Vision Research 49 (2009) 134–139

Contents lists available at ScienceDirect

Vision Research

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

Testing a counter-intuitive prediction of optimal cue combination

Chris M.P. Muller *, Eli Brenner, Jeroen B.J. Smeets

Research Institute MOVE, Faculty of Human Movement Sciences, VU University Amsterdam, Van der Boechorststraat 9, 1081 BT Amsterdam, The Netherlands

ARTICLE INFO

Article history:
Received 3 July 2008
Received in revised form 29 September 2008

Keywords:
Optimal cue combination
Slant
Binocular
Monocular

ABSTRACT

Weighted averaging is said to be optimal when the weights assigned to the cues minimize the variance of the final estimate. Since the variance of this optimal percept only depends on the variances of the individual cues, irrespective of their values, judgments about a cue conflict stimulus should have the same variance as ones about a cue consistent stimulus. We tested this counter-intuitive prediction with a slant matching experiment using monocular and binocular slant cues. We found that the slant was indeed matched with about the same variance when the cues indicated slants that differed by 15° as when they indicated the same slant.

1. Introduction

When several sources of information (cues) about a certain aspect of a scene are present, the human visual system makes separate estimates of that aspect, based on the individual cues, and harmonizes the different estimates through weighted averaging (Landy, Maloney, Johnston, & Young, 1995; Young, Landy, & Maloney, 1993). This harmonizing takes care of the slight differences between the estimates that arise from processing noise. If a monocular cue and a binocular cue both provide information about the slant of a surface, estimating its value to be m and b respectively, and are assigned weights w and $(1-w)$ respectively, then according to the weighted averaging model the perceived slant (S) can be represented by:

$$S = wm + (1-w)b \tag{1}$$

If the errors in m and b are independent, as they are likely to be for slant specified by perspective and binocular disparity, the variance of S can be given by:

$$\sigma_s^2 = w^2\sigma_m^2 + (1-w)^2\sigma_b^2 \tag{2}$$

where σ_m^2 is the variance in the monocular cue and σ_b^2 is the variance in the binocular cue. Weighted averaging is considered to be optimal when the variance of $S(\sigma_b^2$ in Eq. (2)) is minimal. To find the minimal variance of S one determines the value of w for which the derivative of Eq. (2) with respect to w is zero. This gives a value for w expressed as a function of the individual cue variances:

$$W = \frac{\sigma_b^2}{\sigma_m^2 + \sigma_b^2} \tag{3}$$

By substituting w from (3) into (2) we find the lowest possible variance of S expressed in terms of the individual cue variances (see Ernst & Bulthoff, 2004; Hillis, Watt, Landy, & Banks, 2004).

$$\sigma_{S\_optimal}^2 = \frac{\sigma_m^2\sigma_b^2}{\sigma_m^2 + \sigma_b^2} \tag{4}$$

The counter-intuitive aspect of Eq. (4) is that $\sigma_{S\_optimal}^2$ is independent of the slant estimates m and b, so that the precision of the combined percept will be the same for percepts based on consistent cues ($m = b$) as for percepts based on conflicting cues ($m \neq b$) Of course, Eq. (2) and its consequences are only valid if the variability is all in the individual cues. If the variability mainly arises after the cues are combined, the precision of that later stage may (or may not) be influenced by the visible consequences of the conflict (Muller, Brenner, & Smeets, 2007). However, it has previously been shown that the variability under similar conditions mainly depends on the variability in the single cues (Knill & Saunders, 2003; Muller et al., 2007). Similarly, if there is any variability in the weights (Brenner, Granzier, & Smeets, 2007), we can expect a larger variability in the cue conflict conditions. The reason for this is that the variability in the perceived slant as a consequence of variability in the weights is proportional to the difference between the slant judged from the separate cues. If the judged slant is the same for both cues then variability in the weights makes no difference. Intuitively, one would therefore expect the precision to be higher if there is a close agreement between the cues.

In order to use Eq. (4) we need to have estimates of the reliability of the individual cues (σ_m^2 & σ_b^2). It is difficult to find conditions that are equivalent for isolated cues and combined cues, because

* Corresponding author.
E-mail address: C.Muller@fbw.vu.nl (C.M.P. Muller).

