ELSEVIER

Amodal completion impairs stereoacuity discrimination

Fang Hou [a,1], Hongjing Lu [b,1], Yifeng Zhou [a,*], Zili Liu [b,*]

[a] *School of Life Sciences, University of Science and Technology of China, Hefei, Anhui 230027, PR China*
[b] *Department of Psychology, University of California, 1285 Franz Hall, Box 951563, Los Angeles, CA 90095, USA*

Abstract

Visual stimulus configuration can influence elementary visual processes. We provide empirical evidence to demonstrate this effect in stereoscopic depth discrimination. Two vertically aligned bars were presented in stereo such that one of them was closer to the human observer. Observers discriminated which of the two was closest. In the first, "occluded" condition, a horizontal bar, positioned closest in depth to the observer, was added to the display such that the two vertical bars perceptually completed to form a whole by connecting together behind the horizontal bar. In the second, control condition, the horizontal bar was placed furthest away from the observer such that there was a visible gap between the two vertical bars, which could no longer complete perceptually. We measured observers' psychometric functions using the method of constant stimuli, and found that their discrimination sensitivity d' was smaller when the two vertical bars perceptually completed than when they did not. We used a simple model to illustrate that when the two vertical bars perceptually completed, they also tended to be perceived as coplanar in the fronto-parallel plane. This consequence of completion made it more difficult to discriminate any difference in depth between the two vertical bars.
© 2005 Elsevier Ltd. All rights reserved.

Keywords: Amodal completion; Stereodiscrimination; Perceptual completion; Bayes; Discrimination sensitivity; Bias

1. Introduction

The hypothesis that visual perception is an inference process (Helmholtz, 1924) places a great deal of emphasis on the role of prior knowledge about the visual world. In modern terms, Bayesian prior probability distribution presumably plays as important a role as stimulus information in determining the eventual percept that the visual system settles on (Kersten & Yuille, 2003; Knill & Kersten, 1991; Knill & Richards, 1996). However, despite its fruitful outcomes (Feldman, 2000; Mamassian & Landy, 2001; Weiss, Simoncelli, & Adelson, 2002), this approach remains controversial.

In a prominent article, Nakayama and Shimojo (1992) argued that, when stimulus information is ambiguous and consistent with several different configurations of the physical world, the final percept is completely determined by the likelihood of the physical layout, assuming a generic viewpoint. In other words, although the prior probability of viewpoint is assumed existent and uniformly distributed, the prior probability of the physical layout is not considered or, even if considered, plays no role.

In the current study, we will demonstrate that the prior probability distribution of physical layout is not only influenced by stereoscopic depth discrimination (and vice versa), but that it can also alter discrimination sensitivity. Before describing our study in detail, however, we will first review the background literature concerned with the influences of prior probabilities.

2. Background

2.1. Influence from amodal completion

From the perspective of signal detection theory, influences from priors may alter either discrimination bias,

---

* Corresponding authors. Fax: +1 310 206 5895 (Z. Liu).
  E-mail addresses: zhouy@ustc.edu.cn (Y. Zhou), zili@psych.ucla.edu (Z. Liu).
[1] These authors contributed equally to this work.

discrimination sensitivity, or both. From this perspective, alteration of bias is likely to occur at a later, decision stage; whereas alteration of sensitivity is likely to occur at an earlier, signal encoding stage. In this regard, alteration of sensitivity is stronger evidence of top–down influence. That is why, when only decision bias was found to have been altered, it was called "top-but-not-very-far-down" by Mussap and Levi (1995).

Mussap and Levi (1995) studied the possible top–down influence from amodal completion (i.e., perceptual completion behind occlusion) on 2D vernier acuity, i.e., the discrimination of two nearly collinear vertical bars with a small offset sideways (Westheimer, 1976; Wülfing, 1892). Subjects judged the direction of the offset (left, right, or zero). It was found that discrimination sensitivity d' was not influenced by amodal completion. Only the decision bias β was influenced when no feedback was provided, in which case subjects were more likely to say that the two vertical bars were collinear (i.e., zero offset).[2] This was therefore called by Mussap and Levi (1995) a "top-but-not-very-far-down" process. This result might be interpreted as follows. Although binocular disparity was used to manipulate the presence or absence of occlusion, the direction of the vernier offset was sideways in the fronto-parallel plane, not in depth. Furthermore, it has been shown that two parallel bars with an offset sideways do not strongly complete with each other (Kellman & Shipley, 1991). Given that depth difference is relatively more uncertain than difference in the fronto-parallel plane (Harris, McKee, & Watamaniuk, 1998), a stronger effect may be expected by manipulating occlusion for an offset that is in stereoscopic depth rather than in the fronto-parallel plane.

