



City Research Online

City St George's, University of London

Citation: Tomassini, A., Morgan, M. J. & Solomon, J. A. (2010). Orientation uncertainty reduces perceived obliquity. *Vision Research*, 50(5), pp. 541-547. doi: 10.1016/j.visres.2009.12.005

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/16919/>

Link to published version: <https://doi.org/10.1016/j.visres.2009.12.005>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Orientation Uncertainty Reduces Perceived Obliquity

Alessandro Tomassini, Michael J. Morgan & Joshua A. Solomon*

Optometry Department, City University, London, UK.

The influence of prejudice on perception should be greatest when certainty about stimulus identity is least. We exploited this relationship to reveal visual biases for the cardinal orientations: vertical and horizontal. Specifically, when we increased the variance of orientations in an array of grating patches, estimates of the mean became less oblique. This result is consistent with a stable prior, or prejudice, for those orientations most prevalent in natural scenes.

1. INTRODUCTION

Many contemporary theorists describe vision as form of Bayesian inference (Knill, Kersten, & Yuille, 1996). That is, our perceptions result from the combination of prior beliefs with data we gather from the environment. As an example, consider the convex appearance of concave faces (Gregory, 1970). To a Bayesian, this phenomenon suggests a belief that faces are concave with a very low probability (Yellott & Kaiwi, 1979).

The influence of prior knowledge is demonstrably greatest when certainty about stimulus likelihood is least. Possibly the earliest example was the finding that when depth cues become impoverished, the distance of any object appears closer to 2 m than it really is (Gogel, 1969). However, it is not clear whether this bias has anything to do with prior knowledge. The 2 m distance has no known behavioural relevance.

*To whom correspondence should be addressed. E-mail: J.A.Solomon@city.ac.uk

In the laboratory, behavioural relevance can be manipulated. For example, using virtual reality, Kording and Wolpert (2004) taught observers to compensate for a shift in apparent hand positions. The ability to compensate for subsequent shifts was found to vary with the fidelity of the virtual image. When that image was blurred, pointing behaviour was biased toward the shift they learned.

Rather than bias observers in the laboratory, two previous studies (Stocker & Simoncelli, 2006; Weiss, Simoncelli, & Adelson 2002) exploited the smooth and slow motions known to dominate the statistics of natural flow fields (Roth & Black, 2007), and fit a Bayesian model to the relationship between stimulus contrast and perceived speed (cf. Hammet, Champion, Thompson & Morland, 2007; Thompson, Brooks, & Hammet, 2006). Our approach is similar, except instead of relying on noise intrinsic to the visual system, we decided to directly manipulate uncertainty by adding variance to the stimulus.

We derived our predictions from the predominance of approximately horizontal and vertical contours in our environment (Coppola, Purves, McCoy, & Purves 1998; Switkes, Mayer & Sloan, 1978). It has already been established that humans have greater sensitivity and acuity for simple stimuli having these cardinal orientations (Appelle, 1972). A preference for motion along cardinal axes has also been demonstrated (Mansfield, 1974). These behavioral “oblique effects” are thought to be consistent with physiological studies showing that relatively few neurons are tuned to oblique orientations (Andrews & Schluppeck, 2000). What we sought to determine was whether perception actually would shift toward the cardinal orientations when confidence in the true stimulus orientation was low.

2. METHODS

In our initial test of this hypothesis, we asked five normal, naïve observers to align two “comparison” dots with their estimate of the average orientation in an array of Gabor patterns. The orientation of each Gabor was randomly selected from a Gaussian distribution (see Fig. 1). We expected estimates of the mean to be closer to the closer cardinal axis. No efforts were made to restrict the contents of the room from view. We did not want to discourage observers from adopting any typical prior they might have for the orientation content of indoor scenes. Observers were asked to fixate on the centre of the iMac display on which the stimuli were presented, but this fixation was not enforced in any way.