0042-6989/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2008.10.006

removing one cue usually requires manipulations that could also influence the precision of the other cue (Muller et al., 2007). We therefore use a method that does not depend on finding such conditions. In order to compare equivalent cue conflict and cue consistent stimuli, we chose stimuli in which both the monocular and the binocular cue indicated a slant of either 50° or 65° (top tilted backwards). There are four possible combinations of these angles, two of which consist of a consistent cue pair: (50,50), (65,65), and two of which consist of a conflicting cue pair: (50,65) and (65,50). We can rewrite Eq. (4) as:

$$\frac{1}{\sigma^2_{S\_optimal}} = \frac{1}{\sigma^2_m} + \frac{1}{\sigma^2_b} \tag{5}$$

So for the four above-mentioned combinations of slants, and assuming optimal cue combination, we can write:

$$\frac{1}{\sigma^2_{50,50}} = \frac{1}{\sigma^2_{m50}} + \frac{1}{\sigma^2_{b50}} \tag{6a}$$

$$\frac{1}{\sigma^2_{65,65}} = \frac{1}{\sigma^2_{m65}} + \frac{1}{\sigma^2_{b65}} \tag{6b}$$

$$\frac{1}{\sigma^2_{50,65}} = \frac{1}{\sigma^2_{m50}} + \frac{1}{\sigma^2_{b65}} \tag{6c}$$

$$\frac{1}{\sigma^2_{65,50}} = \frac{1}{\sigma^2_{m65}} + \frac{1}{\sigma^2_{b50}} \tag{6d}$$

The variances in the left terms of the equations are variances that we will measure. The variances in the right terms are those of the two underlying cues. The sum of the right parts of the two equations for the cue consistent stimuli ((6a) and (6b)) is equal to the sum of the right parts of the two equations for the cue conflict stimuli ((6c) and (6d)). Therefore, the sum of the left parts must also be equal:

$$\frac{1}{\sigma^2_{50,65}} + \frac{1}{\sigma^2_{65,50}} = \frac{1}{\sigma^2_{50,50}} + \frac{1}{\sigma^2_{65,65}} \tag{7}$$

This can be rewritten in the form of Eq. (4), which is easier to understand in terms of variances.

$$\frac{\sigma^2_{50,65}\sigma^2_{65,50}}{\sigma^2_{50,65} + \sigma^2_{65,50}} = \frac{\sigma^2_{50,50}\sigma^2_{65,65}}{\sigma^2_{50,50} + \sigma^2_{65,65}} \tag{8}$$

This is the formalization of the counter-intuitive prediction that we will test. We will refer to the left hand term of Eq. (8) as the combined variance of the conflict trials and to the right hand term of Eq. (8) as the combined variance of the consistent trials. Intuitively, one would expect the combined variance for matching conflict stimuli to be larger than that for matching consistent stimuli, as formalized in Eq. (9).

$$\frac{\sigma^2_{50,65}\sigma^2_{65,50}}{\sigma^2_{50,65} + \sigma^2_{65,50}} > \frac{\sigma^2_{50,50}\sigma^2_{65,65}}{\sigma^2_{50,50} + \sigma^2_{65,65}} \tag{9}$$

The aim of the present study is to test whether the intuitive prediction (Eq. (9)) is valid, in which case we can reject the counter-intuitive prediction (Eq. (8)). To do so, we had observers match the slant of a probe surface to the slant of a reference surface, and analyzed the variance in the slant settings.

2. Methods

2.1. Setup

Our setup consisted of an Apple G5 computer that generated the images and processed the responses, a 57 cm (diagonal) Sony Trinitron monitor (resolution 1096 × 686 pixels), and Crystal Eyes stereo shutter spectacles. The images were generated at a refresh rate of 160 Hz (80 Hz per eye), using only the red gun because the spectacles work best for red images. Observers sat 1 m from the screen, so that the screen was approximately 27° × 17°.

