Distortions of perceived length in the frontoparallel plane: Tests of perspective theories

UTA WOLFE and LAURENCE T. MALONEY
New York University, New York, New York

and

MIMI TAM
New York University, New York, New York

The perceived length of a line segment in a frontoparallel plane is sometimes affected by the presence of other line segments in the visual field. Perspective theories attribute such interactions to size-constancy scaling: The configuration of line segments present in the visual field includes depth cues that trigger size scaling of each line segment. In three experiments, we test this claim for a range of simple configurations composed of two line segments joined at a point. These configurations include the inverted T configuration of the bisection illusion, as well as the L configuration of the horizontal–vertical illusion. We conclude that the available depth cues, even when supplemented by known biases in perspective interpretations, do not account for observed distortions in judgments of relative length.**

The importance of visual illusions lies in what they tell us about visual perception (Coren & Girgus, 1978; Gregory, 1970, 1997). In this article, we will be concerned with biases in the visual estimation of the length of line segments presented in a frontoparallel plane. We are interested in testing theories that are intended to predict such biases, and in doing so, we will explore a wider range of stimuli than just the traditional illusion configurations. Each of the stimuli will consist of two line segments joined at a point. We refer to these stimuli as *two-line configurations*, examples of which are shown in Figures 1A and 1B.

The two configurations in Figure 1 have been extensively studied: When the vertical and horizontal lines are of the same length, observers typically judge the vertical line in each configuration to be longer. Although there is some inconsistency in terminology, the corresponding illusions are known as the *bisection illusion* (Figure 1A) and the *horizontal–vertical illusion* (Figure 1B), respectively (Coren & Girgus, 1978, pp. 27–29). The magnitudes of the two illusions vary with the viewing conditions, but the former illusion is usually larger than the latter.

Several authors have proposed that the horizontal–vertical illusion is a consequence of a fixed distortion in the visual field. The basis for the distortion has been sought in (1) imperfections in the refractive properties of the eye (Avery & Day, 1969; Thompson & Schiffman, 1974; Valentine, 1912), (2) nonhomogeneous photoreceptor spacing in the retina (Begelman & Steinfeld, 1971), (3) nonuniformities in retinal pigment distribution (Bayer & Pressey, 1972), and (4) the oval shape of the visual field (Künnapas, 1955). Fixed distortions of length (separation) are found even when the visual field contains only two points and the comparison of length is against a memory standard (McGraw & Whittaker, 1999).[1] Williams and Enns (1996) have argued that at least two independent factors contribute to the horizontal–vertical illusion.

Such distortions likely do affect judgments of length. However, they cannot explain visual illusions of length where the length of a given line segment in the visual field is affected by the presence or absence of other line segments. If, for example, the configuration of Figure 1A is rotated 90°, the reader will likely still see the previously vertical segment as the longer one.[2] No fixed spatial distortion could account for this pattern of results. A common explanation of the bisection illusion (which gives it its name) is that the presence of the bisection point where the lines meet leads to an apparent shortening of the bisected (horizontal) line (Oppel, 1855). These sorts of interactions are well established for judgments of collinearity and angle (e.g., Greene, 1998).

Even if we consider only a tiny subset of possible "visual field contents," the number of possible interactions is large, posing a severe obstacle to characterizing them through experimental measurement. Oppel (1855), for example, studied simple configurations of collinear items and measured how adding and subtracting items affected the apparent separation of the remaining items. Adding a single bisecting point between two other points
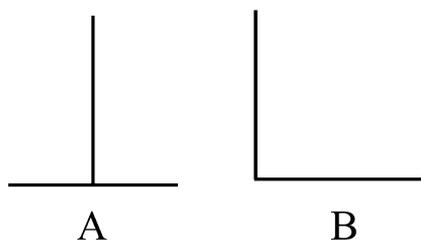
Figure 1. Two visual illusions of length: (A) the bisection illusion and (B) the horizontal–vertical illusion. In each, the vertical line appears to be longer than the horizontal line. For most observers, the illusion magnitude is larger for the configuration in panel A than for that in panel B.

leads to a decrease in their apparent separation—the bisection illusion mentioned previously. But adding more than one intermediate point sometimes had the opposite effect: The apparent separation of the endpoints increased. There is no quantitative predictive theory of even the interactions among collinear items in the visual field.

One puzzling aspect of observed interactions is that the visual system seems to be working very hard to "get it wrong." Of course, the experimenter has no right to demand that the visual system "justify" its operation, but it would nonetheless be satisfying if an interaction theory would not only predict observed interactions in judgments of length but also offer some insight into why such interactions occur at all.