Stereoscopic discrimination in depth as a function of amodal completion was studied by Yin, Kellman, and Shipley (2000). A colored circular disk was presented in depth, either in front of or behind a gray rectangle. Behind this rectangle and disk, a third object, a colored oval shape completed itself amodally. Subjects discriminated whether the disk was in front of or behind the rectangle. It was found that discrimination sensitivity d' was reduced when the disk and the amodally completed oval shape shared the same color compared to when they had different colors or when the oval shape was absent. This indicates that when the oval and the disk shared the same color, the two shapes were grouped together to form a single surface in a single depth plane, therefore, making it more difficult to discriminate the disk's veridical depth relative to the rectangle's depth. There is, however, one aspect of the stimulus design in this study that could be improved. Since only

the contour of the disk provided its stereoscopic depth information relative to the rectangle, when the disk was behind the rectangle, the disk "pulled" with it its surrounding region of the rectangle also away from the observer (there was no hole in the rectangle). This is analogous, assuming that the observer's viewing direction is top–down, to a circular Frisbee disk sitting on top of a rectangular mosquito net. This "behind" condition was in contrast to the condition when the circular disk was in front of the rectangle, with the latter's surface being perfectly planar. This stimulus difference might be why subjects were less accurate when the disk was behind than in front of the rectangle. This might be also why subjects were a little biased against choosing the disk as being behind.

Liu, Jacobs, and Basri (1999) also studied stereoscopic depth discrimination under amodal completion. They assumed that the stronger two planar surfaces were grouped together via amodal completion, the harder it would be to discriminate stereoscopic depth differences between these two surfaces. They found that amodal completion with convex contours made stereodepth discrimination more difficult than when concave contours were presented. Although, they conducted a pilot experiment to verify their assumption, the number of subjects were small (i.e., three). Clearly, additional experiments are needed to verify this assumption.

2.2. Contextual effects in stereodiscrimination

So far, we have reviewed the literature concerned with the influence of amodal completion on vernier and stereoacuity discrimination. More generally speaking, stereoacuity discrimination appears to be influenced by stimulus configuration, which is often referred to as contextual effects. In what follows, we will review contextual effects in stereoacuity discrimination. The overarching theme of the review is that contextual effects can be understood as cue interactions (Landy, Maloney, Johnston, & Young, 1995), of which amodal completion is an example cue that can be in conflict with stereodepth information.

Mitchison and Westheimer (1984) presented two parallel vertical bars in depth and asked subjects to discriminate which was closer. They found that discrimination threshold was greatly elevated when the two bars were connected by two horizontal bars to form a square. This is possibly because the monocular linear perspective cues of the square indicated a square in the fronto-parallel plane. In fact, even when the two bars were connected by a single horizontal bar to form a letter 'H', discrimination threshold was elevated (McKee, 1983). Perhaps for a similar reason, Mitchison and Westheimer (1984) also found an elevated depth discrimination threshold between two columns of dots when they were flanked by additional columns of dots to form a slanted plane. These additional columns provided additional binocular disparity information to potentially aid depth discrimination between the middle two columns. However, because all the dots formed into a square grid

---

[2] $d' = Z(\text{hit-rate}) - Z(\text{false-alarm-rate})$, $\beta = \text{normpdf}(Z(\text{hit-rate}))/\text{normpdf}(Z(\text{false-alarm-rate}))$, where normpdf is the normal probability density function. An intuitive way to understand the bias β is that it is the ratio of the y-coordinates of the two normal distributions when the x-coordinate is at the decision criterion. For example, when the criterion is set where the two normal distributions intersect, the decision is bias free $\beta = 1$.

that acted as a monocular cue, dots may have been perceived to be in a fronto-parallel plane that accordingly caused discrimination to become much worse. Cue conflict can explain a more recent study of stereoacuity discrimination (Vreven, McKee, & Verghese, 2002). A stereoprobe was presented before a stereosurface that was either curved or planar and defined either by uniform luminance or by random dots. Subjects judged whether the probe was in front of or behind the surface. Since uniform luminance is a shape-from-shading cue indicating a planar as opposed to a curved surface, the following results appear completely consistent with a cue conflict interpretation regarding discrimination threshold. The threshold was greater for a curved surface defined by uniform luminance than defined by random-dots. The discrimination threshold was also greater for a curved than for a planar surface when both were defined by uniform luminance. Finally, the threshold was greater for a larger than for a smaller surface when both were curved and defined by uniform luminance, even though the latter contained additional stereoinformation from the contours.

Stimulus configuration not only impaired stereoacuity discrimination, but also was found to improve it. Discrimination threshold of a test line was lowest when a reference line was presented in the fixation plane (Badcock & Schor, 1985; Blakemore, 1970; McKee, Levi, & Bowne, 1990; Ogle, 1953). When a plane was added to a stereoacuity stimulus, discrimination appeared to be made in reference to this plane that improved discrimination (Glennerster & McKee, 1999). Indeed, additional reference lines were sufficient to improve discrimination by a factor of 10 (Kumar & Glaser, 1992). It appears that multiple, explicit, and redundant reference lines in stereo enhanced encoding accuracy of relative depth so that discrimination could be enhanced.