2.2. Stimuli

There were two 8 cm diameter slanted virtual surfaces, that were only visible because four asynchronously refreshing limited lifetime (100 ms) discs were projected onto them. The slant of one of the surfaces (reference) was determined by the computer. The slant of the other surface (probe) could be manipulated by the observer by moving the mouse. The simulated discs had a diameter of 1.2 cm. The centers of the surfaces were separated from each other by 12 cm. Fig. 1 shows a stereogram of a single frame of the stimulus in a situation where the reference and the probe look identical. The surfaces' slants were specified by monocular cues (the elliptical shape of the discs and their distribution over the surfaces) and by binocular disparity.

The short lifetime of the discs prevented observers from following changes in individual discs' positions or shapes while they manipulated the probe surface's slant, which may have otherwise provided additional information about slant. In order to independently control the slant that was indicated by each of the two cues, we determined the position and shape of the images as seen from a point between the eyes (for the monocular slant) and rendered images that would give the same shape from that viewpoint but with disparities that are in accordance with the binocular slant (Hillis et al., 2004). Our conflict stimuli had quite a large slant conflict (15°), but not so large that observers would switch between two percepts (van Ee, van Dam, & Erkelens, 2002). For the probe surface, the two cues were always consistent. For the reference surface, the cues could either be consistent or in conflict.

2.3. Task

On each trial observers were presented with the two slanted surfaces at the same time, and were asked to match the slant of the (cue consistent) probe surface on the right to the slant of the reference surface on the left. They modified the probe surface's slant by moving the computer mouse. They indicated that they were satisfied with their setting by clicking the computer mouse, which also started the next trial. On each trial we recorded the slant that observers set as well as the time they took to make that setting.

2.4. Observers and conditions

Twelve observers (two authors; 10 naive) took part in a single experimental session. All observers had normal stereo acuity. The study was approved by the ethics committee of the Faculty of Human Movement Sciences.

We presented six different reference surface conditions (see Fig. 2), the four conditions described in Section 1 (50,50; 65,65; 50,65; 65,50) and two intermediate cue consistent conditions (55,55; 60,60). The purpose of the latter two will be explained in Section 2.5.

In each session the six different conditions were repeated 40 times, in random order, giving a total of 240 trials. The session was preceded by five practice trials that were not recorded.

2.5. Analysis

We analyzed the variance in the set slants in each of the six conditions for each observer. In order to remove trials in which observers pressed the button accidentally, we defined an outlier as a setting that exceeded the observers average slant for that

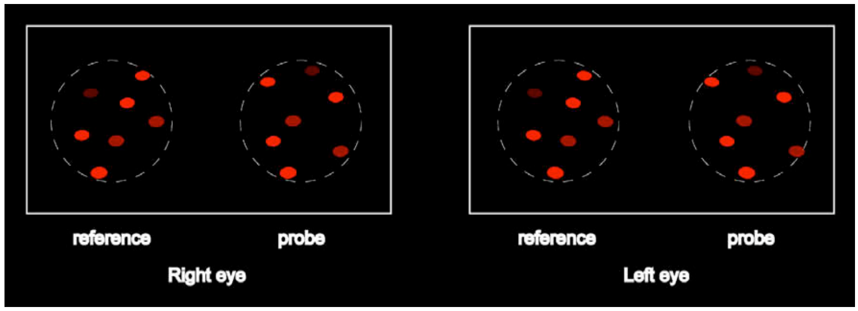

Fig. 1. Stereogram of one stimulus frame. The two images on the right are the left eye images of the reference and probe surfaces. The images on the left are the right eye images of the reference and probe surfaces. In the actual stimulus, only four discs were visible at a given time on each surface. The extra three discs on each surface, shown in varying luminance, illustrate the fact that due to the discs' limited life, the subjective percept was of there being about seven discs at a given time. The dotted lines indicate the boundaries within which the discs appeared; the solid lines help cross fusing this image, none of the lines were present in the actual display.

