlation as strongly as the initial correlation. Thus, learning from the first half lasted longer than just a few trials. Error bars are the standard deviation of 2000 bootstrap estimates (Efron & Tibshirani, 1993) of the vertical separation of the two curves, fitted using the same procedure as with the original data in Experiment 1. Small horizontal offsets in bottom panels are to make the data easier to see.

learning from the first half of the session contributes to the shifting of curves in the second half. All five of the trainees tested showed lower reliance on POSN during the second half of the session. We conclude that binocular disparity trades with the new (long term) POSN cue, not just with short term position-dependent biases (also variously called "hysteresis", "set", "momentum", or "priming").

5. Experiment 4: retinal position versus world position cues

Fixation was controlled (by instruction) in Experiments 1–3, so it is unclear whether the recruited cue was position

on the retina or position on the display screen. This distinction is interesting because if the POSN cue was primarily retinal position, it means that learning can occur relatively early within the hierarchy of visual processing, in contrast with a proposed bias for high-level learning in the case of learning to discriminate (Ahissar & Hochstein, 2004).

Stimuli were similar to those of Experiment 1 except that the cube appeared at one of three horizontally separated locations (but always next to the fixation mark for that trial). For trainees in the *retinal* training group, stereo and occlusion cues specified a rotation direction according to the side of fixation at which the cube appeared. For trainees in the *world* training group, these same long-trusted cues specified a consistent rotation direction at each of the three locations on the display screen (see Fig. 4).

Eight trainees were tested in each training condition (retinal or world). Retinal training led to highly significant ($p < .001$) cue recruitment in seven of the eight trainees. World training did not lead to significant cue recruitment in any trainee, nor did a significant effect emerge in the average across trainees. Errors of fixation cannot account for the data because fixation errors would have reduced the overlap in retinal position, but not world position, for training and test stimuli; this could attenuate learning of a retinal POSN cue but would be unlikely to augment it. While this experiment does not rule out the possibility that world position was also learned as a cue, it clearly shows that retinal position was the primary basis for the POSN cue's effectiveness in our experiments.

Much is known about the representation of stereomotion stimuli in retinotopic areas of visual cortex, based on both psychophysics (Anstis & Harris, 1974; Mayhew & Anstis, 1972; Nawrot & Blake, 1991) and direct measurement of neural activity (Backus, Fleet, Parker, & Heeger, 2001; Bradley, Chang, & Andersen, 1998; DeAngelis & Newsome, 1999; Dodd, Krug, Cumming, & Parker, 2001; Huk, Ress, & Heeger, 2001). Neurons in area MT of the macaque monkey are jointly tuned for retinal location, motion direction, and disparity (Bradley et al., 1998; DeAngelis & Newsome, 1999; Dodd et al., 2001). Schlack and




Fig. 4. Experiment 4 display configuration, that made rotation direction (as specified by long-trusted cues) contingent on either retinal or world position during training. Both fixation marks were visible at all times and the trainee was instructed which mark to fixate on each trial by a flashing circle (shown for right fixation mark). The rotating cube always appeared to the left or right of fixation (at one of the locations shown by the dashed marks). To train *retinal* position, the cube on training trials had rightward rotation if it appeared to the right of fixation, and leftward rotation if to the left of fixation (or vice versa, counterbalanced across trainees). To train *world* position, the cube on training trials had rightward rotation at the center location and leftward rotation at the two side locations (or vice versa). Different groups received retinal and world training, respectively.

Albright (2005) recently demonstrated associative learning that caused a static 2D pattern to elicit activity similar to real motion in macaque MT. Area MT (or V5) is therefore a candidate site for the neuronal changes that implement recruitment of the POSN cue in humans, and it may be possible to study the learning in an animal model.

In our experiments, two values of the new cue (POSN) were paired with two perceptual outcomes (directions of perceived 3D rotation). We do not yet know whether training at a single position has any effect; it seems likely that such training would bias appearance at the trained position to resemble the training stimuli, but the effect at other positions could either be generalization (bias in the trained direction), or an opposite bias (contrast effect), or neither. Simultaneously presented objects that rotate in 3D at different locations do influence each others' appearance (Eby, Loomis, & Solomon, 1989).