We now return to the hypothesis that stereoacuity discrimination under amodal completion can be considered as cue conflict (see also Bülthoff, Bülthoff, & Sinha, 1998; Gregory, 1968). That is to say, amodal completion suggests that planar surface regions that are meant to be amodally completed tend to be perceived as coplanar, whereas binocular disparity information suggests that these regions are not coplanar. To make this hypothesis precise, we will in the next section develop a simple Bayesian model, which predicts that amodal completion impairs stereoacuity discrimination.

3. A simple Bayesian model

We developed a simple Bayesian model as a way to understand why amodal completion might impair stereoacuity discrimination. In order to put the model in the proper context, we first describe the specifics of the experimental task.

3.1. The task

In the experiment, two vertical bars had a depth difference (Fig. 1). A subject discriminated whether the upper bar was closer or further away than the lower bar. In the occluded condition, a horizontal bar was positioned closest to the subject in such a way that the two vertical bars could amodally complete behind the occluder. In the non-occluded condition, the horizontal bar was positioned furthest away from the subject, so that there was a visible gap between the two vertical bars. Therefore, the vertical bars could not complete, amodally or modally. To avoid any accidental alignment, the horizontal bar, when being furthest away, was also slightly shortened in the vertical dimension. As a result, even a monocular stimulus image could not provide a cue to suggest that the horizontal bar was in front of the two vertical bars (Krech & Crutchfield, 1958). The hypothesis being tested was whether discrimination sensitivity would be reduced in the occluded condition, relative to the non-occluded condition.

3.2. Discrimination sensitivity d′

The model was set up in the manner of cue interaction, as discussed in Section 2.2. The formulation was Bayesian,


Fig. 1. Schematic of the two experimental conditions: occluded (left) and non-occluded (right). Left: the horizontal bar is positioned in front of the two vertical bars, which amodally complete behind the occlusion. Right: the horizontal bar is positioned behind the two vertical bars, which no longer perceptually complete as the gap between the two is clearly visible. Note that the height of the horizontal bar was reduced slightly in the non-occluded condition relative to the occluded condition, so that even in each monocular retinal image the two vertical bars could not amodally complete.

where the likelihood function was estimated from the disparity cues and depth positions of the vertical bars and the "prior probabilities" specified the context of the depth position of the horizontal bar (occluding vs. non-occluding). Our use of the term "prior probabilities" was different from the normal usage that refers to variations in degree of belief, prior to data. Our usage was closer to "model selection" (Grenander & Miller, 1994), i.e., selecting from several priors an appropriate prior in a given context. Here, the context was the depth position of the horizontal bar, which was assumed to be known as the prior.

Specifically, when the horizontal bar was in front, the prior assumed that the two vertical bars amodally completed. This implied that the two vertical bars were in the same depth plane, at a depth position μ. Let z_u and z_l represent the depth coordinate of the upper and lower vertical bars, respectively. Assuming that all distributions are Gaussian in shape, this prior can be expressed as

$$p(z_u|\mu) \propto \exp\left[-\frac{(z_u - \mu)^2}{2\sigma_p^2}\right],$$
$$p(z_l|\mu) \propto \exp\left[-\frac{(z_l - \mu)^2}{2\sigma_p^2}\right], \tag{1}$$

where σ_p is the standard deviation of Gaussian prior. We further assumed that the prior probability distribution of variable μ is

$$p(\mu) \propto \exp\left[-\frac{(\mu - \mu_o)^2}{2\sigma_\mu^2}\right], \tag{2}$$

where μ_o is a constant.

Assume that the relative depth difference between the two vertical bars, as provided by binocular disparity, is 2δ. This can be formulated into the likelihood in Bayesian term as (without loss of generality, $z = 0$ is assumed to be the midpoint in depth between the two vertical bars)

$$p(\delta|z_u) \propto \exp\left[-\frac{(z_u - \delta)^2}{2\sigma^2}\right],$$
$$p(\delta|z_l) \propto \exp\left[-\frac{(z_l + \delta)^2}{2\sigma^2}\right], \tag{3}$$

for each of the two vertical bars, respectively. The posterior probability distribution of the two vertical bars will therefore be:

$$p(z_u|\delta) = \int_{-\infty}^{+\infty} p(\mu)p(z_u|\mu)p(\delta|z_u)\,d\mu$$
$$\propto \exp\left\{-\frac{\left[z_u - \frac{\sigma^2\mu_o + (\sigma_p^2 + \sigma_\mu^2)\delta}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}\right]^2}{2\frac{\sigma^2(\sigma_p^2 + \sigma_\mu^2)}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}}\right\}, \tag{4}$$