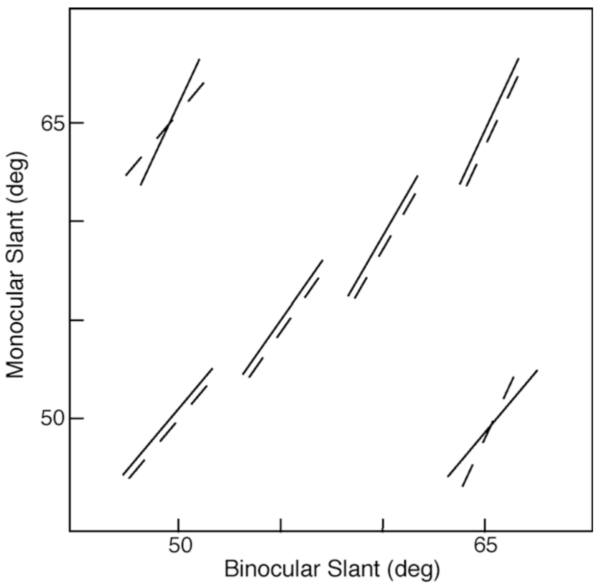

Fig. 2. Graphical representation of the reference slants in the six conditions. Solid lines indicate the monocular slant and dashed lines indicate the binocular slant.

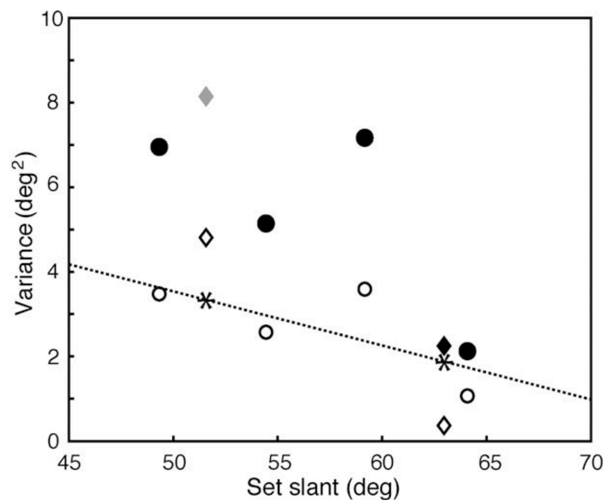

Fig. 3. Illustration of how we estimated the variances in perceiving the slants of the six reference surfaces (open symbols) from the variances in the settings (filled symbols). Data for one observer. For trials with consistent cues, each reference variance (open circle) is simply half the measured variance (corresponding filled circle). For trials with conflicting cues (diamonds) we interpolated values at the set slant from a linear fit to the open circles (stars). We subtracted these values from the variances measured in the cue conflict trials (filled diamonds) to obtain estimates of the variance in perceiving the reference surface slant (open diamonds). The gray and black diamonds represent the conflict conditions in which the binocular and monocular cue indicate a 50° slant, respectively. Their positions show that this observer gave most weight to the binocular cue.

condition by more than 3 standard deviations (as calculated without that setting). This resulted in us excluding 3% of the data from further analysis. The measured variances in the cue conflict and cue consistent conditions cannot be compared directly, because the reliability of each cue depends on its slant, and these dependencies differ for the two cues (Banks, Hooge, & Backus, 2001; Knill, 1998). We therefore evaluate the relationship shown in Eq. (8).

In order to obtain the variances that are used in Eq. (8) we have to consider that in our experiment the measured variance of the set slants is the sum of both the variance in perceiving the reference surface slant and the variance in perceiving the probe surface slant. As Eqs. (8) and (9) only deal with the variance in perceiving the reference slant, we had to separate these two variances. We did so by first estimating the contribution of perceiving the cue consistent probe surface. We then subtracted this estimate from the total variance to obtain an estimate of the variance in perceiving the reference surface. Fig. 3 illustrates the steps that we followed to do so (as described below).

In the cases in which both the reference surface and the probe surface were cue consistent stimuli, we assume that the variance in perceived slant is equal for both the reference surface and the probe surface. So, for the cue consistent conditions, halving the measured variances (filled circles in Fig. 3) gives us the variance we are interested in; the variance in perceiving the reference surface slant (open circles in Fig. 3).