6. General discussion

Experiments 1–3 clearly demonstrate that both the new POSN cue and the long-trusted disparity cue influenced the appearance of test stimuli after training. The data are consistent with a model in which POSN and disparity had independent effects on appearance. This idea is an appealing one: it corresponds to a working assumption by the visual system that a newly discovered cue should be treated as depending on the world, and not on other cues. If that is the case, then probit analysis (Finney, 1971) is particularly appropriate for measuring the effect of a newly recruited cue. Doshier et al. (1986) used this analysis to demonstrate that long-trusted cues have additive effects as disambiguators of a Necker cube's perceived rotation. In fitting the data using probit analysis, we assumed that the two cues (POSN, disparity) in a given test stimulus had independent (additive) effects according to the levels of the cues (our experiments used two levels for POSN and four levels for disparity). Note that trial-to-trial error in the measurement of a cue, that might plausibly have added to variation in the effectiveness of the disparity cue, is not represented explicitly in this model; instead it contributed to the (normally distributed) noise term.

An additional theoretical viewpoint is worth mentioning in connection with this interpretation of the results. The system may, by default, use a new cue in the manner of a naive Bayes classifier (Lewis, 1998). This type of classifier assumes that various indicator variables (cues) are conditionally independent. That is to say, the system assumes that the POSN cue does not depend on disparity, but only on the rotation of the cube. Is it plausible that the system would make such an assumption when learning a new cue? There are several reasons why this might be the case. First, in the real world, it is sometimes true that multiple cues are generated by a given world property, with the cues being conditionally independent, so the system may assume this. Indeed, long-trusted cues are often combined in a weighted average ("weak fusion", Clark & Yuille, 1990; Landy et al., 1995), suggesting that an assumption of independence is Bayes-optimal in

many natural viewing situations. Second, the naive Bayes strategy is simple to implement and performs reasonably well for many practical problems (Lewis, 1998). Third, as previously noted, Haijiang et al. (2006) have already demonstrated that a new cue can be effective in the absence of the long-trusted cues with which it was paired during training. Thus the newly recruited cue was treated by the system as being informative in its own right.

Bistable stimuli that contain conflicting cues may prove useful to measure perceptual learning in future cue recruitment experiments. They also allow the experimenter to test models of perceptual cue combination: we have ruled out the possibility that disparity, when present, prevents utilization of the newly recruited POSN cue. The basis for the POSN cue's effect is primarily retinal position rather than position in the world, which suggests that the learning of this cue could be studied electrophysiologically in retinotopically organized areas of visual cortex.

Acknowledgments

This work was supported by NIH Grant R01 EY 013988 and by the Human Frontier Science Program. We thank two anonymous reviewers for their very helpful comments.