At a viewing distance of 57 cm, each comparison dot had an angular subtense of 0.2° . It was white with a luminance of 221 cd/m^2 . The two dots were presented at a viewing eccentricity of 5° , on opposite sides of fixation. Observers adjusted the azimuth of the dots using two keys on the keyboard, and clicked the mouse button when it appeared to match that of the average orientation. There was an inter-trial interval of 1 s. In a control experiment we fixed the azimuth of the dots, and four naïve observers adjusted the orientation of the Gabor array using the same two keys. Three of these latter observers also participated in the main experiment.

Each Gabor in the array had a spatial frequency of 6.9 c/deg , a spatial phase of either $\pi/2$ or $-\pi/2$ (randomly chosen for each Gabor pattern), a space constant (σ) of 0.072° , a mean luminance of 111 cd/m^2 , and a contrast of 0.99. Prior to each trial, the Gabor patterns were placed, one at a time, in a $5.7^\circ \times 5.7^\circ$ square. The placement of each Gabor pattern was random, with the constraint that no two Gabors could have centres closer than 0.43° (i.e. 6σ). The number of Gabor patterns required to fill each square was 132 ± 4 . The entire array appeared within a Gaussian window, the space constant of which was 1.4° . In separate blocks, we used Gabor arrays having 0.1 s, 0.5

s, and response-terminated displays. Only response-terminated displays were used in the control experiment.

3. RESULTS

3.1. Main Experiment: Response bias

In this study, we were primarily concerned with perceptual biases. In particular, we wanted to know whether the variance of orientations affected their apparent mean. However, what we measured were response biases, i.e. differences between the true mean orientation and the azimuth of the comparison dots. (NB: We use positive numbers to represent clockwise tilts. Thus positive response biases indicate responses that are clockwise with respect to unbiased responses.) Perceptual and response biases are not necessarily the same; for example, observers could have a perceptual bias towards the cardinal axis, but a response bias in the direction of making the comparison dots less vertical than the gratings. We shall start by analysing response biases only.

Our observers were remarkably precise in their estimates of orientation. When all the data were pooled without regard to observer, display duration, orientation variance or mean physical tilt, the standard deviation (SD) of response bias was just 9.8° . Nonetheless, there were a few trials, even with long durations and low orientation variance (as in Fig. 1a), for which the bias was strangely large. Perhaps on these trials, observers mistakenly clicked the mouse button, indicating alignment, before they had actually moved the comparison dots from their random starting positions. We decided to establish a rather conservative criterion for removing these outliers from the data set. Thus we kept all data within 8 SDs of zero bias. With this criterion, exactly 5 trials were discarded, and the SD of the remaining 4,315 fell to 9.3° .

Each point in Fig. 2 shows the average response bias of our five naïve observers, collapsed across display duration and orientation variance. Error bars contain 2 standard errors (SEs), i.e. $\frac{2}{\sqrt{n_i}}$, where b_i denotes observer i 's response bias when the Gabor orientations were selected from a distribution with mean s , and n_i represents the number of trials observer i performed in that condition.

As a rule, these data points fall in the shaded regions of this figure, indicating a tendency to give the comparison dots an alignment that was *more oblique* than the mean orientation in the stimulus array. Exceptions to this rule, which occur at mean physical tilts of $\pm 55^\circ$ with respect to vertical, suggest that biases toward (or away from) the vertical and horizontal axes may not be equal.

The smooth curve in Fig. 2 satisfies the equation

$$b = \frac{a \sin(\gamma s)}{1 + \cos(\gamma s)} \quad (1)$$

where s is the mean physical tilt of the stimulus and r is the response bias. The parameter a determines the maximum bias, and the parameter γ determines how much stronger biases away from the vertical axis are than biases away from the horizontal axis. (NB: $\gamma = 0$.) When $\gamma = 1$, these two biases are equal, and the smooth curve will cross the line of zero bias when the stimulus array has a mean physical tilt that is halfway between vertical and horizontal (i.e. $\pm 45^\circ$). Below, where the effect of orientation variance is examined, it will be useful to summarise observer performances using the best-fitting values of a and γ . Fig. 2 shows the curve that best fits the entire data set. The mean distance between it and each data point is 1.1 SEs.

