Vision Research 70 (2012) 34–43

Contents lists available at SciVerse ScienceDirect

Vision Research

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

Misperception of aspect ratio in binocularly viewed surfaces

Paul B. Hibbard [a,*], Ross Goutcher [b], Lisa M. O'Kane [c], Peter Scarfe [d]

[a] School of Psychology, University of St. Andrews, St. Mary's Quad, South Street, St. Andrews, Fife KY16 9JP, UK
[b] Psychology, School of Natural Sciences, University of Stirling, Stirling FK9 4LA, UK
[c] School of Psychology and Neuroscience, University of St. Andrews, St. Mary's Quad, South Street, St. Andrews, Fife KY16 9JP, UK
[d] Department of Experimental Psychology, University of Cambridge, Downing Street, Cambridge CB2 3EB, UK

ARTICLE INFO

Article history:
Received 20 February 2012
Received in revised form 3 August 2012
Available online 19 August 2012

Keywords:
Horizontal–vertical illusion
Slant
Binocular disparity
Aspect ratio

ABSTRACT

The horizontal–vertical illusion, in which the vertical dimension is overestimated relative to the horizontal direction, has been explained in terms of the statistical relationship between the lengths of lines in the world, and the lengths of their projections onto the retina (Howe & Purves, 2002). The current study shows that this illusion affects the apparent aspect ratio of shapes, and investigates how it interacts with binocular cues to surface slant. One way in which statistical information could give rise to the horizontal–vertical illusion would be through prior assumptions about the distribution of slant. This prior would then be expected to interact with retinal cues to slant. We determined the aspect ratio of stereoscopically viewed ellipses that appeared circular. We show that observers' judgements of aspect ratio were affected by surface slant, but that the largest image vertical:horizontal aspect ratio that was considered to be a surface with a circular profile was always found for surfaces close to fronto-parallel. This is not consistent with a Bayesian model in which the horizontal–vertical illusion arises from a non-uniform prior probability distribution for slant. Rather, we suggest that assumptions about the slant of surfaces affect apparent aspect ratio in a manner that is more heuristic, and partially dissociated from apparent slant.

1. Introduction

Forming an accurate understanding of the shapes and positions of objects in the world is a difficult computational problem, due to the ambiguity of the information reaching the eyes. The image created on the retina by any given object is consistent with any number of possible physical projections from the real world. A simple and much considered example of this problem (e.g. Thouless, 1931a, 1931b) occurs when an observer views a flat, elliptical surface. Although the surface has a fixed size and shape (e.g. circular), the size and shape of the image of the disc projected onto the retina is also determined by its distance from the observer, and its slant. This is a problem because the same retinal image would be created by a surface of a different size or shape, if it were presented at a different distance, or with a different slant. There is no way to unambiguously determine the actual shape and size of the surface from the shape and size of the elliptical contour projected onto the retina in this simple example.

However, in most everyday situations, there will be many useful sources of visual information for determining the distance and slant of the surface. For example, motion parallax, binocular disparity, linear perspective and texture can all provide information about the three-dimensional structure of the environment (Cutting, 1997). Here, we consider the information that is available from binocular cues, and how this is used in the perception of the slant and shape of surfaces.

In addition to the immediate sensory information available to an observer, prior knowledge or assumptions about the environment may also usefully influence our interpretation of depth. Bayes' rule (see for example Knill & Richards, 1996) provides the optimal means of combining uncertain visual information with such prior knowledge of the environment. Consider for example the problem of estimating the slant of a surface from binocular visual information. In this case, Bayes' equation can be formulated as follows:

$$p(S|B) \propto p(B|S) \cdot p(S) \qquad (1)$$

Here $p(B|S)$ is the likelihood function for the binocular disparity cue, and expresses the probability that the current disparity estimates B would have been obtained given a surface of true slant S. $p(S)$ is the prior, expressing the probability of encountering any particular slant S, irrespective of the current visual information. This therefore embodies our prior statistical knowledge of the structure of our environment. Finally, $p(S|B)$ is the posterior, expressing the probability that the particular slant is S, given the available binocular disparity information. The posterior is the totality of information

* Corresponding author. Fax: +44 1334 463042.
E-mail addresses: pbh2@st-andrews.ac.uk (P.B. Hibbard), ross.goutcher@stir.ac.uk (R. Goutcher), lisaokane1@gmail.com (L.M. O'Kane), ps611@cam.ac.uk (P. Scarfe).

0042-6989/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.visres.2012.08.003

available for making decisions about (i.e. estimating) the slant of the surface. In order to act upon this information, a decision rule must be applied. Two commonly used decision rules are to calculate the mean of the posterior distribution, or its maximum. The former is also referred to as minimum mean squared error estimation, since it minimizes this error in estimation, as defined by the difference between the true and estimated values. The latter, referred to as maximum likelihood estimation (MLE), selects the most likely value of the slant of the surface given all the available information (see for example Brainard & Freeman, 1997).

The advantage of a Bayesian approach is that, provided prior assumptions accurately reflect the observer's environment, inaccuracy in estimation can be reduced. Both precision and accuracy can be improved through the application of an appropriate prior. This theoretical approach also provides a compelling explanation of many biases that have been observed in such situations. For example, Bayesian priors have been used to account for perceptual biases is the cases of motion (Stocker & Simoncelli, 2006; Weiss, Simoncelli, & Adelson, 2002), distance (Yang & Purves, 2003), depth (Burge, Fowlkes, & Banks, 2010) and orientation (Girchick, Landy, & Simoncelli, 2011). It is thus possible to provide putative explanations of perceptual biases as reflections of the nature of the underlying priors, which are, in turn, based on the statistical regularity of information in the world. This logic has been used to relate natural scene statistics to observed psychophysical phenomena. One particularly relevant example of this is the explanation of the horizontal–vertical illusion proposed by Howe and Purves (2002).

The horizontal–vertical illusion is an example of the more general phenomenon that the apparent length of a line depends on its orientation. Typically, horizontal lines are perceived as shorter in extent than vertical lines of the same physical length. Indeed, the apparent length of a line varies systematically with orientation, with lines around 20–30° from vertical having the greatest apparent length (Cormack & Cormack, 1974; Craven, 1993; Pollock & Chapanis, 1952; Shipley, Nann, & Penfield, 1949). This phenomenon has been demonstrated in a wide variety of stimuli. These include the apparent distance between pairs of points; the relative lengths of horizontal and vertical lines, for example when presented in a 'T' (the bisection illusion) or an 'L' (the horizontal–vertical illusion) configuration and the apparent aspect ratio of simple geometrical figures such as ellipses and rectangles (Fick, 1851 (cited by Künnapas (1955)); Sleight & Austin, 1952; McManus, 1978). For example, Sleight and Austin (1952) found that observers presented with a perfect circle displayed what they referred to as a 'classical' illusion. That is, the circle appeared to be elongated in the vertical direction. Their results were more mixed when observers judged the shape of rectangles. However, using a much larger group of participants, McManus (1978) found that on average squares tended to appear elongated in the vertical direction by a factor of 1.58%.