For the trials with cue conflicts the variances cannot be assumed to be the same for both surfaces. We therefore first estimated the contribution of perceiving the (cue consistent) probe surface slant, and then subtracted that from the total to be left with only the contribution of the reference slant. The two additional cue consistent stimuli mentioned earlier, with cue consistent reference slants of 55° and 60°, were included to help us do this. We assume that the relationship between the variance and the slant is linear within the slant range of our experiment; between 50° and 65°. This allows us to fit a line to the four (halved) variances (open circles in Fig. 3) and used this to interpolate the variance in perceiving the probe surface slant for the set slants in the cue conflict trials (stars in Fig. 3). These variances were subtracted from the measured cue conflict variances (filled diamonds in Fig. 3) to obtain our estimate of the variance in perceiving the reference surface slant on conflict trials (open diamonds in Fig. 3). Note that if the matching itself introduces variability, such variability is attributed equally to the two surfaces in the cue consistent condition. Any increase in such variability when the reference surface contains a cue conflict will be attributed to perceiving the reference surface by our calculation.

Following these steps we obtained an estimate of the variance in perceiving the reference surface slant for each condition and observer. The estimated variances were used to calculate the combined variances as shown in Eq. (8). We report the average values that we found for the two terms of Eq. (8), and use a paired *t*-test to evaluate whether there is a consistent difference between the terms across observers.

Observers might trade response precision for response speed. We therefore also determined the median time that each observer took to match the reference slant for cue consistent and cue conflict trials. We used a paired *t*-test to compare the median response times for the two kinds of trials across observers.

Finally, to ascertain that we could expect to see a change in the combined variance if observers stopped combining the cues when they were in conflict, we calculated what the combined variance of each observer would have been if they had used only the most reliable cue, rather than a combination. This step relies on the assumption that observers were combining optimally when the cues were not in conflict.

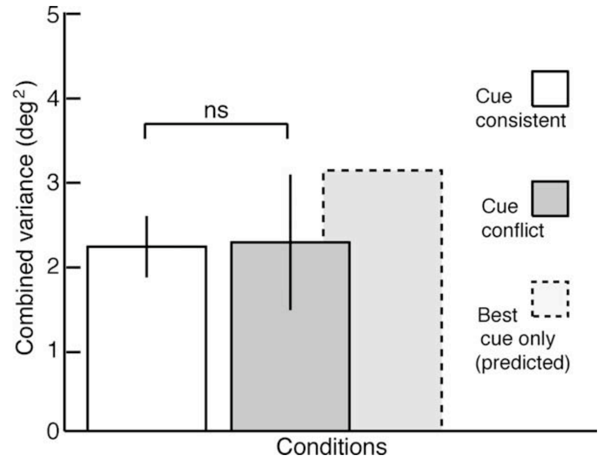

Fig. 5. Combined variances in the cue consistent (white bars) and cue conflict (dark gray bars) conditions, as well as the predicted combined variances (light gray, dashed bars) in the cue conflict conditions if observers had only used their best cue. Averages across the 10 observers, with the between observer standard errors.

3. Results

The average variance per observer was about 10 deg^2, which is within the range found in other studies (e.g. Hillis et al., 2004; Muller et al., 2007). One observer's data was excluded because her variances were about three times as high as the values for the other 11 observers. The histograms in Fig. 4 show the distributions of the settings in the two cue conflict conditions for the remaining 11 observers. Each distribution has a single peak between the values indicated by the two conflicting cues, which is consistent with observers relying on a weighted average on each trial. Since our reasoning is based on the variances in the settings all being due to random variability, we examined the distribution of the values in more detail.

We used an Anderson–Darling test (Anderson & Darling, 1952) to examine whether any of the distributions clearly deviated from a normal distribution when expressed in the units in which the averaging could occur (Eq. (2); we assume averaging in slant), and found that 5 of the 66 distributions (involving 3 of the 11 observers) deviated significantly from a normal distribution (at a 5% level). On the basis of chance, we would only expect three distributions that deviate from normality if the underlying distributions are normal (given the 5% level of the test). The difference is no cause for concern, but since one observers' data failed the test of normality for three of his six conditions, that observer's data was also removed. We report data for the remaining 10 observers in the rest of this paper.