In summary, the data show a large response bias away from the cardinal axes.

3.2. Control Experiment: Response bias

We were surprised to find such large response biases away from the closest cardinal axes. Of course, this bias has no bearing on our hypothesis, which concerns the effect of orientation variance. Nonetheless, we wondered about the origin of this more general bias. It too might have had a basis in uncertainty. If observers were even less certain about the azimuth of the comparison dots than they were about the mean of the Gabor array, then the comparison dots would be even more susceptible to perceptual biases. If such biases did indeed favour the cardinal axes, then the comparison dots would have to be further from these axes in order to appear aligned with the Gabor arrays.

Another possibility is that the bias away from the cardinal axes has nothing at all to do with uncertainty. Perhaps observers are simply reluctant to report more cardinal orientations. We tested this possibility in a control experiment that used exactly the same stimuli (and some of the same observers) that were used in the main experiment. The only difference was, this time observers rotated the Gabor array until its mean appeared to be aligned with the azimuth of the dots.

Results of the control experiment (Fig. 3) are consistent with this latter possibility. Each point shows the average response bias collapsed across orientation variance. Like those from the main experiment (in Fig. 2), these data points fall in the shaded regions of this figure, indicating response biases away from the closest cardinal axis.

3.3. Effect of orientation variance on response bias

Fig. 4 illustrates the effects of stimulus uncertainty on response bias. The top panels contain the same data as Fig. 2, here segregated on the basis of stimulus variance. The data in both panels form the same general pattern, indicative of response biases away from the closest cardinal axis, but the high-variance biases are

not as large as the low-variance biases. Observers, therefore must have found the mean orientations high-variance arrays to be closer to the closest cardinal axis than the mean orientations of otherwise identical arrays with low orientation variance.

The bottom panels of Fig. 4 contain the same data as Fig. 3 segregated on the basis of stimulus variance. Again, the pattern in both panels is indicative of response biases away from the closest cardinal axis, but here the high-variance biases are larger than the low-variance biases. This means that the observers rotated the high-variance Gabor arrays farther from the closest cardinal axis when attempting to make its average orientation parallel to the azimuth of the dots. Thus, the effect of variance found in the main experiment is confirmed by the results of the control experiment: the average orientation of high-variance arrays appears closer to the closest cardinal axis than that of low-variance arrays.

To quantify this, our central result, we fit Equation (1) to the data in each panel of Fig. 4. Although data from the control experiment (bottom panels) suggest an even larger effect, the statistics we quote below were all derived from the main experiment (top panels), in which more responses were collected from more observers. When the parameter y was fixed at the best-fitting value for the entire set of data (0.45, see Fig. 2), best fits of Equation (1) suggest twice as much response bias in the former conditions. (When y was free to vary, the ratio of best-fitting a values was even larger: 2.8:1.0.)

To assess the significance of our central result, we divided the bias of each response by the smooth curve seen in Fig. 2, and performed a full-factorial three-way ANOVA (2 stimulus variances x 3 durations x 5 observers) on the quotients. Significant main effects were found for stimulus variance [$F(1, 4290) = 15.08$, $p < 0.0001$] and observer [$F(4, 4290) = 14.03$, $p < 0.001$]. There was no significant effect

of display duration [$F(2, 4290) = 0.62, p = 0.54$], nor were there any significant interactions ($p > 0.1$).