$$p(z_l|\delta) = \int_{-\infty}^{+\infty} p(\mu)p(z_l|\mu)p(\delta|z_l)\,d\mu$$
$$\propto \exp\left\{-\frac{\left[z_l - \frac{\sigma^2\mu_o - (\sigma_p^2 + \sigma_\mu^2)\delta}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}\right]^2}{2\frac{\sigma^2(\sigma_p^2 + \sigma_\mu^2)}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}}\right\}, \tag{5}$$

respectively. Therefore, the posterior estimation of the vertical bars are two Gaussians of identical shape, and the MAP (maximum a posteriori) estimation of the positions of the two bars are at $\frac{\sigma^2\mu_o + (\sigma_p^2 + \sigma_\mu^2)\delta}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}$ and $\frac{\sigma^2\mu_o - (\sigma_p^2 + \sigma_\mu^2)\delta}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}$, respectively. According to the standard signal detection theory, discrimination sensitivity d' between these two Gaussians is

$$d' = \frac{2\delta}{\sigma}\frac{\sqrt{\sigma_p^2 + \sigma_\mu^2}}{\sqrt{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}} \leqslant \frac{2\delta}{\sigma}. \tag{6}$$

We are now going to demonstrate that, under the non-occluded condition, discrimination sensitivity d' is $2\delta/\sigma$, under the following assumptions. When the horizontal bar is furthest away, the visible gap between the two vertical bars indicates that no perceptual completion between the two vertical bars is possible. We assume that the prior in this case is non-informative, i.e., it is a Gaussian with $\sigma_p \to \infty$. From Eq. (6), discrimination sensitivity d' is therefore $\frac{2\delta}{\sigma}$. Therefore, we have shown that under the assumptions above, the prior expectation of amodal completion between the two vertical bars leads to impaired discrimination sensitivity d'.

Conceptually, the impairment of d' can be understood as follows. The amodal completion prior with a finite standard deviation σ_p scales down the distance between the two likelihood Gaussian functions by a factor of $\frac{\sigma_p^2 + \sigma_\mu^2}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2} \leqslant 1$. Therefore, even though the standard deviation of each of the two Gaussian functions is also scaled down, it is down scaled less (by a factor of $\sqrt{\frac{\sigma_p^2 + \sigma_\mu^2}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}}$). Consequently, the net result is that d' is scaled down by a factor of $\sqrt{\frac{\sigma_p^2 + \sigma_\mu^2}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}}$.

More generally, when the amodal completion prior is assumed to be locally distributed (as opposed to be uniformly distributed as $\sigma_p \to \infty$), so long as the net effect is the distance between the two likelihood distributions scales down more than their standard deviations do, discrimination is expected to decrease. We have used Gaussian functions as a concrete example. Another reason to use Gaussians is that d' will otherwise be undefined.

3.3. Decision criterion and the consequent bias β

We now look at the decision criterion that determines any possible bias away from 50 to 50 in choosing the upper or lower vertical bar as being closer to the observer. When this decision criterion is set in a relative scale proportional to the distance between the two posterior probability distributions of the, respectively, two vertical bars, then the bias

is not going to change from the non-occluded to the occluded conditions. This is straightforward to understand because the introduction of a localized prior probability for the occluded condition does not break the symmetry between the two likelihood probability distributions for the two vertical bars. In other words, the two posterior probability distributions remain translationally symmetric for the non-occluded and occluded conditions alike. Formally, this can be demonstrated as follows.

For simplicity, we consider the non-occluded condition first. Recall that $z = 0$ is the midpoint between the two posterior probability distributions, let us assume that the coordinate of the decision criterion is at $z = z_{no} = \lambda\delta$, where $-1 \leqslant \lambda \leqslant 1$. Then by definition, the bias

$$\beta_{no} = \frac{\exp[-(z_{no} + \delta)^2/2\sigma^2]}{\exp[-(z_{no} - \delta)^2/2\sigma^2]} = \exp\left(-\frac{2\lambda\delta^2}{\sigma^2}\right). \quad (7)$$

In the occluded condition, the midpoint between the two posterior probability distributions is now $z = \frac{\sigma^2\mu_o}{\sigma^2+\sigma_p^2+\sigma_\mu^2}$. Then the decision criterion should be at

$$z = z_o = \frac{\sigma^2\mu_o}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2} + \lambda\frac{\sigma_p^2 + \sigma_\mu^2}{\sigma^2 + \sigma_p^2 + \sigma_\mu^2}\delta. \quad (8)$$