A variety of explanations for this phenomenon have been proposed. One suggestion is that the bias reflects differences in the horizontal and vertical extent of the visual field of a binocular observer (Künnapas, 1955). Since the binocular visual field is wider than it is tall, a horizontal line will cover a smaller proportion of the extent of the field than will a vertical line of the same length. Differences in the apparent magnitude of horizontal and vertical extents then reflect the fact that extent is measured relative to the size of the field of view. Support for this account comes from the finding that the illusion is smaller under monocular viewing, when the visual field is more isotropic (Prinzmetal & Gettleman, 1993). Another type of explanation implicates properties of image formation in the retina, including imperfections in refractive properties (Avery & Day, 1969; Thompson & Shiffman, 1974; Valentine, 1912), the spacing of photoreceptors (Begelman & Steinfeld, 1971) or the distribution of retinal pigments (Bayer & Pressey, 1972). It has also been suggested that the phenomenon might be related to differences in the way that eye-movements are made in different directions. Because eye-movements in the vertical direction require more effort than those in the horizontal direction, a movement of the same extent will appear longer, simply as a result of the greater effort that would be required to make a saccade from one end of the line to the other (Schiffma & Thompson, 1974; Stacey, 1969).

A limitation of such accounts is that, while they all provide a clear route via which the phenomenon arises, they do not provide a theoretical explanation for why it should occur. In each case one may counter that, since the particular property of the visual system that is implicated in the phenomenon is stable (the shape of the binocular field; asymmetries in the refractive properties of the eyes; the distributions of photoreceptors and retinal pigments; the energy expended to make eye-movements), it might reasonably be taken into account when judging the lengths of lines. Indeed, it is well known that the visual system is able to calibrate itself readily to changes in its environment, for example in prism adaptation, and other adaptation effects (Webster, 2011), and this calibration might be expected to allow the above factors to be taken into account.

An alternative class of explanations, of particular relevance to the current discussion, provides a functional account of the phenomenon. These explanations suggest that biases reflect the fact that the possible three-dimensional layout giving rise to the retinal image is taken into account. If the goal of the observer is to determine the likely length of a line in 3D space on the basis of the retinal image, then one would only expect this to be unaffected by orientation if equal physical extents, at different orientations in the world, tended to project lines of the same extent in the retinal image. Howe and Purves (2002) tested this assumption directly by comparing the projected lengths of lines in the image with the actual 3D length of the projected features in the real world. They found that lines of the same length in the retinal image, but different orientations, tended to be projected by different 3D distances. Here, 3D distance refers to the magnitude of the separation in 3D space between the locations forming the endpoints of the line in the retinal image. In particular, they found that a constant retinal image separation tended to correspond to the smallest 3D separation in the horizontal direction, rose to a maximum for lines oriented approximately 20–30° away from vertical, then fell again slightly towards vertical. This pattern provides an excellent match to empirical data on the effects of orientation on line length. It also provides a very direct account of the phenomenon: the apparent length of a line depends on its orientation because the actual lengths of the projections of lines onto the retinal image depend on orientation. If the same retinal separation tends to correspond to greater distances between two points in the world for vertical image separations that for horizontal separations, then it might come as no surprise that observers see vertical lines as longer than horizontal ones. A similar explanation was proposed by Craven (1993), based on an analysis of the density of zero-crossings in filtered images. Note that these explanations are very different from the explanations criticized above. Other explanations attribute the bias to properties of the visual system that are not taken into account when making judgements of shape. The explanations proposed by Craven (1993) and Howe and Purves (2002) suggest that it is the fact that the visual system takes account of statistical regularities in our environment that gives rise to the bias.

The difference in 3D distances corresponding to equal intervals in images arises from statistical regularities in the 3D environment. Of most relevance here is the presence of a horizontal ground plane. In many scenes, there is a clear relationship between the vertical location of a point in the image, and its distance from the observer: points higher in the image tend to be further away. This relationship can be seen in the existence of height-in-the-field as a powerful depth cue (Cutting, 1997). There is no similar relationship between the horizontal location of a point and its distance. This means that, while on average the depth separation between horizontally separated points will be zero, for vertically separated points it will not. As a result, a given pair of horizontally separated points in the image is likely to correspond, on average, to a smaller 3D separation in the environment than a pair of points that are separated vertically in the image.

This account of the phenomenon is related to other explanations (Gregory, 1963, 1973; Woodworth, 1938) collectively referred to as perspective theories by Wolfe, Maloney, and Tam (2005). These theories explain the bias as a consequence of the differential size scaling of lines of different orientations. According to this view, the arrangement of lines in an image represents a configural depth cue, and some of the lines will therefore be interpreted as if extended in depth out of the fronto-parallel picture plane. This is consistent with Howe and Purves' (2002) explanation, in which vertical lines in the image appear longer because the estimation of their length takes account of the fact that they are likely to be slanted away from frontoparallel.

An important limitation of this specification is that it does not take account of perceptual information about depth. Such information would be available in most situations, and would be expected to affect the apparent aspect ratio of the surface. Von Collani (1985) provided clear evidence that perspective depth cues in photographs do in fact affect the strength of the horizontal–vertical illusion. When an 'L' shape was presented against a scene containing clear perspective cues to a receding surface (e.g. a road) a larger effect was found than when it was presented against a scene without variation in depth (a wall). One complication here, raised by Gregory (1998), is that in many situations in which the horizontal–vertical illusion is present, the stimulus does not appear to be slanted in depth. For example, Von Collani (1985) showed just as strong an illusion for the 'L' shape presented on its own, as when it was presented against a background with clear perspective depth cues. Only when a clearly frontoparallel wall was presented in the background did the effect diminish. Thus, while shape biases might arise from considerations of the perspective projection of slanted surfaces, this might be dissociated from the actual apparent slant of the surface. Gregory (1998) argued that, when stimuli are presented in the dark, the paradoxical conflict between apparent shape and apparent depth is removed. Gregory discussed the situation in which the stimuli are drawn with luminous paint and viewed in a dark room. Similar conditions can be created with the presentation of stimuli on a computer monitor. Informally he stated that, in these conditions, stimuli do actually appear to be slanted in depth, in a manner consistent with the interpretation of their shape and size.

Aspect ratio biases, when accounted for in terms of the expected depth structure of the environment, can be considered to reflect prior assumptions about the distributions of slants of surfaces. The goal of the current research is to determine whether aspect ratio biases can be accounted for in terms of such prior assumptions about slant, by formulating the perception of slant and shape as a problem of Bayesian estimation. An implicit assumption here is that, for a given aspect ratio in the image, apparent slant and apparent shape are mutually consistent (Koffka, 1935). Numerous studies in which both slant and shape have been measured have reported results that are inconsistent with this principle of shape-slant invariance (Beck & Gibson, 1955; Eby & Braunstein, 1995; Epstein, Hatfield, & Muise, 1977; Stavrianos, 1945). In contrast, Li and Durgin (2010) tested this assumption in a series of experiments in which observers both made verbal estimates of the slant of surfaces, and judged the relative distance between points, forming an 'L'-shape, in the saggital and frontal planes. Verbal estimates and the slant inferred from the judged aspect ratios were in good agreement, suggesting that biases in apparent aspect ratio reflected biases in the apparent slant of the surfaces. This study was restricted to an analysis of surfaces whose slant was within 24° of horizontal.

Other studies have demonstrated clear biases in the perception of slant, over a wider range of orientations. The particular pattern of bias in judgements of apparent slant appears to be inconsistent with the slant bias explanation of the horizontal–vertical illusion, when this is coupled with the idea of shape-slant invariance. For example, Durgin, Li, and Hajnal (2010) showed a vertical tendency in verbal judgements; the judged slant of surfaces within 15° of horizontal were drawn towards horizontal, while the slant of all other surfaces was drawn towards vertical. The effect of this tendency was that the greatest bias, of around 14°, was found for surfaces with a slant of 54° away from horizontal. This tendency is towards surfaces of vertical orientation, defined with respect to gravity, and was found to be invariant to changes in gaze direction; this is distinct from the idea of a frontal tendency, which would predict that surface slant would be biased towards an orientation that is gaze-normal (Gibson, 1950). Biases in apparent slant, with respect to a gravitationally-defined horizontal, are well modeled as a sine function of actual slant (Durgin & Li, 2012). Todd, Christensen, and Guckes (2010) also reported accurate perception of slant, for surfaces containing binocular depth cue whose orientation was close to frontoparallel.