The preceding statistics indicate significant individual differences. To explore these differences, we averaged the aforementioned quotients separately for each observer.¹ These averages were all greater than zero, indicating at least a trend for response biases away from the closest cardinal axis, however only three of the five means (one for each observer) were significantly greater than zero (one-tailed t-tests, $p < 0.005$). When trials with the larger stimulus variance were excluded, the remaining quotients were significantly greater than zero in four of the five subjects (one-tailed t-tests, $p < 0.001$). These last two analyses suggest a significant effect of stimulus variance in the data of one particular observer. This significance was confirmed with a two-way ANOVA [2 stimulus variances x 3 durations; $F(1, 858) = 12.2, p < 0.0005$]. Analogous tests on quotients calculated from the other observers' responses indicate a significant [$F(1, 858) = 10.3, p < 0.002$] main effect of stimulus variance in one other observer, a marginally significant [$F(2, 858) = 4.09, p < 0.02$] main effect of display duration in another observer, and no significant interactions in any.

Fig. 5 contains another illustration of our central result. For each mean orientation, we have subtracted the average bias recorded when SD was 2° from that recorded when it was 14° . Most of the data fall in the unshaded regions of the figure, indicating a tendency to give the comparison dots an alignment that was relatively less oblique (i.e. closer to the closest cardinal axis) when orientation variance was relatively high.

¹ Each observer's data has been summarised with graphs analogous to those in the new Fig. 4. They have been uploaded as Supplementary Material. One file for the main experiment; another for the control experiment.

3.4. Response variance

We probed observer uncertainty by measuring the SD of their response biases. For each orientation variance and mean physical tilt, each observer's data were pooled without regard to display duration.² From each of these pools, the standard deviation of response bias was calculated. Fig. 6 shows these standard deviations, averaged across observer. As a rule, these averages are larger in the right panel, where they describe trials with greater orientation variance. That is, when the stimuli had high variance, so did the observers' responses. Within each panel, there is also a clear oblique effect. That is, observers appear to have relatively low uncertainty (i.e. greater precision) when reporting orientations close to the cardinal axes.

4. DISCUSSION

4.1. Bayesian theory

To our knowledge, Bayesian theory is the only theory that explicitly links perceptual bias with uncertainty. Specifically, it predicts that perceptual bias should increase with uncertainty. Our immediate concern was the perceptual bias for mean orientation, and whether it would increase as observers became less certain of that mean. The immediately preceding analysis suggests that observers were indeed less uncertain about the mean orientation of the low-variance stimuli than they were about that of the high-variance stimuli, so all that remains is to determine the effect that orientation variance had on the perceived mean.

² For the Gabors with a 2° SD, both uncertainty and response bias seems to have shrunk with display duration. Specifically, for the 0.1 s, 0.5 s and response-terminated displays, average response SDs were 7.3, 6.5 and 4.3°, and average “bias quotients” (i.e. response biases divided by the smooth curve in Fig. 2) were 2.5, 2.0 and 1.4, respectively. On the other hand, the Gabors with a 14° SD did not show that kind of an effect. Response SDs were 9.3, 9.6 and 7.9°, and the corresponding bias quotients were 0.5, 0.9 and 0.6.

As noted above, direct measurements of the perceived mean are impossible because other biases can also affect observer responses. However, if we can assume that all of those other biases are invariant with orientation variance, then the aforementioned difficulty can be circumvented by subtracting the response bias with a low orientation variance from the response bias with a high orientation variance, as in Fig. 5. These differences must therefore indicate a change in perceptual bias. Since the response biases we collected with high-variance stimuli were closer to the cardinal axes than the response biases we collected with low-variance stimuli, the perceptual bias that becomes manifest under conditions of uncertainty is a perceptual bias to see things as more vertical or horizontal than they really are.

4.2. Tilt normalisation

This result is consistent with the widely documented tendency to see slightly tilted stimuli as upright (Howard, 1982). It was first described by Gibson and Radner (1937), who argued that this “tilt normalisation” was a gradual effect. Tilted stimuli *eventually* come to appear less tilted than they really are. Given this argument, it may seem strange that we did not find any significant effect of display duration on response bias, however it seems that the literature is not entirely clear on the temporal aspects of tilt normalisation.