Two closely related interaction theories, both intended to explain observed distortions of length, also provide insight into why we get it wrong. They are the perspective theory of Woodworth (1938) and the *misapplied size-constancy scaling* theory of Gregory (1963, 1973). We refer to these theories collectively as *perspective theories*[3] (terminology due to Gregory, 1997). For convenience in describing them, we will confine our attention to visual

scenes ("configurations") comprising a few line segments confined to a plane perpendicular to the line of sight.

Woodworth (1938) and Gregory both assumed that a line segment enters into visual processing by two routes: first, as an isolated line segment that must be assigned a length and location within three-dimensional (3-D) space, and second, as part of a configuration that triggers a depth interpretation leading to size-constancy scaling. The flow of information in their theories is diagrammed in Figure 2. The length and location of any single line is affected by the presence or absence of other line segments, but *only* because the latter alters the depth interpretation and the size scaling applied to the former. The upper pathway in Figure 2, which selects the depth interpretation, is the distinctive element shared by the perspective theories.

Woodworth (1938) proposed that observers consciously perceive the stimuli, which are confined to the frontoparallel plane, as extended in depth, out of the frontoparallel plane. A consequence of his perspective theory is that to the extent that the observer perceives the vertical line in Figures 1A and 1B as longer than the horizontal line, he or she must also perceive at least one of the lines as displaced in depth from its true location. The key idea of Woodworth's theory is that the observer is simply interpreting the configurations of Figures 1A and 1B as line drawings of 3-D scenes. The illusion results from confusion between two possible interpretations of a picture (Sedgwick & Nicholls, 1993; Sedgwick, Nicholls, & Brehaut, 1995). Indeed, in a complex line drawing with linear perspective cues, it would seem to be an error to describe the resulting distortions of length as illusions rather than as examples of veridical picture interpretation. For the simple configurations of Figures 1A and 1B, and related illusions involving only a handful of line segments, it is less obvious that the illusions of length are simply consequences of picture interpretation.
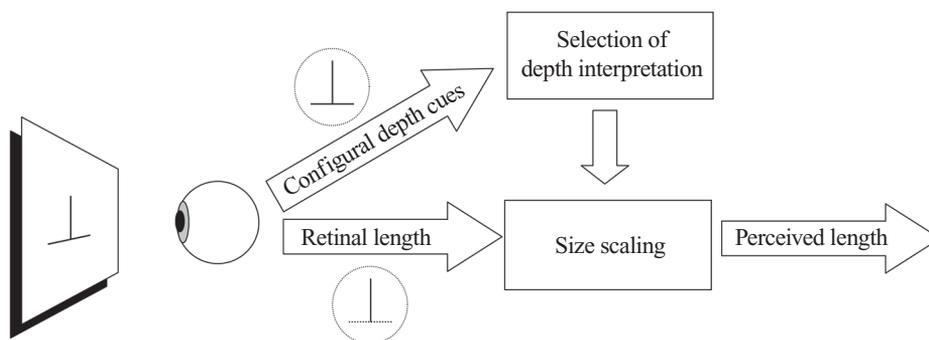

Figure 2. Information flow in perspective theories. The stimulus configuration enters into visual processing by two distinct routes. The diagram illustrates the estimation of the length of the vertical line segment in the inverted T configuration (Figure 1A). The overall configuration selects a depth interpretation that controls size scaling of each line segment (upper pathway). The retinal length of each segment is scaled according to the depth interpretation (lower pathway), and the resulting scaled length is the length perceived by the observer. In Gregory's perspective theory, the depth interpretation that controls size-constancy scaling need not be the same as the depth interpretation (not shown) that is consciously perceived.

The misapplied size-constancy scaling theory of Gregory (1963, 1970, 1973, 1997; see also Girgus & Coren, 1975) is similar in many respects to Woodworth's theory. Size scaling is responsible for the visual illusion by the same mechanism that Woodworth proposed, but in Gregory's version this scaling need *not* be accompanied by conscious perception of depth signaled by the cues that led to size scaling. Gregory's hypothesis is remarkable in that it explicitly allows for the possibility that the visual system arrives at an inconsistent state where the perceived lengths of line segments, their perceived locations in 3-D space, and the retinal images correspond to no single physical scene. The consciously perceived depth and the depth that triggers conscious size estimates through constancy mechanisms are not the same; the visual system has "dissociated." Milner and Goodale (1995) described how judgment of size based on conscious estimates of scene attributes can deviate from estimates inferred from visually guided motor responses and argued that no single underlying visual representation can account for both. Gregory's remarkable theory, in effect, proposes that such inconsistencies can be found in visual processing alone even if visually guided motor responses are neglected. This cannot happen in Woodworth's perspective theory. Such dissociations may be common in scenes that are too simple to have a pictorial depth interpretation but that nevertheless contain pictorial depth cues. They may contribute to many visual illusions.