The bias is therefore

$$\beta_o = \frac{\exp\left[-\frac{\left(z_o - \frac{\sigma^2\mu_o - (\sigma_p^2+\sigma_\mu^2)\delta}{\sigma^2+\sigma_p^2+\sigma_\mu^2}\right)^2}{2\sigma^2\frac{\sigma_p^2+\sigma_\mu^2}{\sigma^2+\sigma_p^2+\sigma_\mu^2}}\right]}{\exp\left[-\frac{\left(z_o - \frac{\sigma^2\mu_o + (\sigma_p^2+\sigma_\mu^2)\delta}{\sigma^2+\sigma_p^2+\sigma_\mu^2}\right)^2}{2\sigma^2\frac{\sigma_p^2+\sigma_\mu^2}{\sigma^2+\sigma_p^2+\sigma_\mu^2}}\right]} = \exp\left(-\frac{2\lambda\delta^2}{\sigma^2}\right)$$

$$= \beta_{no}. \quad (9)$$

We have demonstrated therefore that the bias will not change from the non-occluded to occluded conditions, no matter where the prior probability distribution of the amodal plane ($z = \mu_o$) is centered, so long as this distribution is localized (i.e., σ_μ is finite) when the decision criterion is set on a relative basis.

When the decision criterion is set on an absolute depth value regardless of the occlusion condition, then the bias is different between the occluded and non-occluded conditions. This is because, in the occluded condition, the amodal plane ($z = \mu_o$) is at an arbitrary position μ_o whereas, in the non-occluded condition, the midpoint between the two vertical bars is at $z = 0$. However, we are going to argue that $\mu_o \neq 0$ is not a highly likely situation. Recall that $z = 0$ is the midpoint between the two vertical bars in the non-occluded condition, and $\frac{\sigma^2\mu_o}{\sigma^2+\sigma_p^2+\sigma_\mu^2}$ is the midpoint in the occluded condition. Based on symmetry argument, it is equally likely for $\mu_o > 0$ and $\mu_o < 0$. Therefore, the expected value of μ_o is $\mu_o = 0$. So long as $\mu_o = 0$, even if the decision criterion is set at an arbitrary constant but unchanged from the non-occluded to the occluded conditions, then the bias will

remain unchanged. This again follows the symmetry argument. It can also be verified by, in Eqs. (8) and (9), letting $\mu_o = 0$ and $z_{no} = z_o =$ an arbitrary constant.

As will be shown in the next section, human subjects had a bias of more often choosing the lower vertical bar as being closer, consistent with the "looking from above" hypothesis (Mamassian & Landy, 1998). This suggests that when the two vertical bars are amodally completed into a single plane, this plane is slanted away from the observer toward the ground plane. Moreover, this bias of more often choosing the lower bar as closer was robust enough to be unchanged in the non-occluded condition. Nevertheless, this bias was not accounted for in the simple model presented in this section, since both vertical bars were assumed to be parallel to the fronto-parallel plane.

In summary, a simple Bayesian model has been presented in this section that provided an explanation as to why the occluded condition may show a reduction in discrimination sensitivity d', whereas decision bias β is not expected to change between the non-occluded and occluded conditions.

4. Experiment

4.1. Stimuli and procedure

Before the experiment proper, there was a practice session. A fixation point 2.22×2.19 min of arc in size was presented in the center of the display. Two textured rectangles ($0.81° \times 0.80°$ each) were then presented, one above and one below the fixation point. One rectangle was in front of the zero-disparity plane, the other behind. The absolute value of the disparity of each rectangle was 4.43 min of arc. Subjects decided which rectangle was in front of the other. Feedback was provided by a computer beep after an incorrect response. Accuracy was computed every 20 trials. The practice session stopped when a subject reached at least 95% accuracy.

In each experimental trial, two vertical bars and one horizontal bar were displayed at the center of the screen for unlimited time (Fig. 2), while the fixation was not shown. A subject decided if the upper vertical bar was in front of or behind the lower vertical bar. The stimulus disappeared after the subjects' response via a key press, without feedback. The next trial started in 1 s.

Each of the two vertical bars was $0.41° \times 1.00°$ in size. One was in front of, and the other behind the zero-disparity depth plane, by the same amount of disparity (e.g., $+0.236$ and -0.236 min of arc). Five levels of disparity were used for psychometric measurement: ±0.236, ±0.354, ±0.709, ±0.945, and ±1.182 min of arc. The horizontal bar was positioned in depth either closest to the subject and therefore served as an occluder (disparity $+0.15°$) or furthest away from the subject (disparity $-0.15°$), allowing the gap between the vertical bars to be clearly visible. When the horizontal bar was an occluder, its size was $1.623° \times 0.402°$, making it large enough to close the gap

Fig. 2. Sample experimental stimuli, a subject decided whether the upper vertical bar was in front of or behind the lower vertical bar. Left: the horizontal bar is in front of the two vertical bars. Middle: the horizontal bar is behind the two vertical bars. In this condition the height of the horizontal bar was reduced slightly so that a gap was visible between the horizontal bar and each of the two vertical bars to prevent amodal completion. Right: the texture images used to wrap around the vertical or horizontal bars, respectively, (the two vertical texture images were identical).

between the two vertical bars. In this condition the two vertical bars presumably amodally completed behind the occluder. When the horizontal bar was presented behind the vertical bars, its height was reduced to 0.359° from 0.402° so there was a visible gap between the horizontal bar and each of the two vertical bars. The texture wrapping the horizontal bar was scaled accordingly, so that disparity information provided by the horizontal bar was approximately unchanged (smaller area but denser texture). We reduced the height to ensure that there would be no accidental alignment of the boundaries of the vertical and the horizontal bars, which otherwise would have provided monocular cues to suggest a horizontal occluder (Krech & Crutchfield, 1958; Nakayama & Shimojo, 1992).