Possibly the most similar antecedents to the current study are those in which observers had to compare a briefly flashed oriented stimulus with a continuously present comparison line. In one of these antecedents (Andrews, 1965), it was a briefly flashed line segment’s tilt that observers adjusted so as to appear parallel with that of a continuously visible thread. Like our current results, Andrews’s indicate that alignment was perceived when the tilt of the briefly flashed stimulus was more similar to that of the closest cardinal axis than the continuously present comparison’s was.

Andrews's (1965) results can be considered supportive of Gibson and Radner's (1937) argument for a gradual normalisation process because he found that response biases decreased, and eventually changed sign, as the duration of the stimulus flash increased. However, two years later, Andrews (1967) again reported that response bias depended on flash duration, but in this case the bias changed in the other direction. Thus, it seems there is no consensus, even within laboratories, regarding the effect of display duration on tilt normalisation.

4.3. A Bayesian model for tilt normalisation

Tilt normalisation may be explained in Bayesian terms. A diagram of our model appears in Fig. 7. The top half of the figure shows a caricature of Bayes' Theorem. If visual estimates of the mean were based on one randomly selected Gabor pattern, then the likelihood function should have an SD that is slightly larger than that used to create the stimulus. This is because the visual system has only limited fidelity. To obtain the curves shown in Figs. 5 and 6, we assumed that the visual system supplied additional jitter, with an SD between 8 and 10° to each Gabor pattern in each array. (We have made no attempt to model the curves in Figs. 2 and 3. Differences between perceptual and response bias will be ignored in this modeling exercise.)

This amount of "early noise" is greater than that inferred from experiments on variance discrimination, which used somewhat different stimuli (Morgan, Chubb & Solomon, 2008). Nonetheless, more direct estimates of this early noise from two additional control experiments (described in Appendix B) support these values. Similar experiments using Morgan et al's stimuli have yet to be performed.

Other models of the oblique effect exist (Heeley et al, 1997; McMahon & MacLeod, 2003), but ours is relatively simple. We assume that the visual system adds more noise to oblique orientations than it does to vertical and horizontal. Specifically,

we assume that the SD of the Gaussian noise added to each element is a parabolic function of that element's orientation:

$$\sigma = \sigma_0 + \sigma_1 \theta^2 \quad (2)$$

where the element's orientation θ is in degrees from horizontal. (A description of why we chose this particular parabola appears in Appendix A.)

Of course, estimates of mean orientation can be based on more than one Gabor pattern. Just like the SE in statistics, the SD of the likelihood function for the mean should decrease with the square-root of N , the number of Gabor patterns used in the estimate. To obtain the curves shown in Figs. 5 and 6, we assumed $N = 3$. Regardless how many Gabors observers use, any increase in their variance must be reflected in the spread of the likelihood function. "Uncertainty" is the name given to this spread by Bayesian scholars (e.g. Knill & Pouget, 2004).

If our observers did use just 3 Gabor patterns in their estimates of mean orientation, they would have been much less efficient than Morgan et al's (2008), who seem to have been able to use between 4 and 10. We are not entirely certain, but one reason for such inefficiency may be that our observers were all unpracticed students. Morgan et al's, on the other hand, were all professional psychophysicists. Another reason is that Morgan et al did not use a Gaussian window, which seriously limited the visibility of all but a few of our Gabor patterns.

Inefficiency alone cannot tell us why perceptual biases change with uncertainty. As can be seen in Fig. 7, the likelihood function remains centred on the true mean orientation, which is indicated by the dashed line. Biases arise when the likelihood function is combined with a set of prior beliefs, which in our model arise from the statistics of natural scenes. We very loosely fit a parabolic curve to measurements of broadband, indoor environments (Switkes et al, 1978). These

measurements appear as blue dots in the bottom half of Fig. 7. The parabola can be described by the formula
$$P(\theta) = \frac{1}{N} \left(1 + a \cos^2(\theta - \theta_0) \right)$$
, where θ is orientation in degrees from horizontal. This formula, which produces a minimum at 49.3° , provides an excellent fit to Switkes et al's data when the steepness parameter $a = -0.014$. When $a = 0$ the distribution becomes flat. To obtain the curve shown in Figs. 5 and 6, we had to exaggerate the anisotropy of Switkes et al's measurements a bit: we used $a = -0.04$. (A description of why we chose $N = 3$ and $a = -0.04$ appears in Appendix A.)