References

- Adams, W. J., & Mamassian, P. (2004). Bayesian combination of ambiguous shape cues. *Journal of Vision*, 4, 921–929.
- Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends in Cognitive Sciences*, 8, 457–464.
- Anstis, S. M., & Harris, J. P. (1974). Movement aftereffects contingent on binocular disparity. *Perception*, 3, 153–168.
- Backus, B. T., & Banks, M. S. (1999). Estimator reliability and distance scaling in stereoscopic slant perception. *Perception*, 28, 217–242.
- Backus, B. T., Banks, M. S., van Ee, R., & Crowell, J. A. (1999). Horizontal and vertical disparity, eye position, and stereoscopic slant perception. *Vision Research*, 39, 1143–1170.
- Backus, B. T., Fleet, D. J., Parker, A. J., & Heeger, D. J. (2001). Human cortical activity correlates with stereoscopic depth perception. *Journal of Neurophysiology*, 86, 2054–2068.
- Backus, B. T., & Matza-Brown, D. (2003). The contribution of vergence change to the measurement of relative disparity. *Journal of Vision*, 3, 737–750.
- Banks, M. S., & Backus, B. T. (1998). Extra-retinal and perspective cues cause the small range of the induced effect. *Vision Research*, 38, 187–194.
- Bradley, D. C., Chang, G. C., & Andersen, R. A. (1998). Encoding of three-dimensional structure-from-motion by primate area MT neurons. *Nature*, 392, 714–717.
- Brunswik, E. (1956). Perception and the representative design of psychological experiments. In (pp. 92, 96, 123–131). Berkeley, CA: University of California Press.
- Buell, T. N., & Hafter, E. R. (1991). Combination of binaural information across frequency bands. *Journal of the Acoustical Society of America*, 90, 1894–1900.
- Carter, O. L., & Pettigrew, J. D. (2003). A common oscillator for perceptual rivalries? *Perception*, 32, 295–305.
- Clark, J. J., & Yuille, A. L. (1990). *Data fusion for sensory information processing systems*. Boston: Kluwer.
- DeAngelis, G. C., & Newsome, W. T. (1999). Organization of disparity-selective neurons in macaque area MT. *Journal of Neuroscience*, 19, 1398–1415.
- Dodd, J. V., Krug, K., Cumming, B. G., & Parker, A. J. (2001). Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *Journal of Neuroscience*, 21, 4809–4821.
- Doshier, B. A., Sperling, G., & Wurst, S. A. (1986). Tradeoffs between stereopsis and proximity luminance covariance as determinants of perceived 3D structure. *Vision Research*, 26, 973–990.
- Eby, D. W., Loomis, J. M., & Solomon, E. M. (1989). Perceptual linkage of multiple objects rotating in depth. *Perception*, 18, 427–444.
- Efron, B., & Tibshirani, R. J. (1993). An introduction to the bootstrap. In *Monographs on statistics and applied probability* (57). New York: Chapman & Hall.
- Finney, D. J. (1971). *Probit analysis*. Cambridge, UK: Cambridge University Press.
- Geisler, W. S., & Kersten, D. (2002). Illusions, perception and Bayes. *Nature Neuroscience*, 5, 508–510.
- Haijiang, Q., Saunders, J. A., Stone, R. W., & Backus, B. T. (2006). Demonstration of cue recruitment: change in visual appearance by means of Pavlovian conditioning. *Proceedings of the National Academy of Sciences of the United States of America*, 103, 483–488.
- Hebb, D. O. (1949). *Organization of Behavior*. New York: Wiley.
- Helmholtz, H. v. (1910/1925). *Treatise on Physiological Optics* (Vol. III). New York: Optical Society of America.
- Hochberg, J. E., & Krantz, D. (2004). Brunswik and Bayes: A Review of the Essential Brunswik: Beginnings, Explications, Applications, by Kenneth R. Hammond, Thomas R. Stewart (pp. 785–787). New York: Oxford University Press, 2001. *Contemporary Psychology: APA Review of Books*, 49.
- Huk, A. C., Ress, D., & Heeger, D. J. (2001). Neuronal basis of the motion aftereffect reconsidered. *Neuron*, 32, 161–172.
- Johnston, E. B., Cumming, B. G., & Parker, A. J. (1993). Integration of depth modules: stereopsis and texture. *Vision Research*, 33, 813–826.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271–304.
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian Inference*. Cambridge, UK: Cambridge University Press.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Research*, 35, 389–412.
- Lewis, D. D. (1998). Naive Bayes at forty: the independence assumption in information retrieval. In *Tenth European Conference on Machine Learning* (pp. 4–15). Berlin: Springer.
- Long, G. M., Toppino, T. C., & Kostenbauder, J. F. (1983). As the cube turns: evidence for two processes in the perception of a dynamic reversible figure. *Perception & Psychophysics*, 34, 29–38.
- Long, G. M., & Toppino, T. C. (1994). A new twist on the rotating-trapezoid illusion: evidence for neural-adaptation effects. *Perception*, 23, 619–634.
- Maloney, L. T. (2002). Statistical decision theory and biological vision. In *Perception and the Physical World* (pp. 145–189). Chichester, UK: Wiley.
- Mamassian, P., & Goutcher, R. (2005). Temporal dynamics in bistable perception. *Journal of Vision*, 5, 361–375.
- Mayhew, J. E., & Anstis, S. M. (1972). Movement aftereffects contingent on color, intensity, and pattern. *Perception & Psychophysics*, 12, 77–85.
- Nawrot, M., & Blake, R. (1991). The interplay between stereopsis and structure from motion. *Perception & Psychophysics*, 49, 230–244.
- Schlack, A., & Albright, T. D. (2005). Associative Learning in cortical visual area MT of macaque monkeys. *Society for Neuroscience Abstracts*, 31.
- van Ee, R., Adams, W. J., & Mamassian, P. (2003). Bayesian modeling of cue interaction: bistability in stereoscopic slant perception. *Journal of the Optical Society of America A-Optics Image Science and Vision*, 20, 1398–1406.
- Wallach, H., & Austin, P. (1954). Recognition and the localization of visual traces. *American Journal of Psychology*, 67, 338–340.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5, 598–604.
- Young, M. J., Landy, M. S., & Maloney, L. T. (1993). A perturbation analysis of depth perception from combinations of texture and motion cues. *Vision Research*, 33, 2685–2696.