The experiment had four blocks, with 80 trials each. Two blocks had the horizontal bar in front of the vertical bars, serving as the occluder (the occluded condition); the other two blocks had the horizontal bar behind, serving as the control (the non-occluded condition). At the beginning of each block, subjects were informed about the relative depth between the horizontal and the vertical bars. Half of the subjects ran the four blocks in the order: occluded, non-occluded, non-occluded, and occluded. The other half of the subjects ran in the counter-balanced order: non-occluded, occluded, occluded, non-occluded. The experiment lasted for about 40 min per subject.

4.2. Apparatus

Subjects wore red and green stereofilters to view the stimulus from a distance of 3 m. The average luminance of the stimulus to each eye was 1.67 cd/m^2 (measured with a UDT 161 photometer). The stimuli were displayed on a Sony G220 monitor with a resolution of 1600×1200 pixels driven by a Matrox G450 graphics card. The vertical

refresh rate of the monitor was 75 Hz. The monitor was calibrated with a TES-1330 A digital photometer.

4.3. Subjects

Thirty-six students from the City of Hefei, China, who were naïve to the purpose of the experiment and with normal or corrected-to-normal vision participated. The disparity threshold of each participant was smaller than 40 s of arc as measured by the Titmus® stereotest.

4.4. Results

A three-way ANOVA was performed to analyze discrimination sensitivity d' as a function of the following three factors: amodal completion (occluded vs. non-occluded), depth difference (the five levels of disparity difference between a vertical bar and the zero-disparity plane), and the two experimental sequences. The main effect of amodal completion (occluded vs. non-occluded) was statistically significant ($d' = 1.65$ vs. 1.88) ($F(1,34) = 4.56$, $p = 0.04$).[3] This suggests that stereodiscrimination was worse with amodal completion than without (Fig. 3). As expected, the main effect of depth difference was also statistically significant ($F(4,136) = 69.08$, $p < 0.001$). The main effect of experimental sequence, however, was not significant ($F(1,34) < 1$).

Although Fig. 3 shows, as expected, monotonic functions of discrimination sensitivity d' as a function of binocular disparity, this function is clearly nonlinear. This nonlinearity contrasts with the linear function between d' and disparity as predicted by the simple Bayesian model shown in Eq. (6). This demonstrates one limitation of the simple model, which did not consider any possible effect of ceiling performance. Fig. 3 shows another limitation of the simple model, as follows. Again from Eq. (6), one can find the following:[4]

$$\sqrt{\sigma_{\mathrm{p}}^2 + \sigma_{\mu}^2} = 2\delta \frac{d'_{\mathrm{o}}/d'_{\mathrm{no}}}{\sqrt{(d'_{\mathrm{no}})^2 - (d'_{\mathrm{o}})^2}}, \tag{10}$$

where d'_{o} and d'_{no} are discrimination sensitivity for the occluded and non-occluded conditions, respectively. One would expect that the variance of the priors σ_{p}^2 and σ_{μ}^2 to be constant. However, numerical evaluations of the Eq. (10) from experimental results in Fig. 3 gave rise to values ranging from 1.88 to 2.08 min of arc. So the two vari-

---

[3] The bitmap texture images, as shown in Fig. 2 right, were used to create textures for the stimuli. Each texture image, however, had a thin dark (though not black) outline, due to the authors' error. This created an effective tiny "gap" between each vertical bar and the horizontal bar, even in the occluded condition, since the dark outline was similar in luminance to the black background. This may have reduced the effect. In this sense, our effect was a conservative measure. The dark outline was three pixels in thickness, amounting to 0.7 min of arc.

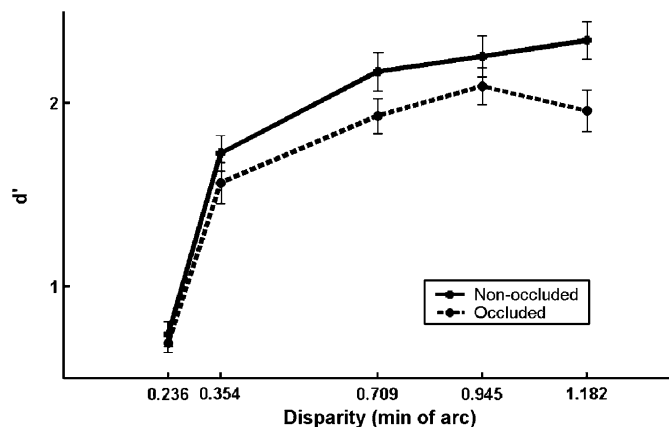[4] We thank one anonymous reviewer for pointing this out and deriving it.