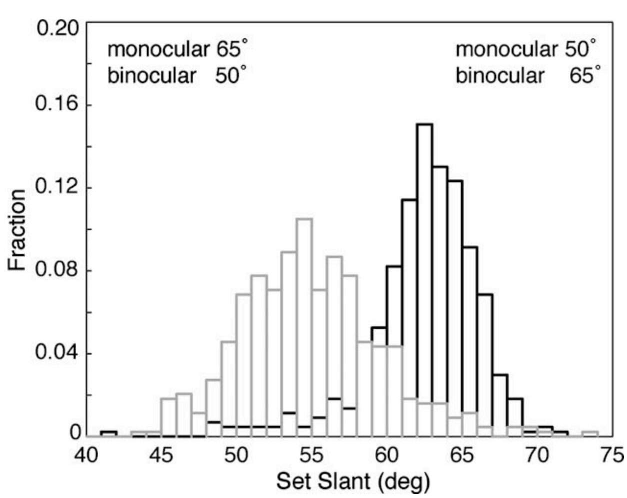

Fig. 4. Distribution of slant settings in the two cue conflict conditions for 11 observers (1° bins).

Fig. 5 shows the averaged combined variances in the cue consistent and cue conflict conditions for the 10 observers. The combined variances are the values we obtain for the right and left terms of Eq. (8). A paired *t*-test showed no significant difference between these two values ($p = 0.95$). Therefore, we cannot reject the counter-intuitive prediction formulated in Eq. (8), that the bars are equal in height. We determined the weight that observers assigned to each cue in the cue conflict conditions and used these to determine expected percentage increases in variability if only the best cue had been used. The observers' combined variance for the two cue consistent conditions was then multiplied by this percentage to obtain the expected combined variance if observers switched to using only the best cue when the cues are in conflict. The predicted combined variance for the best-cue-scenario is clearly larger.

Observers all took longer to respond on cue conflict trials than on cue consistent trials, which indicates that they were influenced by the conflict. This difference is in line with the intuition that matching the surfaces in cue conflict trials is more difficult than doing so in cue consistent trials. A paired *t*-test showed that the average time to respond was significantly higher in the cue conflict conditions (Fig. 6; $p < 0.001$).

4. Discussion

The intuitive prediction that the variances in matching cue conflict trials would be systematically higher (Eq. (9)) is not supported by our data. The data in Fig. 5 clearly shows that there is no systematic difference between the variances for the two kinds of trials. We therefore cannot reject the counter-intuitive prediction that the variances in perceiving the reference surface slant in cue conflict trials and cue consistent trials are equal (Eq. (8)). It is clear that if observers had switched to only using their best cue, the combined variance would have been higher in the cue conflict conditions. A similar increase in combined variance as predicted for switching to the best cue only would be found for a standard deviation of 6% in the cue weights.

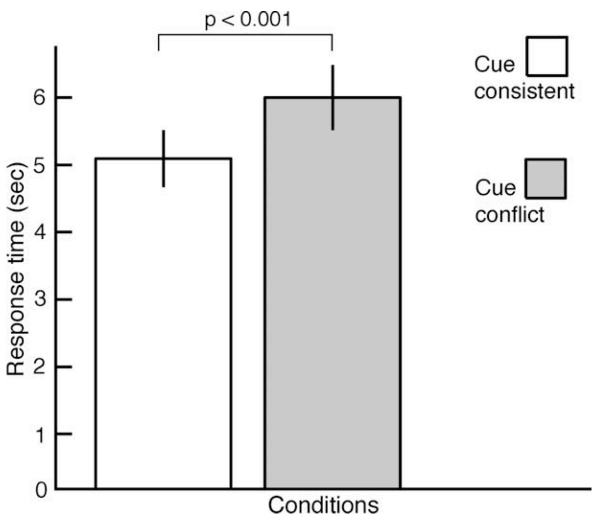

Fig. 6. Median response times, averaged across observers, for the cue consistent conditions (white bar) and the cue conflict conditions (gray bar) with the between observers standard errors.

Observers consistently took longer to match cue conflict trials. This justifies our choice of reference slants for the stimuli, because it means that observers must have noticed the conflict in some way. However, this difference in response time could be a confounding factor.