The goal of the current study is determine the effects of changes in slant on the perception of surface aspect ratio, and whether aspect ratio biases can be explained in terms of apparent slant. We present observers with a task in which they are shown surfaces with a rectangular or elliptical outline in the retinal image, and asked to determine the aspect ratio of the surface relative to a square or a circle. That is, they are asked to consider whether the surface is taller or shorter than it is wide. The slant of the surfaces is defined by both binocular and texture cues. By measuring the point of subjective equality in this task, at which the surface appears to be square or circular, we are able to infer its apparent slant (Li & Durgin, 2010). This allows us to determine the extent to which slant cues affect the apparent shape of the surface. A given aspect ratio in the image is expected to appear as a surface that is progressively taller as the slant away from frontoparallel is increased.

This task also allows us to determine whether the horizontal–vertical illusion can be interpreted as a bias in apparent slant. When a surface, with a circular outline, is presented with a frontoparallel orientation, it is predicted that it will appear taller than it is wide. This bias could reflect a bias in the apparent slant of the surface, consistent with the statistical regularities that typify our environment. If so, then it would be expected that a surface with a circular outline in the retinal image would appear circular if binocular cues specified that the surface was slanted so that the top of the surface was closer than the bottom. This would counteract the bias attributed to prior assumptions about surface slant.

We present two experiments that address these questions. In the first experiment, we measure biases in the perception of aspect ratio, under conditions of uncertainty about shape. Under these conditions, we expect the effect of any prior to be relatively large. In the second experiment, we measure biases in apparent aspect ratio across a range of disparity-defined slants. These results are used to determine whether the aspect-ratio bias is affected by manipulations of slant in a manner consistent with the influence of prior assumptions on apparent slant along the lines proposed by Howe and Purves (2002).

2. Experiment one: Bias in the estimation of aspect ratio

Experiment one reports biases in the estimation of aspect ratio for simple rectangular stimuli. Rather than using figures with well-defined edges, we used sparse random dot stimuli in which the locations of the dots were confined to lie within a rectangular region. This increases uncertainty in the visual information provided about shape, and should thus increase the influence of any prior. The observer's task was to decide whether the rectangular region containing the dots was taller or shorter than it was wide. The use of sparse stimuli would be expected to considerably increase uncertainty in the shape of the object, and thus the effect of any priors on its estimation. Given that there has been some variability in results reporting a horizontal–vertical illusion for geometric shapes (McManus, 1978; Sleight & Austin, 1952) it is important to establish a robust effect for the kinds of stimuli used in this study.

2.1. Material and methods

2.1.1. Apparatus

Stimuli were generated, and the experiment controlled, using MATLAB and the Psychophysics toolbox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) running on a PC computer. Stimuli were presented on a 21-in. Sony Trinitron CRT monitor, viewed from a distance of 80 cm. Observers completed the experiment in a dark laboratory, with their head positioned in a chinrest to control the viewing distance.

2.1.2. Stimuli and procedure

Stimuli were presented against a black background on the monitor, and consisted of red, anti-aliased circular dots, presented within a rectangular region centered in the observer's visual field. The diameter of the dots was 5 arcmin. The background luminance of the screen was 0.14 cd m^{-2}, and the luminance of the dots was 21.2 cd m^{-2}. Dots were presented in a rectangular region of width 6 cm. The height of the region was varied from trial-to-trial. Seven aspect ratios were presented, ranging from 0.7 to 1.2. One hundred dots were presented randomly, following a uniform distribution, within the specified area. Each aspect ratio was presented ten times within a block of trials. Two blocks were run, giving a total of 20 repetitions of the seven aspect ratios. At the beginning of each block of trials, the observer was presented with a central fixation cross. When the observer pressed a response key, this was replaced by a stimulus, which was presented for 1 s. After this time, the stimulus disappeared, and was replaced by the fixation cross until the observer made a response. Observers responded in a two-alternative forced choice task by making a key press to indicate whether the rectangular region containing the dots was taller or shorter than it was wide. The next stimulus appeared once a response had been given.

It should be noted that these data were collected as part of a wider project to investigate the effect of motion adaptation on perceived shape (Hibbard et al., 2010). Observers also completed the above procedure after a period of motion adaptation. The data presented here are baseline conditions in observers that had not been presented with any adapting stimuli. The results are included as they represent a clear demonstration of the aspect ratio biases that are of interest.

2.2. Results and discussion

A typical psychometric function is shown in Fig. 1a. This shows the proportion of ''taller'' responses as a function of the actual aspect ratio of the rectangle. For each observer, we calculated the point of subjective equality, which indicates the physical aspect ratio corresponding to an apparently square rectangle. This was done by fitting a Weibull function to the data, using the psignifit toolbox (Wichmann & Hill, 2001a, 2001b). Ninety five percentage confidence intervals on the estimated PSEs were calculated using bootstrap analysis with 5000 repetitions. PSEs for individual observers are shown in Fig. 1b. Eleven of the fourteen observers showed a significant classical effect (perceptual elongation in the vertical direction relative to the horizontal direction). One observer showed a significant non-classical effect (perceptual elongation in the horizontal direction relative to the vertical direction). The remaining two observers showed no significant bias.

The mean aspect ratio of the apparently square rectangle was 0.92, with a standard deviation of 0.07. This is significantly different from 1 ($t(13) = 4.23$, $p < 0.001$). Thus, to appear square to the average observer, a rectangle had to be shorter than it was wide, in line with the standard horizontal–vertical illusion. This 8% error is considerably larger the 1.58% error reported by McManus (1978). We attribute this difference to the uncertainty in shape in our stimuli, which consisted of areas sparsely populated by randomly located dots. We calculated the effective slant, representing the slant of a surface that would project a retinal image with an aspect ratio equal to that of the average apparently square rectangle. This gives a value of 22.4°. That is, if the bias were to be attributed purely to a misperception of slant, then the results are consistent with the surface appearing to have a slant of 22.4°.

3. Experiment two: Aspect ratio judgements for slanted surfaces

The second experiment attempted to manipulate the apparent slant of the stimuli, using binocular disparity cues, to determine the effect of experimentally defined slant on aspect ratio judgements. If the biases found in experiment one arose because of the apparent slant of the stimulus, then a change in apparent slant should result in a change in bias. This means that we should be able to use cues to slant from binocular disparity to null the bias, such that the apparent slant of the surface is frontoparallel. At this point, an ellipse that had an aspect ratio of 1 in the image would appear circular. The purpose of this experiment was to test this directly. In so doing, we make use of measures of apparent aspect ratio, as a function of the geometrically specified slant of the surface, to infer the observer's prior assumptions about the distribution of slant. If aspect ratio biases are a direct result of the apparent slant of the surface, then it should be possible to use other cues, in this case binocular disparity, to counteract this bias. Alternatively, it is possible that the effective slant of the surface, from the point-of-view of biases in its apparent shape, might be dissociated from its perceived slant (Gregory, 1998). If biases cannot be nulled, then apparent shape and apparent slant must be, at least partially, dissociated.