We used the maximum of the product of likelihood and prior for our model's estimates of mean orientation. The curve in Fig. 5 shows how these *maximum a posteriori* estimates are drawn closer to the closest cardinal axis when uncertainty is high.

4.4. Conclusion

A human observer could implement Bayes' rule simply by giving more weight to those orientations with higher prior probability. This preferential weighting could arise from a predominance of neurons tuned to these (cardinal) orientations. Indeed, evolution may have produced anisotropic distributions of stimulus preference, simply so that organisms could behave rationally in conditions of uncertainty. Specifically, our results show that when observers are uncertain about orientation, they have a prejudice against visual tilt.

ACKNOWLEDGMENTS

This research was supported by a grant from the Engineering and Physical Sciences Research Council (EP/E064604). The first paragraph of the Introduction appeared as part of News From The Field in *Attention, Perception & Psychophysics*, 71/2, p. 215 (2009).

APPENDIX A

Two parameters were adjusted to obtain the model fits shown in Figs. 5 and 6: the sample size N and the steepness parameter a . Before describing how we selected their values, we first describe how Switkes et al's (1978) data were fit. The data in their Fig. 2a re-appear in our Fig. 7. When reflected about the vertical (i.e. 90°) axis, those data form a parabola, such that the relative frequency of occurrence can be summarised with an expression having the form $f(\theta) = 1.26 - 0.004\theta^2$, where θ denotes the degrees from horizontal. A least-squares fit to the data has

The minimum of this parabola occurs at 49.3° , and halfway between horizontal and the intercardinal axis (i.e. at 22.5°) the relative frequency of occurrence is 1.26. The formula

ensures that both of these properties will remain independent of the steepness parameter a .

The value $a = -0.04$ was selected simply because it produced a good-looking Fig. 7. The prior (red curve) was neither too flat in the top panel nor too steep in the bottom panel. Next, we fit the data in Fig. 6 with a parabolic function mapping Gabor orientation to internal noise. Specifically, we found the a that minimised the mean squared difference between the SD of our model's Gaussian likelihood

$$(A1)$$

and the data in Fig. 6 for each (integer) value of N between 1 and 10. In the preceding equation μ and σ^2 represent the mean and variance of the Gabor orientations, respectively.

Equation (A1) fits the data in Fig. 6 best when

The root-mean-squared (RMS) error

between the fit and our data was 0.11 (Naperian) log units. However, a better fit to the data in Fig. 5 could be obtained with lower values of N . We felt that $N = 3$ was the best compromise. When $N = 3$, the data in Fig. 6 are best fit when

In this case the RMS error between the fit and our data is 0.14 log units.

Finally, with the other parameters fixed, we examined the effect of halving a . This modification resulted in a greater RMS error between our model's predictions and the data in Fig. 5.

APPENDIX B

Two control experiments were conducted to obtain more direct estimates of the putatively early noise that limits orientation acuity. In both experiments, two otherwise identical, differently oriented Gabor patterns appeared at fixation with an inter-stimulus-interval of 1.5 s. The orientation of the first Gabor was selected randomly from a uniform distribution over all possible orientations. It was flashed for 0.15 s. The second Gabor remained visible until the observer responded. In one experiment (A), the observer (author JAS) had to decide whether the second was rotated clockwise or anti-clockwise of the first. In the other experiment (B), the observer rotated the second Gabor using the keypad, until it matched his memory of the first Gabor's orientation. The two experiments were conducted in ABBA fashion.