Both perspective theories assume that (1) perceived length interactions that occur when an item is added or subtracted from the visual field can only be due to changes in depth interpretation, and (2) these depth interpretations are triggered by aspects of the configuration that normally affect perceived depth via the upper pathway in Figure 2.

Neither Woodworth nor Gregory has proposed an explicit theory of how the depth interpretation is, in general, selected. If perspective theories are to count as explanations of the distortions in perceived length in Figures 1A, 1B, and other two-line configurations, they must be complemented by a specification of the mechanism that selects the depth interpretation. To achieve this end, we must first analyze the depth cues that are present in such simple configurations and that hypothetically trigger depth interpretations via the upper pathway in Figure 2.

Figures 1A and 1B, viewed binocularly, share many traditional depth cues: binocular disparity, motion parallax (if the observer is moving), and elevation of the vertical line relative to the horizontal, among others (Levine, 2000). The configuration of Figure 1A contains an occlusion cue not present in Figure 1B: the inverted T junction where the vertical line meets the horizontal. In the first experiment, we test whether the simple presence or absence of this cue is responsible for the difference in illusion magnitude in the two configurations of Figure 1.

The bisection illusion is often described as the combination of two illusions, one the horizontal–vertical, and the other connected to the collinearity effects described

by Oppel (1855; see also Coren & Girgus, 1978). This way of speaking is imprecise, since one phenomenon can never serve to explain another. What is almost certainly intended is that the mechanism underlying the distortion observed in the horizontal–vertical illusion and that underlying the collinearity effects described by Oppel are both at work in the configuration of Figure 1A. The nature of these mechanisms is currently a matter of conjecture. In Experiment 1, we examine whether a size-constancy scaling mechanism triggered by the presence or absence of an evident depth cue can account for the differences between Figures 1A and 1B.

EXPERIMENT 1

To recapitulate, one depth cue that is present in Figure 1A but not in Figure 1B is the apparent occlusion of the vertical line by the horizontal line.[4] In this experiment, we test whether the presence or absence of the occlusion cue (the inverted T junction) accounts for the difference in the magnitude of illusion for 11 two-line configurations, 8 of which are intermediate between the configurations of Figure 1A (the bisection illusion) and Figure 1B (the horizontal–vertical illusion).[5]

Cues such as the inverted T junction have been extensively studied in the computational vision literature. It is an example of the sort of simple depth primitive used in computational models of line drawing analysis associated with "blocks world" (see Cohen & Feigenbaum, 1982; Huffman, 1971). The information contained in such primitives often proved to be sufficient to reconstruct simple 3-D line drawings of polyhedra, indicating that the cue does carry useful information.

Method

Observers. Seven observers (5 female and 2 male) were paid to participate in the experiment. All had normal or corrected-to-normal vision, were naive about the purpose of the experiment, and required 2–2.5 h to complete eight blocks with rest breaks.

Stimuli and Apparatus. The stimuli were displayed on a Sony GDM-G500 21-in. monitor running at $1,280 \times 1,024$ resolution under the control of a Matrox G450 graphics card and a Dell 410 Workstation running Red Hat 6.1 Linux software. The sequence and time of stimulus displays were controlled by a special-purpose program written in C by one of the authors, using the X11R6 Windows graphical interface and a special-purpose graphics driver from Xi Graphics. The monitor employed had a display area that was very close to flat. Only a small central region was used in presenting stimuli. The monitor was adjusted so that vertical distances measured in pixels and horizontal distances measured in pixels were identical and 100 pixels in either direction spanned 2.73 cm. We can therefore refer to both horizontal and vertical distances in centimeters or pixels interchangeably.