3.1. Material and methods

3.1.1. Apparatus

Stimuli were generated, and the experiment controlled, using MATLAB and the Psychophysics toolbox running on a PC computer as before. Stimuli were presented on the same 21″ Sony Trinitron monitor as in experiment one, this time viewed from a distance of 40 or 100 cm. Experiments were performed in a fully-lit laboratory. Misperception of distance is known to be a significant factor in the misperception of depth from binocular disparity (Brenner & van Damme, 1999; Johnston, 1991). We wished to minimize any such effects, so that apparent distance was as accurate as possible. This should then provide the clearest assessment of the effects of slant on apparent shape. Binocular disparity was controlled by viewing the stimuli using CrystalEyes liquid–crystal goggles. Stimuli were presented in red, to minimize the cross-talk between the two eyes' views.

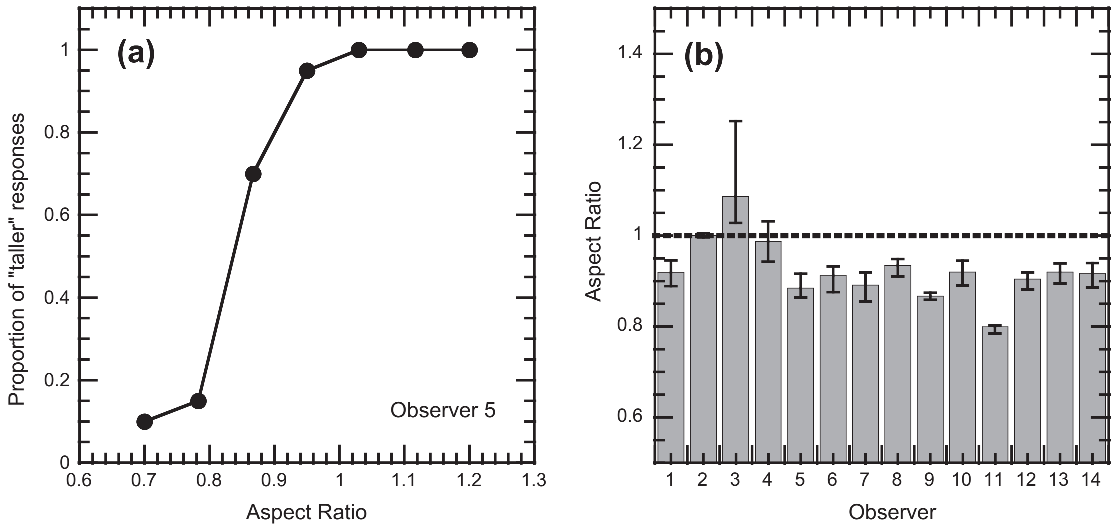

Fig. 1. (a) A typical psychometric function, showing the relationship between the physical aspect ratio of stimuli, and their judged aspect ratio. (b) The aspect ratio of the apparently squared rectangle for each of the 14 observers. Error bars show ±95% confidence limits.

3.1.2. Stimuli and procedure

Each stimulus was created by placing 1000 dots within a predefined elliptical area on the screen. The dots were distributed uniformly across the surface in 3D space. The height of the stimulus was fixed at 8° of visual angle. This corresponded to a height of 5.6 cm at the 40 cm viewing distance, and 14.0 cm at the 100 cm viewing distance. The width of the ellipse was under the observer's control. Moving the mouse leftwards reduced the width of the ellipse, moving it rightwards increased the width. The same dots remained on the screen, and were thus moved horizontally as the settings were made. The positions and sizes of the dots were determined by perspective projection. At the center of the surface, the size of the dots was 9.2 arcmin. The background luminance of the screen was 0.14 cd m^{-2}, and the luminance of the dots was 21.2 cd m^{-2}. The observer's task was to vary the aspect ratio until the surface appeared as a circular disc. When the observer was happy with their setting, they pressed a mouse button to remove the stimulus from the screen and record their setting. The next trial started when the observer pressed the spacebar on the keyboard.

The observer made 10 settings for 13 slants (0°, ±5°, ±10°, ±15°, ±30°, ±45°, ±60°) within a block of trials. Three blocks were run in total, to give 30 repetitions at each slant. Trial order was randomized with blocks.

3.2. Results and discussion

The mean aspect ratios set by each observer, as a function of the distance and the slant of the surface, are shown in Fig. 2. To emphasize the effects of the specified slant on aspect ratio settings, the data are presented as both the aspect ratio in the retinal image, and the aspect ratio of the three-dimensional surface. It should be emphasized that the observers' task was to consider the aspect ratio of the three-dimensional surface.

Dotted lines in Fig. 2a and b show the on-screen aspect ratios consistent with the projection of a circular surface, across the range of slants tested experimentally. The largest on-screen aspect ratio was usually set for frontoparallel surfaces. A significant bias was again observed in all cases: an ellipse that was shorter than it was wide was seen as circular. For all observers, at both distances, the set image aspect ratio for frontoparallel surfaces was significantly less than 1.0 (one-sample t-tests, in all cases $p < 10^{-5}$).

Smaller image aspect ratios were set for both positive and negative slants away from frontoparallel. This shows that observers are taking binocular information into account when making their judgements. Fig. 2c and d replots the data as the geometrically specified aspect ratio of the 3D surface. When the surface was close to frontoparallel, observers tended to overestimate the height of the surface, such that a surface that was shorter than it was wide appeared circular; for larger slants, the height tended to be underestimated.

These results are not consistent with the predictions of a simple Bayesian model of slant, as specified in Eq. (1), in which biases in the apparent aspect ratio of the surface result from biases in the apparent slant. There ought to be some value of the slant specified by disparity for which the apparent slant, once combined with the prior, is zero. At this point, the image aspect ratio would be set to 1. The fact that such a point was not found would lead us to predict that it was not possible to present a surface so that it appeared frontoparallel. For our viewing conditions, this explanation of the aspect ratio bias is very unlikely, since it would be inconsistent with the well-established vertical tendency discussed earlier (Durgin, Li, & Hajnal, 2010). This is because, if the apparent slant of surfaces is biased towards vertical, we would expect aspect ratio to be judged accurately for frontoparallel surfaces. This is shown in Fig. 2e and f, in which the data are replotted as the equivalent slant, i.e. the slant at which the set image aspect ratio would be projected by a surface that was circular. Note that the data point for 0° of slant is plotted twice, with both a positive and negative value. This is because these data are consistent with a surface that is slanted away from frontoparallel, but there is no way to determine the appropriate direction of effective slant from the data. This discontinuity when the data are plotted in this way suggests that biases in apparent slant cannot give a complete account of biases in aspect ratio. Nevertheless, it is clear that simulated slant was taken into account by the observers in making their setting, as is evident from the positive slope for both directions of simulated slant. The misperception of aspect ratio is best interpreted as an additional bias, over and above that which might be predicted from any plausible misperception of slant.