Observers probably took longer to answer because the two surfaces never looked exactly the same in the cue conflict trials (Muller et al., 2007) so that they were satisfied less quickly. Taking more time could, however, in itself help observers to make more precise settings. If this is the case, then the failure to find the intuitive relationship between the cues (Eq. (9)) may just be the result of observers making their decision when they reach a certain level of precision.

5. Control experiment

We conducted a control experiment in which observers could not vary the time they took to make their decisions, so they could not take longer to make their settings to compensate for a lower instantaneous precision.

The control experiment was essentially the same as the main experiment, but the time that observers could look at the surfaces for making the slant judgments was limited. This way observers could not decrease their variance for the conflict trials by taking more time.

5.1. Methods

We used the same setup and stimuli as in the main experiment. The reported data is for the same 10 observers for whom we reported the data in the main experiment (Fig. 5). The stimuli in this control experiment were also identical to those in the main experiment, but we used a two-alternative forced choice paradigm. On each trial observers were presented with two slanted surfaces, shown side by side for 1.5 s. Observers were asked to judge whether the slant of the probe surface was larger or smaller than that of the reference surface. Pilot sessions with two observers and various stimulus durations showed that 1.5 s was short enough to lead to more variable settings in all conditions. This was important because finding more variable responses in both conditions is an indication that viewing time limits performance in both conditions, and thus that the full presentation time is used. The slant of the probe surface was varied according to a staircase procedure. The same six reference surfaces that we used in the main experiment were used as the reference for six staircases that were all interleaved during one session. The initial slant of the probe surface was 57.5° for all staircases. If an observer judged the probe surface to be more slanted than the reference surface, the probe surface was 2° less slanted on the next trial for that condition. If it was judged to be less slanted, it was 2° more slanted on the next trial for that condition.

5.2. Analysis

We determined the proportion of "more slanted" responses for every combination of slants of the reference surface. We fit a cumulative Gaussian distribution to these values for each condition (weighted by the square root of the corresponding number of presentations) to estimate the variance in the settings. From these variances we again calculated the combined variances of Eqs. (8) and (9) using the procedure outlined in Fig. 3.

5.3. Results

The results of the control experiment are shown in Fig. 7. The combined variances are larger than in the main experiment in both conditions, confirming that the precision was influenced by limiting the time. As in the main experiment (Fig. 5), the combined variances for the cue consistent conditions (white bar) and for the cue conflict conditions (gray bar) are not significantly different ($p = 0.96$). This supports the counter-intuitive prediction made by optimal cue combination as formulated in Eq. (8). Again, the predicted value for using the best cue only is clearly larger.

5.4. Discussion

The results of the control experiment suggest that the fact that observers took longer to make their settings on cue conflict trials than on cue consistent trials in the main experiment did not influence the precision of their settings. Observers probably did not take more time in the cue conflict trials than in the cue consistent trials to compensate for being less certain about the slant, but more likely postponed their decision because the match looked less satisfactory.

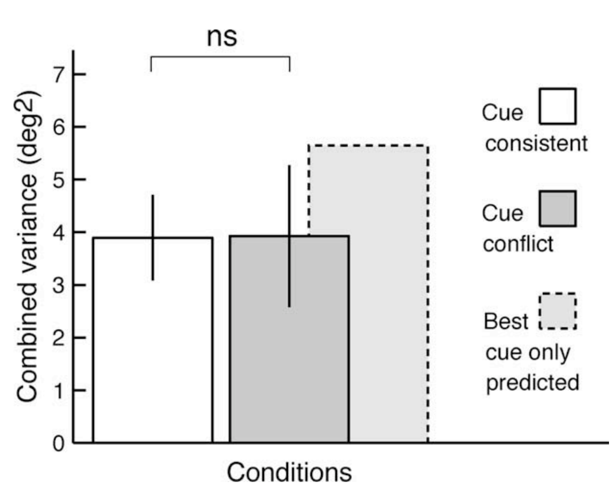

Fig. 7. Results for the control experiment. The bars show the combined variances (averaged across 10 observers) in the cue consistent conditions (white bars) and cue conflict conditions (gray bars), as well as the predicted values for using the best cue only (light gray bars). with the between observer standard errors.