Differences between the correct answer and each of the observer's 80 responses in experiment (B) had an SD of . Note that if these errors are samples from a Gaussian distribution, then the 95% confidence interval for this SD is , where is the inverse of the cumulative chi-square distribution, with M degrees of freedom; specifically

$$. \quad (B1)$$

The psychometric function for experiment (A) can be predicted using Signal-Detection Theory (Green & Swets, 1966) and our estimate for the SD of Gaussian noise added by the visual system to the orientation of the first Gabor. (We assume that any noise added to the second, response-terminated Gabor is negligible.) Percent correct should be $\frac{1}{2} [1 + \Phi(\frac{\Delta\theta}{\sigma})]$, where Φ represents the cumulative distribution function of any standard normal random variable X , and $\Delta\theta$ is the (acute) angle between the two Gabors' orientations.

As can be seen in Fig. B1, the actual data collected in experiment B conform well to this prediction. Thus we can be reasonably confident that $8 - 10^\circ$ is not an overestimate for the SD of early noise that perturbs the apparent orientation of each Gabor.

REFERENCES

- Andrews DP (1965) Perception of contours in the central fovea. *Nature*, 205:1218-1220.
- Andrews DP (1967) Perception of contour orientation in the central fovea part1: Short lines. *Vision Res.*, 7:975-997.
- Andrews TJ, Schluppeck D (2000). Ambiguity in the perception of moving stimuli is resolved in favour of the cardinal axes. *Vision Res.*, 40:3485-3493.
- Appelle S (1972) Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals. *Psychol. Bull.*, 78:266-278.
- Coppola DM, Purves HR, McCoy AN, Purves D (1998) The distribution of oriented contours in the real world. *Proc. Natl. Acad. Sci.*, 95:2621-2623.
- Gibson JJ, Radner M (1937). Adaptation, after-effect, and contrast in the perception of tilted lines. I. Quantitative studies. *J. Exp. Psychol.*, 20:453-467.

- Gogel WC (1969) The sensing of retinal size. *Vision Res.*, 9:1079-1094.
- Green DM, Swets, JA (1966) *Signal detection theory and psychophysics* (Wiley).
- Gregory RL (1970) *The intelligent eye* (McGraw-Hill).
- Hammet ST, Champion P, Thompson P, Morland, AB (2007) Perceptual distortions of speed at low luminance: Evidence inconsistent with a Bayesian account of speed encoding. *Vision Res.*, 47:564-568.
- Heeley, DW, Buchanan-Smith, HM, Cromwell, JA, Wright, JS (1997) The oblique effect in orientation acuity. *Vision Res.*, 37:235-242.
- Howard IP (1982). *Human visual orientation* (New York: John Wiley).
- Knill DC, Kersten D, Yuille A (1996) *Perception as Bayesian Inference*, eds Knill DC, Richards, W (Cambridge Univ. Press), pp 1-21.
- Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *TINS*, 27:712-719.
- Körding KP, Wolpert DM (2004) Bayesian Integration in Sensorimotor Learning. *Nature*, 427:244-247.
- Mansfield RJQ (1974). Neural Basis of Orientation Perception in Primate. *Science*, 186:1133-1135.
- McMahon, MJ, MacLeod, DIA (2003). The origin of the oblique effect examined with pattern adaptation and masking. *J. Vis.*, 3:230-239.
- Morgan M, Chubb C, Solomon JA (2008). A 'dipper' function for texture discrimination based on orientation variance. *J. Vis.*, 8(11):9.
- Roth S, Black MJ (2007) On the spatial statistics of optical flow. *Int. J. of Computer Vision*, 74:33-50.
- Stocker AA, Simoncelli EP (2006) Noise characteristics and prior expectations in human visual speed perception. *Nature Neurosci.*, 9:578-585.

Switkes E, Mayer MJ, Sloan JA (1978) Spatial frequency analysis of the visual environment: Anisotropy and the carpentered environment hypothesis. *Vision Res.*, 18:1393-1399.