It should be noted that there were a number of unmodeled cues in our stimuli that might have biased the interpretation of the surface slant towards frontoparallel. These include focus cues, motion parallax and surface microtextures, and the presence of a frame

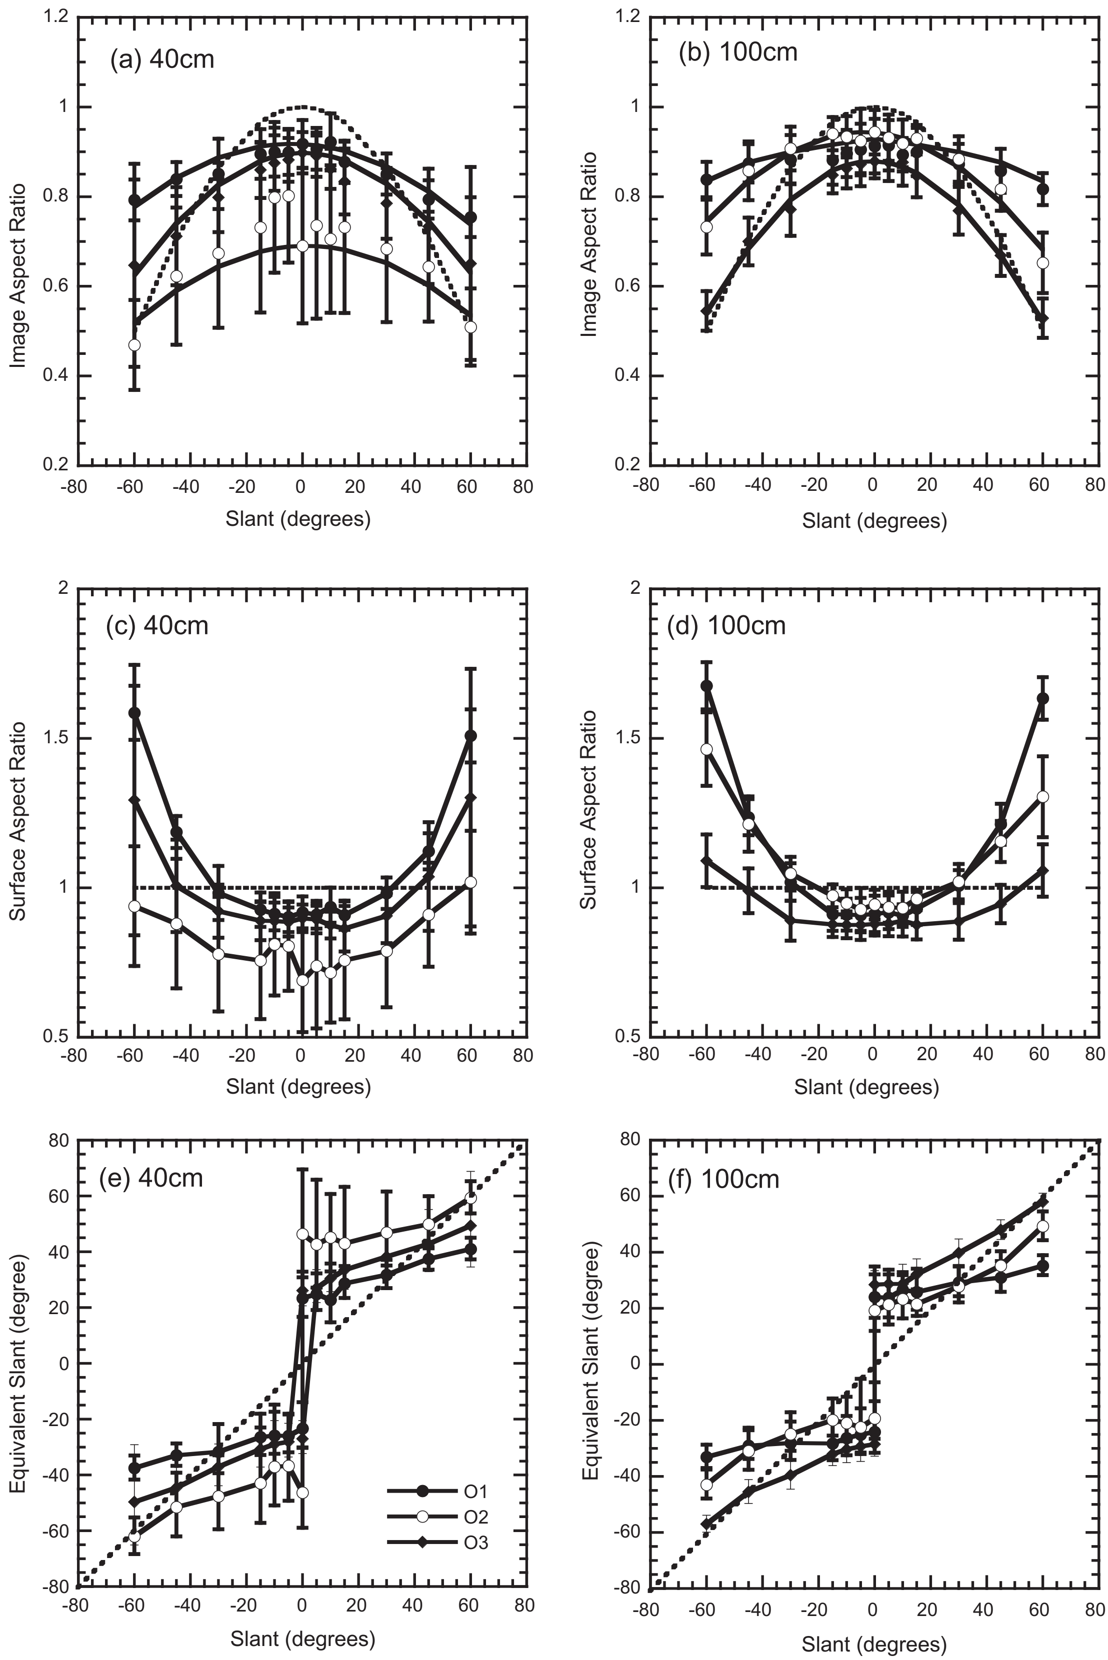

Fig. 2. Image aspect ratio settings made in experiment two. The image aspect ratio that appeared to be circular is plotted as a function of the slant of the surface for (a) the 40 cm and (b) the 100 cm distance. In each case, the plotted points are the mean of 30 settings, and the error bars show ±1 standard deviation. Dotted lines show veridical performance. The solid lines show the best fit of the description of the data shown in Eq. (4). Data are replotted as the aspect ratio of the surface in (c) and (d). (e) and (f) show the equivalent slant for each setting.

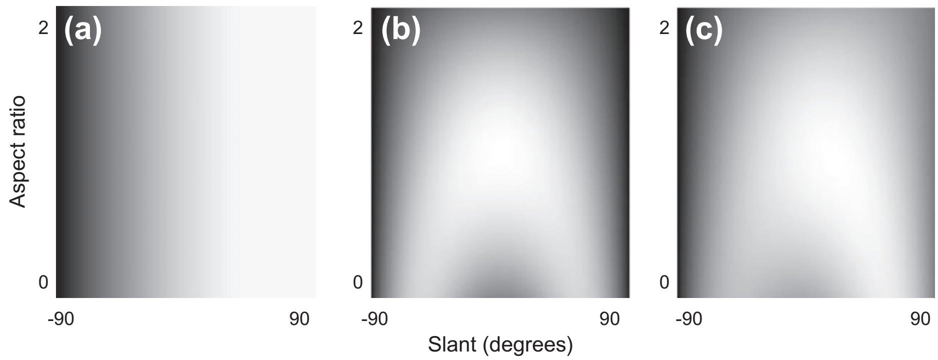

Fig. 3. (a) The prior probability distribution for the Bayesian model. The horizontal dimension represents the slant of the surface, the vertical dimension its aspect ratio. Higher probabilities are represented by brighter pixels. (b) and (c) show the likelihood and posterior probability functions in the same format. The parameters used to generate these distributions are given in the text. All plots show the log of probability, to emphasize the shapes of the distributions.

around the monitor screen (Eby & Braunstein, 1995). All of these could have contributed to a bias towards frontoparallel, which would bias the set retinal aspect ratio towards a value of 1 (Braunstein, 1976).