All stimuli consisted of two *dotted-line segments*, one horizontal and one vertical, against a white background (mean luminance: 98 cd/m²). Each line segment consisted of 11 equispaced black dots (0.5 mm in diameter). In an initial control experiment ($n = 2$) following the same general procedure as described below, we measured the illusion magnitude for the configurations shown in Figure 1 using both solid and dotted lines. No significant difference was found between the illusion magnitude in the dotted- and the solid-line conditions (data not shown), illustrating that configura-
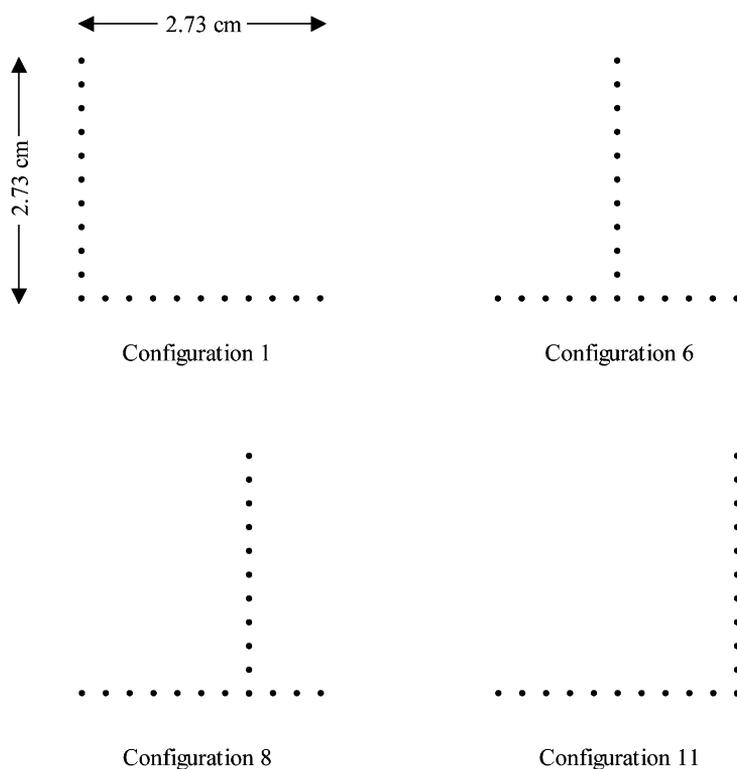
Figure 3. Four of the 11 configurations used in Experiment 1. Two perpendicular dotted lines, each composed of 11 equispaced points, are joined at a single point. In Configuration *n*, the vertical line is joined at the *n*th point of the horizontal, numbered 1 to 11 from left to right. Configuration 1 is the horizontal–vertical illusion configuration (Figure 1B). Configuration 6 is the bisection illusion configuration (Figure 1A).

tions with dotted-line segments are appropriate stimuli for the study of the present illusion. We used dotted rather than solid lines for two reasons: (1) to avoid aliasing artifacts in later experiments when we will draw lines at orientations other than horizontal and vertical, and (2) to remove any ambiguity in defining the length of the vertical line that can be interpreted as terminating at the top or bottom edge of the horizontal line. The difference is small for our stimuli (roughly 1% of the length of the horizontal lines) but not negligible. The length of the dotted line, from one extreme dot to the other, is unambiguous. Observers in our experiments were instructed to judge the distances from one extreme dot to the other (for both the horizontal and vertical extents), to eliminate the possibility that they might judge the vertical line as ending at the dot above the horizontal line (i.e., at the 10th instead of the 11th dot from the top). The control experiment mentioned above also illustrated that the illusion magnitude was not reduced when using dotted-line stimuli, thus making it unlikely that the subjects did not consider the full vertical extent of the dotted line.

The observer viewed the stimulus binocularly[6] from a distance of 76 cm. The horizontal line was always 2.73 cm in length (100 pixels, 2.1° of visual angle), whereas the length of the vertical line varied under control of the experimental program, as described below. In Experiment 1, the nonhorizontal line was always vertical. Its lower endpoint coincided with one of the 11 points in the horizontal line, resulting in a total of 11 stimulus configurations that differed only in the horizontal displacement of the vertical line. Configuration 1 corresponds to the L-shaped stimulus often used in studies of the horizontal–vertical illusion. In Configurations 2–11, the vertical line is progressively shifted to the right by 10 pixels (or

by one point of the horizontal line) with Configuration 6 being the inverted T stimulus, which is typically used in studies of the bisection illusion and the horizontal–vertical illusion. Four examples are shown in Figure 3. Seven of the 11 stimuli (Configurations 1, 3, 5, 6, 7, 9, and 11) are shown along the abscissa of the data plots (Figure 4).

Procedure. A trial began with presentation of a small fixation cross in the center of the display area for 500 msec, followed by a blank screen for 500 msec. The stimulus was then presented for 500 msec, followed by a blank screen until the observer responded by pressing a key. The next trial would then follow immediately. The center of the horizontal line of each stimulus configuration was always 1.1 cm (40 pixels) below where the fixation cross appeared.