Thompson P, Brooks K, Hammet, ST (2006) Reducing stimulus contrast: speed can go up as well as down. *Vision Res.*, 46:782-786.

Weiss Y, Simoncelli EP, Adelson, EH (2002) Motion Illusions as Optimal Percepts. *Nature Neurosci.*, 5:598-604.

Yellott JJ, Jr, Kaiwi JL (1979) Depth inversion despite stereopsis: the appearance of random-dot stereograms on surfaces seen in reverse perspective. *Perception*, 8:135-142.

FIGURE LEGENDS

Fig. 1. Example stimuli and typical result in the main experiment. Each little oriented pattern is a Gabor. In (a) the Gabors are tilted $-75^\circ \pm 2^\circ$ clockwise with respect to vertical. On average, observers aligned the two white spots with an angle that was 4° farther from the nearest cardinal axis than the mean of this stimulus (i.e. -71°). In (b) the tilts are $-75^\circ \pm 14^\circ$. On average, observers were unbiased in their alignments of the two white spots with the mean of this array.

Fig. 2. Response bias versus tilt of the Gabor array. Results from the main experiment have been collapsed across observer, display duration and orientation variance to illustrate the general trend, which is that most data fall in the shaded regions, indicating response biases away from the closest cardinal axis. In all figures, each error bar contains two standard errors of its respective mean. In this figure, the smooth curve adheres to Equation (1), with parameter values and .

Fig. 3. Response bias versus azimuth of the dots. When data from the control experiment are collapsed across observer and orientation variance, they once again fall in the shaded

regions, indicating response biases away from the closest cardinal axis. The smooth curve adheres to Equation (1), with parameter values $\alpha = 0.5$ and $\beta = 0.5$.

Fig. 4. The effect of orientation variance on bias. Top and bottom rows show this effect in two complementary experiments. All responses are biased away from the closest cardinal axis. When the Gabors were fixed and observers rotated the dots (top panels), larger biases were recorded with the low-variance arrays. When the dots were fixed and observers rotated the Gabor array, larger biases were recorded with the high-variance arrays. To obtain the smooth curves, we used Equation (1). A different value of the amplitude parameter a was allowed for each panel in the figure, but y was not allowed to vary with orientation variance.



Fig. 5. Another look at the effect of orientation variance on bias in the main experiment. For each mean orientation, we have subtracted the average bias recorded when SD was 2° from that recorded when it was 14° . The solid curves illustrate the behaviour of the Bayesian model illustrated in Fig. 7.

Fig. 6. The effect of orientation variance on the precision of alignment. Here, an indication of observer uncertainty is the SD of response bias. It is clearly larger with the 14° (right panel) stimuli than it is with the 2° (left panel) stimuli. The solid curves illustrate the uncertainty in the Fig. 7's Bayesian model. Specifically, they show the SD of its Gaussian likelihood.


Fig. 7. Bayesian interpretation of results. (A) Bayesian inference predicts a shift in perceived orientation when a wide likelihood is multiplied by an anisotropic prior. (B) A parabola was loosely fit to the energy content of broad-band filtered, indoor scenes (data points from Switkes et al, 1978). Bayesian theory predicts that perception should reflect the posterior density; the product of a likelihood defined by a sample of stimulus orientations and the prior.

Fig. B1. Control experiments confirm early noise with an $\sim 9^\circ$ SD. Smooth curves contain predictions for 2AFC accuracy (i.e. was the second Gabor rotated CW or ACW with respect to the first Gabor) based on the maximum and minimum of the 95% confidence interval for the SD of errors in an orientation matching task that used identical stimuli. Data points show measured 2AFC accuracies. Error bars contain (binomial) 95% confidence intervals.


5-Figure 1





5-Figure 2





5-Figure 3




Bias (deg CW)



Bias (deg CW)




5-Figure 5




5-Figure 6

Uncertainty (deg)

2-deg S.D.



14-deg S.D.



Mean of Gabors (deg CW of Vertical)

5-Figure 7