One possible Bayesian model of apparent slant and aspect ratio is outlined in Fig. 3. Fig. 3a shows the prior, which was created by assuming a uniform probability distribution over aspect ratio, and a normal distribution for slant. This embodies the idea that biases in apparent aspect ratio reflect expectations about the distribution of depth in the natural environment. The prior probability distribution is thus given by:

$$p(S,A) \propto \exp\left(\frac{(S-S_P)^2}{2\sigma_P^2}\right) \quad (2)$$

where S is slant, A is aspect ratio, and S_P and σ_P are the mean and standard deviation of the slant distribution. The peak is at 65°, with a standard deviation of 7.5°. These figures were used to approximate the distribution measured by Yang and Purves (2003). Fig. 3b shows the likelihood function. It was assumed that the estimation of slant from disparity, $\hat{S}$, and the aspect ratio of the projection of the surface onto the retinal image, $\hat{a}$, are unbiased and subject to Gaussian error. The likelihood function is then given by:

$$p(\hat{S},\hat{a}|S,A) \propto \exp\left(-\frac{(\hat{S}-S)^2}{2\sigma_S^2} - \frac{(\hat{a}-A\cos S)^2}{2\sigma_A^2}\right) \quad (3)$$

where σ_s and σ_A are the standard deviations for slant and aspect ratio and S and A are the true values of slant and aspect ratio of the surface. In Fig. 3b, values of $\sigma_A = 5\%$ and $\sigma_s = 5°$ were used, and the surface was frontoparallel with an aspect ratio of 1. The posterior probability distribution, shown in Fig. 3c, is then obtained by multiplying together the prior and likelihood, according to Eq. (1). Note that the peak in the posterior probability distribution is at a positive value of slant when the surface is frontoparallel, reflecting the fact that the peak in the prior is at a positive value.

This model estimates the shape and slant of the surface simultaneously, given the shape of the surface in the retinal image, and information about slant; this means that any biases in apparent slant, caused by a non-uniform prior for slant, would have a direct effect on apparent shape. To simulate the results of our psychophysical experiments, we estimated the image aspect ratio that would result in an apparent surface aspect ratio of 1, as a function of both the slant of the surface, and the reliability of slant information. σ_s was varied between 2.5° and 7.5°. These estimates were obtained by varying the image aspect ratio in steps of 0.01, and for each value estimating the surface aspect ratio using a maximum likelihood criterion. The image aspect ratio giving a surface aspect ratio closest to one was then selected. These results are plotted in Fig. 4. For each level of reliability, there is a slant at which an image aspect ratio of one corresponds to a surface aspect ratio of 1. This is the point at which the surface appears frontoparallel. As the reliability of slant information decreases, this slant moves progressively further away from frontoparallel, reflecting an increased influence of the prior. It is important to note that the largest image aspect ratio seen as a circular surface would be 1 in all cases. It is not therefore possible to account for the entirety of the bias in aspect ratio judgments using a Bayesian model of this type, in which biases area attributed to prior assumptions about the distribution of slant, since we typically found no value of slant at which an image aspect ratio of 1 appeared as a surface with an aspect ratio of 1.

The notable features of our data are that there is an overall aspect ratio bias, and that the visually specified slant of the surface affected the set aspect ratio. This pattern of responses is captured by the following equation:

$$\hat{A} = k\cos(gS) \quad (4)$$

Here, k is a constant that captures the overall aspect ratio bias, and g reflects the gain with which the slant information is used. The best fit of this equation is shown by the solid lines in Fig. 2a and b. Note

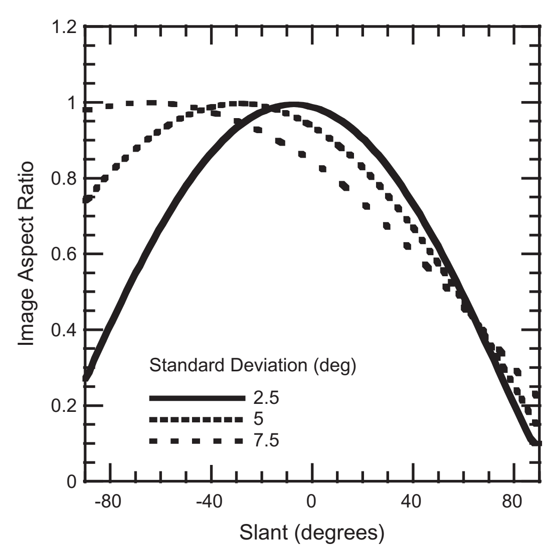

Fig. 4. The image aspect ratio that would appear as an aspect ratio of 1, as estimated by our model, as a function of slant. Separate lines show model results for difference values of the reliability of the slant information provided by binocular disparity.

that this equation does not provide an explanation of the data, but a description. This might however represent a more heuristic use of statistical information about slant than that expressed in the Bayesian model (Braunstein, 1976; Caudek & Proffitt, 1993). The tendency for surfaces to be slanted away from frontoparallel leads to an overall bias in the estimation of aspect ratio, which is captured by the constant $k < 1$; this bias does not however interact directly with perceptual information about slant. Across both distances, and all observers, average values of $k = 0.87$ and $g = 0.66$ were obtained. The same average estimate of the magnitude of the horizontal–vertical bias in our data was obtained when Durgin and Li's (2012) sine model of apparent slant, rather than a linear model, was used.

4. Conclusions

We observed a considerable bias in the perception of the aspect ratio of rectangular and elliptical shapes. Such biases have been reported previously (Fick, 1851 (cited by Künnapas (1955)); McManus, 1978; Sleight & Austin, 1952), and are examples of a general class of effects related to the horizontal–vertical illusion. A compelling, theoretically motivated explanation of such effects attributes them to a depth bias: vertical extents in the retinal image are likely to correspond to greater 3D separations in the environment than are horizontal extents of the same magnitude. This difference reflects the fact that vertically separated points are likely, on average, to have a greater separation in depth. This difference can be attributed to the ubiquity of the ground-plane in our environment, and has been confirmed empirically (Howe & Purves, 2002).

When applied to the perception of shape, this bias amounts to an implicit slant of the viewed shape around a horizontal axis. However, this bias does not appear to interact with binocular cues to the slant of the surface in a straightforward way, as might be expected from, for example, a Bayesian approach to estimating shape. Specifically, while observers' aspect ratio settings took the slant of depicted surfaces into account, we were able to find no binocular-disparity defined slant at which the effective slant away from the frontoparallel plane was zero, such that an ellipse with an aspect ratio of 1 in the image appeared circular. At all slants tested, the surface that appeared to have a circular outline was shorter than it was tall in the image.

This overall pattern of bias is difficult to account for using either Bayesian approaches, or alternatives such as empirical ranking theory (Howe, Lotto, & Purves, 2006), if biases in apparent shape are to be explained in terms of biases in apparent slant. It is possible to frame the problem as one of shape estimation. Prior information about the distribution of shapes of objects in the environment, and how these are projected onto the retinal image, can then be incorporated as a Bayesian prior and likelihood, or as a probability distribution for the horizontal and vertical extents of retinal projections. Either of these approaches can be used to provide an account of the bias with reference to natural-scene statistics. A difficulty arises however when attempting to incorporate the role played by binocular cues to surface slant. Such information does affect apparent shape, and could readily be included in Bayesian or empirical ranking models. However, if shape biases ultimately arise from the expected slant of surfaces about a horizontal axis, no bias is predicted when binocular cues clearly indicate that the surface is frontoparallel. This is not the pattern of results observed here. Nor is it consistent with the finding that horizontal–vertical illusion is smaller under monocular viewing (Prinzmetal & Gettleman, 1993). Although biases in apparent shape might reflect knowledge of the statistical distribution of depth in the natural environment, they do not appear to stem from biases in apparent slant.