On each trial, the observer judged whether the vertical line was longer or shorter than the horizontal line (a two-alternative forced choice task) and recorded his or her judgment by pressing one of two keys. The observer judged each of the 11 stimulus configurations 160 times (a total of 1,760 trials). The 1,760 trials were broken into eight blocks of 220 trials each (20 judgments of each stimulus configuration), with rest breaks between blocks.

Within each session, we adjusted the vertical length of each of the stimulus configurations according to the observer's responses by a staircase procedure. If the observer judged the vertical line in a stimulus configuration to be longer, the length of the vertical line was decreased by 0.082 cm (3% of the width of the horizontal line) on the next presentation.[7] If he or she judged it to be shorter, it was increased by the same amount. The 11 one-up–one-down staircases were randomly interleaved. The initial length of the vertical line in each configuration was set to a random value.
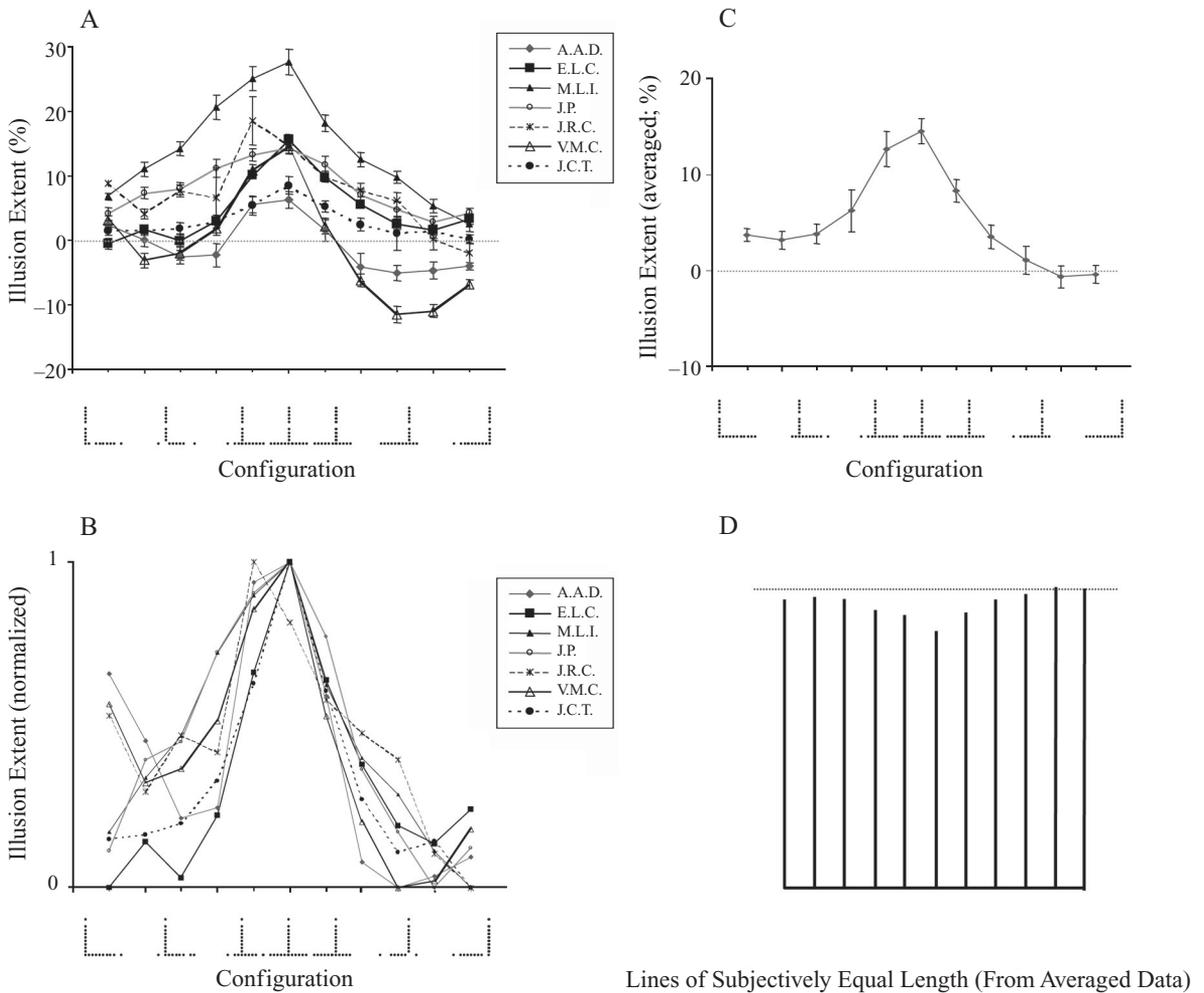
Figure 4. Results of Experiment 1 for the 7 observers. In panels A, B, and C, stimulus configuration is plotted and depicted along the horizontal axis with Configurations 1, 3, 5, 6, 7, 9, and 11 depicted. (A) The illusion extent corresponding to a given stimulus is plotted along the vertical axis. The vertical error bars represent one standard error. Despite large individual differences in distortion magnitude, observers exhibit a very similar pattern in which the distortion varies with stimulus configuration, with slight disruptions in stimulus symmetry causing marked reductions in illusion extent. (B) Illusion extent normalized as in Equation 3 versus stimulus configuration. (C) Illusion extent for each configuration averaged across subjects. (D) Illustration of the line length at each position that, on average, was judged to be equal to the length of the horizontal line. The dashed line marks the vertical distance to the horizontal that is equal to the length of the horizontal.