6. General discussion

In this study, observers were asked to match the slants of two virtual surfaces. In the trials with conflicting cues the cue consistent probe surface could never look completely identical to the reference surface, whereas in the cue consistent trials it could (Hogervorst & Brenner, 2004; Muller et al., 2007). Therefore, having to match two surfaces makes the prediction based on optimal cue combination (that the variances for conflict trials and consistent trials will be equal), even more counter-intuitive than it is for judgments of a single surface. Nevertheless, observers were just as precise in matching the slants of the cue conflict stimuli as they were in matching the slants of the cue consistent stimuli. Moreover, observers were more precise than they would have been if they had switched to using only the most reliable cue in the cue conflict conditions, although this was not statistically significant (paired *t*-tests).

Several studies have shown that cues are combined optimally when the stimuli do not contain noticeable conflicts (Hillis et al., 2004; van Beers, Sittig, & Denier van der Gon, 1999). In our study, observers seem to have registered the conflicts, because they took longer to match reference surfaces for which the cues were in conflict than ones for which they were not (Fig. 6). If we assume that the cues in our study were combined optimally when not in conflict, our results imply that these cues are also combined optimally when they are noticeably in conflict.

Deviations from optimality have been found when the cues do not clearly belong to the same object because the cues are not spatially congruent (Gepshtein, Burge, Ernst, & Banks, 2005). Cue combination can therefore be regarded as robust in the sense that the visual system does not combine cues at all cost (see Knill, 2007). Therefore, one might expect cues to no longer be combined optimally as soon as a conflict is registered. Our results show that this is not the case. At some point, the conflict will become too large for optimal cue combination (van Ee et al., 2002). The conflicts used in our study were quite large (15°, which is about 10 times as large as the SD in the settings), and yet they apparently fall within the range for which cues are combined in a statistically optimal manner.

Acknowledgment

This research was supported by the Netherlands Organization for Scientific Research (NWO; MaGW grant 452-02-007).

References

Anderson, T. W., & Darling, D. A. (1952). Asymptotic theory of certain ‘‘Goodness of Fit” criteria based on stochastic processes. *Annals of Mathematical Statistics, 23*(2), 193–212.

Banks, M. S., Hooge, I. T., & Backus, B. T. (2001). Perceiving slant about a horizontal axis from stereopsis. *Journal of Vision, 1*(2), 55–79.

Brenner, E., Granzier, J. J., & Smeets, J. B. (2007). Combining local and global contributions to perceived colour: an analysis of the variability in symmetric and asymmetric colour matching. *Vision Research, 47*(1), 114–125.

Ernst, M. O., & Bulthoff, H. H. (2004). Merging the senses into a robust percept. *Trends in Cognitive Sciences, 8*(4), 162–169.

Gepshtein, S., Burge, J., Ernst, M. O., & Banks, M. S. (2005). The combination of vision and touch depends on spatial proximity. *Journal of Vision, 5*(11), 1013–1023.

Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision, 4*(12), 967–992.

Hogervorst, M. A., & Brenner, E. (2004). Combining cues while avoiding perceptual conflicts. *Perception, 33*(10), 1155–1172.

Knill, D. C. (1998). Discrimination of planar surface slant from texture: human and ideal observers compared. *Vision Research, 38*(11), 1683–1711.

Knill, D. C. (2007). Robust cue integration: A Bayesian model and evidence from cue-conflict studies with stereoscopic and figure cues to slant. *Journal of Vision, 7*(7 5), 1–24.

Knill, D. C., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research, 43*(24), 2539–2558.

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*(3), 389–412.

Muller, C. M. P., Brenner, E., & Smeets, J. B. J. (2007). Living up to optimal expectations. *Journal of Vision, 7*(3), 1–10.

van Beers, R. J., Sittig, A. C., & Denier van der Gon, J. J. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of Neurophysiology, 81*(3), 1355–1364.

van Ee, R., van Dam, L. C., & Erkelens, C. J. (2002). Bi-stability in perceived slant when binocular disparity and monocular perspective specify different slants. *Journal of Vision, 2*(9), 597–607.

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*(18), 2685–2696.