Fig. 3. Discrimination sensitivity d' as a function of the absolute disparity between the vertical bars and the zero-disparity plane for both the occluded and non-occluded conditions. When the two vertical bars amodally completed behind the occluder (occluded condition), d' was statistically significantly smaller than when the two vertical bars could not perceptually complete. Errors bars represent standard error of the mean.

ances could not be constant. We also looked at subjects' response bias. Subjects indeed showed a bias by selecting the lower vertical bar as being closer more than 50% of the time (occluded: $\beta = 1.12$, non-occluded: $\beta = 1.15$). This is consistent with the "looking-from-above" hypothesis proposed by Mamassian and Landy (1998). However, this effect was not statistically significant. We similarly conducted a three-way ANOVA except on β values rather than d' values. No effect was statistically significant (the two-way interaction between depth difference levels and presentation sequence: $F(4,136) = 1.62$, $p = 0.17$; the main effect of presentation sequence: $F(1,34) = 2.37$, $p = 0.13$; $F < 1$ for all the remaining effects) (Fig. 4). To ensure that the bias analysis was solid, we repeated the three-way ANOVA on $\log(\beta)$ instead of β, (given that $\log(\beta)$ was more likely to have a normal distribution), and obtained similar results whereby no effect reached statistical significance, main


Fig. 4. Discrimination bias β across different disparities between a vertical bar and the zero-disparity plane for both the occluded and non-occluded conditions. No reliable difference was found. A β value greater than one (not statistically reliable) suggests that a subject had a bias favoring the lower vertical bar as being closer to the subject.

effect of disparity level: $F(4,136) = 1.23$, $p = 0.30$; main effect of block sequence: $F(1,34) = 2.36$, $p = 0.13$; two-way interaction between disparity level and block sequence: $F(4,136) = 1.63$, $p = 0.17$; all the remaining F values were smaller than 1 ($F < 1$).

4.5. Discussion

Since no fixation point was displayed in experimental trials and viewing time was unlimited (subjects took 5–10 s on average per trial), one might argue that vergence eye movements may drift toward the horizontal bar.[5] It is true indeed that under the experimental conditions vergence eye movements, and hence the location of the zero disparity fixation plane, could not be controlled. However, even if subjects' vergence eye movements drifted toward the horizontal bar, this could not in itself explain the experimental result. This is because the absolute disparity difference between the horizontal bar and the two vertical bars was the same (0.15°), regardless of whether the horizontal bar was in front of (+0.15°) or behind (−0.15°) the vertical bars.

Another possibility might be that subjects fixated on the closest stimulus, so in the occluded condition they fixated on the horizontal bar, whereas in the non-occluded condition they fixed on the two vertical bars. If this were the case, by the virtue of the two vertical bars being closer to the zero disparity plane in the non-occluded condition, one may expect that discrimination would be better for the non-occluded than for the occluded condition (Glennerster & McKee, 1999; Glennerster & McKee, 2004; Glennerster, McKee, & Birch, 2002). In this regard, one may further expect that discrimination would be better when the horizontal bar is furthest away than when it is closest, independent of amodal completion. However, data in the literature do not seem to support this possibility of fixating at the closest stimuli. In Liu et al. (1999), the horizontal bar was either closest or furthest away from subjects, while its height was 70% of the gap size. Hence, no perceptual completion was possible between the two equivalent "vertical bars" regardless of the position of the horizontal bar. No performance difference was found between the "closest" and "furthest away" conditions. This indicates that when free viewing without a fixation, which was the case in both Liu et al. (1999) and the current experiment, subjects were unlikely to fixate only on the closest stimuli. It is plausible that subjects concentrated on the two vertical bars in order to discriminate their relative depth difference in both conditions.

5. Conclusions

We found that stereoscopic depth discrimination was impaired when monocular cues indicated that two planar surfaces were coplanar without any depth difference. More specifically, in the occluded condition, the two vertical bars

---

[5] We thank one anonymous reviewer for raising this point.

perceptually completed behind a horizontal occluding bar, making relative depth discrimination between the two vertical bars difficult. In comparison, in the non-occluded condition, the horizontal bar was pushed furthest away such that the two vertical bars could not perceptually complete since there was a visible gap between them. In this non-occluded condition, relative depth discrimination became relatively easier.