Analysis. For each observer and each configuration, we fit the staircase data to a psychometric function (the cumulative distribution function [c.d.f.] of a Gaussian variable) by a maximum likelihood procedure. We computed the maximum likelihood estimates of the point of subjective equality (PSE) for the horizontal and vertical lines, and the length corresponded to the 50th percentile on the fitted Gaussian c.d.f. We use the term *illusion extent* to denote the percentage by which the vertical line is shorter than the horizontal line at the PSE. In other words, a PSE of 75 (i.e., a vertical line of 75 pixels is judged to be as long as a horizontal line of 100 pixels) corresponds to an illusion extent of 25%. We calculated standard error estimates by a bootstrap procedure (Efron & Tibshirani, 1993).

Since all observers completed a large number of trials, we were concerned that their PSEs might not be stable across time. We initially compared results by estimating PSEs for each observer on the basis of just the first half of the data and again on just the second half. Comparison of these thresholds showed no consistent trend.

Results and Discussion

Configuration 6 is the bisection illusion configuration of Figure 1A. Configuration 1 is the horizontal–vertical configuration of Figure 1B, and Configuration 11 is its mirror-image. Configurations 2–10 all contain a T-junction occlusion cue. If the presence or absence of the T junction is responsible for the larger illusion in Figure 1A compared with that of Figure 1B, we expect that the illusion magnitudes for stimulus Configurations 2–10 would be identical to one another and greater than those for Configurations 1 and 11.

The results are shown in Figure 4A for each of the 7 observers. Stimulus configuration is plotted along the horizontal axis with Configurations 1, 3, 5, 6, 7, 9, and 11 depicted. The stimulus at the center of the horizontal

axis is the symmetrical inverted T (Configuration 6). Although it can be argued that symmetry/asymmetry is an all-or-nothing property, we will in what follows use the terms *more symmetric* and *more asymmetric* to refer to the configurations toward the center and toward the extremes of the axis, respectively. The illusion extent corresponding to a given stimulus is plotted along the vertical axis. Recall that a *larger* illusion extent implies that the observer considers a *shorter* vertical line as equal in length to the horizontal. The vertical error bars represent 1 standard error. For our group of observers, the illusion extent ranges from 6.2% to 27.7% for the inverted T (average: 15.5%), from −0.6% to 8.8% for the L (average: 3.7%), and from −6.9% to 4.2% for the mirror L (average: −0.4%).

We note, first of all, that there are large differences in distortion magnitude among observers. Possible individual differences are typically ignored in the literature on illusions. Past studies have typically included a large number of observers, each of whom completed a small number of trials; the resulting data were then averaged across observers, and only these averaged effects were reported.

Despite these large interobserver differences in absolute illusion extent, observers exhibit similar patterns of change in illusion extent with change in stimulus configuration. In Figure 4B, we replot the normalized results for each observer, scaled and translated vertically so that the largest illusion magnitude for each observer is 1 and the smallest is 0. Let E_i^k denote the illusion extent for the kth observer and the ith stimulus configuration ($i = 1, 2, \ldots, 11$). Let

$$M_k = \max\left\{ E_1^k, E_2^k, \ldots, E_{11}^k \right\}, \quad (1)$$

and

$$m_k = \min\left\{ E_1^k, E_2^k, \ldots, E_{11}^k \right\}. \quad (2)$$

Then the normalized illusion extents for the kth observer, e_i^k, ($i = 1, 2, \ldots, 11$), are computed by the formula

$$e_i^k = \frac{E_i^k - m_k}{M_k - m_k}. \quad (3)$$

Figure 4C shows the results in Figure 4A averaged across observers. In Figure 4D, we plot the relative line lengths that, on average, were judged to be equal to the length of the horizontal line according to the averaged data in Figure 4C.