One reason we might expect a dissociation between the apparent slant of a surface, and the horizontal–vertical illusion, is if the latter reflects a slant bias in pictorial depth (Koenderink, 1998). It would be entirely possible to maintain separate estimates of the slant of the picture surface, and the slant of objects within the picture. A horizontal–vertical illusion could then be expected even in the presence of reliable and unbiased information about the picture surface. This argument cannot readily explain the biases observed in our second experiment. This was performed in a fully lit laboratory, so that the orientation of the monitor surface was clearly visible; binocular disparity influenced the apparent slant of the surface, and subsequently the apparent aspect ratio of the disc. There was thus a clear dissociation between any pictorial plane, and the apparent slant of the disc.

Biases in the perception of 3D shape, over a wide range of viewing conditions and for stimuli containing binocular depth cues, have been reported many times in the literature. Typically, objects that are relatively close to the observer tend to appear stretched in depth, relative to the horizontal and vertical dimensions, while those that are further way tend to appear relatively squashed in the depth dimension (Bradshaw, Glennerster, & Rogers, 1996; Bradshaw, Parton, & Eagle, 1998; Bradshaw, Parton, & Glennerter, 2000; Brenner & Landy, 1999; Brenner & van Damme, 1999; Collett, Schwarz, & Sobell, 1991; Domini, Caudek, & Tassinari, 2006; Glennerster, Rogers, & Bradshaw, 1996; Johnston, 1991; Liter, Braunstein, & Hoffman, 1994; O'Kane & Hibbard, 2010; Rogers & Bradshaw, 1993, 1995; Scarfe & Hibbard, 2006, 2011; Tassinari, Domini, & Caudek, 2008; Tittle et al., 1995; Todd & Bressan, 1990; Todd & Norman, 1991, 2003). Such biases have been described in terms of the geometries of visual space (Tittle et al., 1995), and the relationship between perceptual and physical space (Wagner, 1985). Such effects might have been expected to cause our stimuli to appear stretched in depth, leading observers to set smaller aspect ratios in our second experiment. However, any such effects would have been expected to be relatively modest at the 1 m viewing distance, at which depth scaling tends to be relatively accurate (Johnston, 1991). Moreover, a slant at which the surface appeared frontoparallel, leading observers to set an image aspect ratio of 1, would still be expected.

It is possible that, rather than any single representation accounting for our perception of 3D shape, multiple representations of important surface properties are independently determined. For example, it has been suggested that independent representations of depth, 3D orientation, curvedness and shape index are available (Koenderink, 1998; Norman et al., 2006; Tittle & Perotti, 1997). Evidence for these multiple representations comes from the fact that performance on tasks dependent on higher-level properties (e.g. shape index, or curvature) cannot always readily be accounted for on the basis of the information available in lower-level properties such as slant (Johnston & Passmore, 1994; Tittle & Perotti, 1997).

It has been argued that a lack of consistency between performance on different perceptual tasks is to be expected (Koenderink, 1998; Tcheang, Gilson, & Glennerster, 2005). One such view that is particularly at odds with the type of model under consideration here is the utilitarian view that the visual system might use a broad range of "tricks and rules of thumb" in order to solve specific tasks (Hibbard, 2008; Ramachandran, 1985). This view would however readily account for the apparent discrepancy between slant and aspect ratio found here; the aspect ratio bias evident in the horizontal–vertical illusion represents a well-founded rule of thumb applied to the estimation of shape (Howe & Purves, 2002), but is not used in a way that ensures a representation of a surface in which all its perceived geometrical properties are necessarily mutually consistent (Koenderink, 1998).

Acknowledgment

This work was funded by BBSRC Grant BB/C005260/1.