We have proposed a simple Bayesian model to interpret the findings above. The model's priors were different in the two conditions, while the likelihood function of stimulus information was unchanged. In the occluded condition, the Bayesian prior assumed that the two vertical bars had zero relative depth difference because of the amodal completion. In contrast, in the non-occluded condition, because the two vertical bars could not complete with each other, the Bayesian prior made a non-committal assumption that each of the two vertical bars was equally likely to be anywhere in depth. In other words, the prior made no preference and let the likelihood function determine the posterior probability distribution. With these assumptions, plus a generic assumption that the probability distribution was Gaussian in shape, we demonstrated that the simple Bayesian model predicted that the sensitivity of relative depth discrimination was impaired under the occluded condition, while the bias was little changed.

References

Badcock, D. R., & Schor, C. M. (1985). Depth-increment detection function for individual spatial channels. *Journal of Optical Society of America A, 2*(7), 1211–1216.

Blakemore, C. (1970). The range and scope of binocular depth discrimination in man. *Journal of Physiology (London), 211*, 599–622.

Bülthoff, I., Bülthoff, H. H., & Sinha, P. (1998). Top–down influences on stereoscopic depth-perception. *Nature Neuroscience, 1*, 254–257.

Feldman, J. (2000). Bias toward regular form in mental shape spaces. *Journal of Experimental Psychology—Human Perception and Performance, 26*(1), 152–165.

Glennerster, A., & McKee, S. P. (1999). Bias and sensitivity of stereo judgements in the presence of a slanted reference plane. *Vision Research, 39*(18), 3057–3069.

Glennerster, A., & McKee, S. P. (2004). Sensitivity to depth relief on slanted surfaces. *Journal of Vision, 4*(5), 378–387.

Glennerster, A., McKee, S. P., & Birch, M. D. (2002). Evidence for surface-based processing of binocular disparity. *Current Biology, 12*(10), 825–828.

Gregory, R. L. (1968). Visual illusions. *Scienticfic American*, 66–76.

Grenander, U., & Miller, M. I. (1994). Representations of knowledge in complex systems. *Journal of the Royal Statistical Society Series B, 56*, 549–603.

Harris, J. M., McKee, S. P., & Watamaniuk, S. N. (1998). Visual search for motion-in-depth: Stereomotion does not 'pop out' from disparity noise. *Nature Neuroscience, 1*(2), 165–168.

Helmholtz, H. v. (1924). *A treaties on physiological optics* (1st ed.). New York: Optical Society of America.

Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology, 23*, 141–221.

Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology, 13*, 1–9.

Knill, D. C., & Kersten, D. (1991). Ideal perceptual observers for computation, psychophysics and neural networks. In R. Watt (Ed.), *Pattern recognition by man and machine*. London: MacMillan Press.

Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. Cambridge University Press.

Krech, D., & Crutchfield, R. S. (1958). *Elements of psychology*. New York: Knopf.

Kumar, T., & Glaser, D. A. (1992). Depth discrimination of a line is improved by adding other nearby lines. *Vision Research, 32*(9), 1667–1676.

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.

Liu, Z., Jacobs, D. W., & Basri, R. (1999). The role of convexity in perceptual completion: beyond good continuation. *Vision Research, 39*, 4244–4257.

Mamassian, P., & Landy, M. S. (1998). Observer biases in the 3D interpretation of line drawings. *Vision Research, 38*(18), 2817–2832.

Mamassian, P., & Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research, 41*(20), 2653–2668.

McKee, S. P. (1983). The spatial requirements for fine stereoacuity. *Vision Research, 23*(2), 191–198.

McKee, S. P., Levi, D. M., & Bowne, S. F. (1990). The imprecision of stereopsis. *Vision Research, 30*(11), 1763–1779.

Mitchison, G. J., & Westheimer, A. (1984). The perception of depth in simple figures. *Vision Research, 24*, 1063–1073.

Mussap, A. J., & Levi, D. M. (1995). Amodal completion and vernier acuity: evidence of 'top-but-not-very-far-down' processes? *Perception, 24*, 1021–1048.

Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science, 257*(5075), 1357–1363.

Ogle, K. N. (1953). Precision and validity of stereoscopic depth perception from double images. *Journal of the Optical Society of America, 43*, 906–913.

Vreven, D., McKee, S. P., & Verghese, P. (2002). Contour completion through depth interferes with stereoacutiy. *Vision Research, 42*, 2153–2162.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience, 5*, 598–604.

Westheimer, G. (1976). Diffraction theory and visual hyperacuity. *American Journal of Optometry and Physiological Optics, 53*(7), 362–364.

Wülfing, E. A. (1892). Ueber den kleinsten Gesichtswinkel. *Zeitschrift für Biologie Bd., 19*, 199–202.

Yin, C., Kellman, P. J., & Shipley, T. F. (2000). Surface integration influences depth discrimination. *Vision Research, 40*, 1969–1978.