Examination of the normalized results (Figure 4B) and averaged results (Figure 4C) indicate that the illusion extent is largest for the most symmetric configurations, decreases rapidly as the configuration becomes slightly asymmetric, and remains the same on average as the asymmetry is further increased. Six of the 7 observers show the largest illusion for Configuration 6, the (only) symmetrical stimulus. For the remaining observer (J.R.C.), the largest illusion is found for the slightly asymmetric Con-

figuration 5, in which the vertical is shifted to the left of the center by a distance corresponding to 10% of the horizontal line's length. In most cases, the steepest decrease in illusion extent can be observed as the configuration changes from the perfectly symmetric Configuration 6 to the slightly asymmetric Configurations 5 and 4 (to the left) and 7 and 8 (to the right). Further increases in asymmetry (i.e., changing to Configurations 3, 2, and 1, and 9, 10, and 11) lead to a slowing or in some cases a reversal of this trend. Only 2 observers (J.P. and M.L.I.) show a perfectly monotonic decrease on both sides of the curve so that the lowest illusion extent is observed at the L and mirror L configurations (1 and 11). For 2 observers (A.A.D. and V.M.C.), the illusion extent reaches its minimum for both the left and right sides at one of the asymmetric inverted T stimuli (Configurations 3 and 9 for 1 observer, and 2 and 9 for the other) but increases again as stimuli change toward the L and mirror L configurations. The curves for the 3 remaining observers (J.C.T., J.R.C., and E.L.C.) show different patterns on the two sides, partly decreasing monotonically (such as J.C.T. and J.R.C. on the left and right, respectively), partly showing the same pattern as A.A.D. and V.M.C. (such as E.L.C. on the left), and partly showing some different nonmonotonic decrease.

On average, the illusion extent decreased by 10.7% as the stimulus changed from the symmetrical Configuration 6 to Configuration 3 [$t(6) = 2.87$, $p = .028$, two-tailed]. The average decrease in illusion extent between Configurations 3 and 1 is 0.1% [$t(6) = 0.81$, n.s., two-tailed]. Similarly, on the right side of the curve, the illusion extent decreased by 13.4% as the stimulus changed from Configuration 6 to Configuration 9 [$t(6) = 2.85$, $p = .029$, two-tailed]. The average decrease in illusion extent between Configurations 9 and 11 is 1.4% [$t(6) = 1.31$, n.s., two-tailed]. It thus appears that small changes in the position of the vertical line can lead to significant decreases in distortion if these positional changes turn a symmetrical stimulus into an asymmetric one. The same positional changes, however, tend to have little effect when they simply increase the asymmetry of an already asymmetric stimulus.

Another striking feature in these data is the difference in illusion extent between configurations with the vertical line to the right of the center (1–5) and those with the vertical line to the left of the center (7–11). On average, the configurations in which the vertical line is a given distance to the right of the center induce smaller illusions than those with the vertical line at the same distance to the left of the center (t test; $p < .05$). This asymmetry might reflect the dominance of the right hemisphere in illusion perception and is consistent with results reported for the Müller-Lyer, Ponzo, Oppel–Kundt, and herringbone illusions under some conditions (Houlard, Fraisse, & Hecaen, 1976; Rasmjou, Hausmann, & Güntürkün, 1999; Rothwell & Zaidel, 1990). However, we have no basis to claim that the asymmetries in these studies and in ours share any common mechanism.[8]

If the T junction is the depth cue whose presence or absence leads to the difference in distortion magnitude observed between the L configuration and the inverted T configuration, any configuration with a T junction (Configurations 2–10) should elicit the same degree of distortion. Furthermore, the L and mirror L stimuli (Configurations 1 and 11), which contain no T junction, should always induce a smaller illusion than any of the T stimuli. The same would be true if the critical depth cue was the L junction, which might be thought of as a coplanarity cue signaling that the vertical and horizontal lines are at the same distance to the observer. The two configurations with such a coplanarity cue (1 and 11) should induce a smaller illusion magnitude than those without such a cue (2–10), which should all elicit the same illusion magnitude. Instead, we report that even slight asymmetries in the T configuration lead to a strong decrease in illusion extent and that L configurations induce the same—in some observers even larger—illusion magnitudes than some of the asymmetric T configurations. It appears that the illusion extent depends critically on the degree of symmetry in the configuration (not commonly thought of as a depth cue) rather than on the presence or absence of a particular junction type (a pictorial depth cue).