References

Avery, G. C., & Day, R. H. (1969). Basis of the horizontal vertical illusion. *Journal of Experimental Psychology, 81*, 376–380.
Bayer, C. A., & Pressey, A. W. (1972). Geometric illusions as a function of pigmentation of the fundus oculi and target size. *Psychometric Science, 26*, 77–79.
Beck, J., & Gibson, J. J. (1955). The relation of apparent shape to apparent slant in the perception of objects. *Journal of Experimental Psychology, 50*, 125–133.
Begelman, D. A., & Steinfeld, J. (1971). Retinal vs. phenomenal determinants of the horizontal–vertical illusion. *Journal of General Psychology, 85*, 77–85.
Bradshaw, M. F., Glennerster, A., & Rogers, B. J. (1996). The effect of display size on disparity scaling from differential perspective and vergence cues. *Vision Research, 36*, 1255–1264.
Bradshaw, M. F., Parton, A. D., & Eagle, R. A. (1998). The interaction of binocular disparity and motion parallax in determining perceived depth and perceived size. *Perception, 27*, 1317–1333.
Bradshaw, M. F., Parton, A. D., & Glennerter, A. (2000). The task-dependent use of binocular disparity and motion parallax information. *Vision Research, 40*, 3725–3734.
Brainard, D. H. (1997). The Psychophysics toolbox. *Spatial Vision, 10*, 433–466.
Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America, A, 14*, 1393–1411.
Braunstein, M. L. (1976). *Depth perception through motion*. New York: Academic Press.
Brenner, M. F., & Landy, M. S. (1999). Interaction between the perceived shape of two objects. *Vision Research, 39*, 3834–3848.
Brenner, E., & van Damme, W. J. M. (1999). Perceived distance, shape and size. *Vision Research, 39*, 975–986.
Burge, J., Fowlkes, C. C., & Banks, M. S. (2010). Natural-scene statistics predict how the figure-ground cue of convexity affects human depth perception. *The Journal of Neuroscience, 30*, 7269–7280.
Caudek, C., & Proffitt, D. R. (1993). Depth perception in motion parallax and stereokinesis. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 32–47.
Collett, T. S., Schwarz, U., & Sobell, E. (1991). The interaction of oculomotor cues and stimulus size in stereoscopic depth constancy. *Perception, 20*, 733–754.
Cormack, E. O., & Cormack, R. H. (1974). Stimulus configuration and line orientation in the horizontal–vertical illusion. *Perception and Psychophysics, 16*, 208–212.
Craven, B. J. (1993). Orientation dependence of human line-length judgments matches statistical structure in real-world scenes. *Proceedings of the Royal Society of London B, 253*, 101–106.
Cutting, J. E. (1997). How the eye measures reality and virtual reality. *Behavior Research Methods, Instruments and Computers, 29*, 27–36.
Domini, F., Caudek, C., & Tassinari, H. (2006). Stereo and motion information are not independently processed by the visual system. *Vision Research, 46*, 1707–1723.
Durgin, F. H., & Li, Z. (2012). Spatial biases and the haptic experience of surface orientation, haptics rendering and applications. In A El Saddik (Ed.), *Tech*. <http://www.intechopen.com/books/haptics-rendering-and-applications/spatial-biases-and-the-haptic-experience-of-surface-orientation>.
Durgin, F. H., Li, Z., & Hajnal, A. (2010). Slant perception in near space is categorically biased: Evidence for a vertical tendency. *Attention, Perception and Psychophysics, 72*, 1875–1889.
Eby, D. W., & Braunstein, M. L. (1995). The perceptual flattening of three-dimensional scenes enclosed by a frame. *Perception, 24*, 981–993.
Epstein, W., Hatfield, G., & Muise, G. (1977). Perceived shape at a slant as a function of processing time and processing load. *Journal of Experimental Psychology: Human Perception and Performance, 3*, 473–483.
Gibson, J. J. (1950). The perception of visual surfaces. *American Journal of Psychology, 63*, 367–384.
Girchick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience, 14*, 926–934.
Glennerster, A., Rogers, B. J., & Bradshaw, M. F. (1996). Stereoscopic depth constancy depends on the subject's task. *Vision Research, 36*, 3441–3456.
Gregory, R. L. (1963). Distortion of visual space as an inappropriate constancy scaling. *Nature, 199*, 678–680.
Gregory, R. L. (1973). A discussion of G.H. Fisher's "Toward a new explanation for the geometrical illusions: Apparent depth of contour proximity?" and the inappropriate constancy-scaling theory. *British Journal of Psychology, 64*, 623–626.
Gregory, R. L. (1998). *Eye and brain: The psychology of seeing*. Oxford: Oxford University Press.
Hibbard, P. B. (2008). Can appearance be so deceptive? Representationalism and binocular vision. *Spatial Vision, 21*, 549–559.
Hibbard, P., Scarfe, P., Robertson, M., & Windeatt, S. (2010). Motion adaptation affects perceived shape. *Journal of Vision, 10*(7), 31.
Howe, C. Q., Lotto, R. B., & Purves, D. (2006). Comparison of Bayesian and empirical ranking approaches to visual perception. *Journal of Theoretical Biology, 241*, 866–875.
Howe, C. Q., & Purves, D. (2002). Range image statistics can explain the anomalous perception of length. *Proceedings of the National Academy of Sciences of the United States of America, 99*, 13184–13188.
Johnston, E. B. (1991). Systematic distortions of shape from stereopsis. *Vision Research, 33*, 1351–1360.
Johnston, A., & Passmore, P. J. (1994). Independent encoding of surface orientation and curvature. *Vision Research, 34*, 3005–3012.
Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception, 36S*, 14.
Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. Cambridge: Cambridge University Press.
Koenderink, J. J. (1998). Pictorial relief. *Philosophical Transactions of the Royal Society of London A, 356*, 1071–1086.
Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt Brace.
Künnapas, T. M. (1955). An analysis of the vertical–horizontal illusion. *Journal of Experimental Psychology, 49*, 134–140.
Li, Z., & Durgin, F. H. (2010) Perceived slant of binocularly viewed large-scale surfaces: A common model from explicit and implicit measures. *Journal of Vision, 10*(13), 13, 1–15.
Liter, J. C., Braunstein, M. L., & Hoffman, D. D. (1994). Inferring structure from motion in two-view and multiple displays. *Perception, 22*, 1441–1465.
McManus, I. C. (1978). The horizontal–vertical illusion and the square. *British Journal of Psychology, 69*, 369–370.
Norman, J. F., Todd, J. T., Norman, H. F., Clayton, A. M., & McBride, T. R. (2006). Visual discrimination of local surface structure: Slant, tilt and curvedness. *Vision Research, 46*, 1057–1069.
O'Kane, L. M., & Hibbard, P. B. (2010). Contextual effect on perceived three-dimensional shape. *Vision Research, 50*, 1095–1100.
Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.
Pollock, W. T., & Chapanis, A. (1952). The apparent length of a line as a function of its inclination. *Quarterly Journal of Experimental Psychology, 4*, 170–178.
Prinzmetal, W., & Gettleman, L. (1993). Vertical–horizontal illusion: One eye is better than two. *Percept ion and Psychophysics, 53*, 81–88.
Ramachandran, V. S. (1985). The neurobiology of perception. *Perception, 14*, 97–103.
Rogers, B. J., & Bradshaw, M. F. (1993). Vertical disparities, differential perspective and binocular stereopsis. *Nature, 361*, 253–255.
Rogers, B. J., & Bradshaw, M. F. (1995). Disparity scaling and the perception of frontoparallel surfaces. *Perception, 24*, 155–180.
Scarfe, P., & Hibbard, P. B. (2011). Statistically optimal integration of biased sensory estimates. *Journal of Vision, 11*(7), 12, 1–17.
Scarfe, P., & Hibbard, P. B. (2006). Disparity-defined objects moving in depth do not elicit three-dimensional shape constancy. *Vision Research, 46*, 1599–1610.
Schiffma, H. R., & Thompson, J. (1974). The role of eye movements in the perception of the horizontal–vertical illusion. *Perception, 3*, 49–52.
Shipley, W. C., Nann, B. M., & Penfield, M. J. (1949). The apparent length of tilted lines. *Journal of Experimental Psychology: General, 39*, 548–551.
Sleight, R. B., & Austin, T. R. (1952). The horizontal–vertical illusion in plane geometric figures. *Journal of Psychology, 33*, 279–287.
Stacey, B. G. (1969). Explanations of the horizontal–vertical illusion and the foreshortening of the receding horizontal. *Life Sciences, 8*, 1237–1246.
Stavrianos, B. K. (1945). The relation of shape perception to explicit judgements of inclination. *Archives of Psychology, 62*, 223–233.
Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience, 9*, 578–585.
Tassinari, H., Domini, F., & Caudek, C. (2008). The intrinsic constraint model for stereo-motion integration. *Perception, 37*, 79–95.
Tcheang, L., Gilson, S. J., & Glennerster, A. (2005). Systematic distortions of perceptual stability investigated using immersive virtual reality. *Vision Research, 45*, 2177–2189.
Thompson, J., & Shiffman, H. R. (1974). The effect on the magnitude of the horizontal–vertical illusion of horizontal retinal eccentricity. *Vision Research, 14*, 1463–1465.
Thouless, R. H. (1931a). Phenomenal regression to the real object. I. *British Journal of Psychology, General Section, 21*(4), 339–359.
Thouless, R. H. (1931b). Phenomenal regression to the real object. II. *British Journal of Psychology, General Section, 22*(1), 1–30.
Tittle, J. S., & Perotti, V. J. (1997). The perception of shape and curvedness from binocular stereopsis and structure from motion. *Perception and Psychophysics, 59*, 1167–1179.
Tittle, J. S., Todd, J. T., Perotti, V. J., & Norman, J. F. (1995). Systematic distortion of perceived three-dimensional structure from motion and binocular stereopsis. *Journal of Experimental Psychology, Human Perception and Performance, 21*, 663–678.
Todd, J. T., & Bressan, P. (1990). The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Percept and Psychophysics, 48*, 419–430.
Todd, J. T., Christensen, J. C., & Guckes, K. M. (2010). Are discrimination thresholds a valid measure of variance for judgments of slant from texture? *Journal of Vision, 10*(2), 20, 1–18.
Todd, J. T., & Norman, J. F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Percept and Psychophysics, 50*, 509–523.
Todd, J. T., & Norman, J. F. (2003). The visual perception of 3-d shape from multiple cues: Are observers capable of perceiving metric structure? *Percept and Psychophysics, 65*, 31–47.

Valentine, C. W. (1912). The effect of astigmatism on the horizontal vertical illusion, and a suggested theory of the illusion. *Journal of Psychology, 5*, 308–330.
Von Collani, G. (1985). The horizontal–vertical illusion in photographs of concrete scenes with and without depth information. *Perceptual and Motor Skills, 61*, 523–531.
Wagner, M. (1985). The metric of visual space. *Perception and Psychophysics, 38*, 483–495.
Webster, M. A. (2011). Adaptation and visual coding. *Journal of Vision, 11*(5), 3, 1–23.
Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience, 5*, 598–604.
Wichmann, F. A., & Hill, N. J. (2001a). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception and Psychophysics, 63*, 1293–1313.
Wichmann, F. A., & Hill, N. J. (2001b). The psychometric function II. Bootstrap-based confidence intervals and sampling. *Perception and Psychophysics, 63*, 1214–1329.
Wolfe, U., Maloney, L. T., & Tam, M. (2005). Distortions of perceived length in the frontoparallel plane: Tests of perspective theories. *Perception and Psychophysics, 67*, 967–979.
Woodworth, R. S. (1938). *Experimental psychology*. New York: Holt.
Yang, Z., & Purves, D. (2003). A statistical explanation of visual space. *Nature Neuroscience, 6*, 632–640.