EXPERIMENT 2

How could we explain the importance of symmetry in Experiment 1 without abandoning the perspective theory model in Figure 2? Configurations composed of two joined line segments have few potential depth cues, a reason why we chose to study them. A second pictorial cue present in the configurations of Figures 1A and 1B is elevation in the visual field (Levine, 2000). In perspective terms, the plane containing the stimulus configuration (the *stimulus plane*) appears to the observer as if it were slanted backward. Other authors have suggested that there is a prior bias to see such a viewing plane (here, the stimulus plane) from slightly above (Mamassian & Landy, 2001)—that is, slanted away from the observer's line of sight.

If the stimulus plane (which is actually a frontoparallel plane, perpendicular to the observer's cyclopean line of sight) is perceived as slanted away[9] from the observer's line of sight, the size-scaling factor appropriate for a vertical line segment will increase relative to that of the horizontal. The perceived size of the line segment is the length of the perspective projection of the line segment from the observer's viewpoint into the slanted plane. This is simply the shadow that the line segment would cast on the plane if a small light source were placed at the observer's viewpoint. As the plane is slanted farther away from the observer's line of sight, the shadow of the vertical segment grows relative to that of the horizontal. In summary, the perspective theory model based on the elevation cue would correctly predict the direction of distortion seen in the bisection illusion of the inverted T configuration (Figure 1A).

However, the elevation cue is also present in the configuration of the horizontal–vertical illusion (Figure 1B). Displacing the vertical line segment horizontally, away from the line of sight and keeping it vertical, increases the length of the projective shadow. For the same degree of illusory slant in the stimulus plane, away from the observer, the projection of a vertical line segment in the L configuration will be *longer* than the projection of the same vertical line segment in the inverted T configuration. As the vertical line is displaced to the left or right, away from the centrally located fixation point, as in Experiment 1, the illusion extent should increase or at least not decrease. This is the opposite of what we found.

Although the elevation cue alone cannot account for the results of Experiment 1, it can do so if combined with a prior bias toward depth interpretations where lines in the scene intersect at right angles. Figure 5 shows a perspective drawing that illustrates this possibility. The horizontal lines in the drawing's lower half represent lines in the ground plane. Lines connecting the lowest horizontal *ground line* to the vanishing point V represent lines orthogonal to the horizontal that also lie in the ground plane and that are receding into the distance. Of these, the only line that is also perpendicular to the horizontal in the 2-D projection originates at the center point C, the point on the ground line with the shortest perpendicular distance from the vanishing point. This vertical line is along the line of sight for an observer located in the cen-

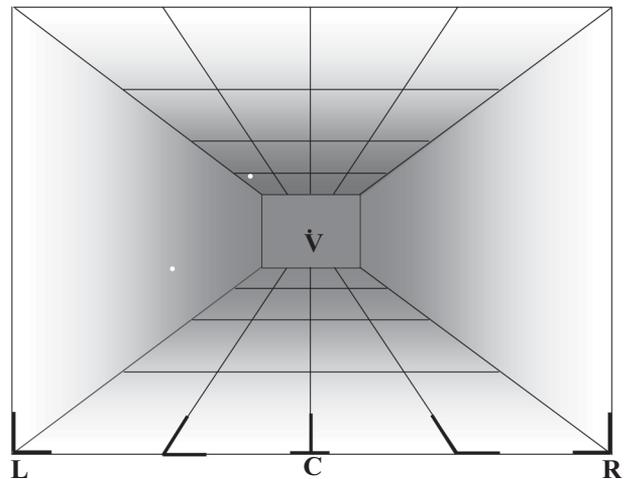**L**   **C**   **R**

Figure 5. Perspective drawing illustrating the model tested in Experiment 2. Horizontal lines in the drawing's lower half represent lines in the ground plane. Lines connecting a horizontal ground line to vanishing point V represent lines orthogonal to the horizontal that also lie in the ground plane and that are receding into the distance. Such receding lines are vertical (i.e., perpendicular to the horizontal in the 2-D projection) if they originate at center point C, but form an angle other than 90º with the horizontal if they originate at points to the left or right of the center. Vertical lines that do not originate at point C (such as those starting at points L and R) do not represent receding lines, but are interpreted as lying in the frontoparallel plane, orthogonal to the horizontal ground line.
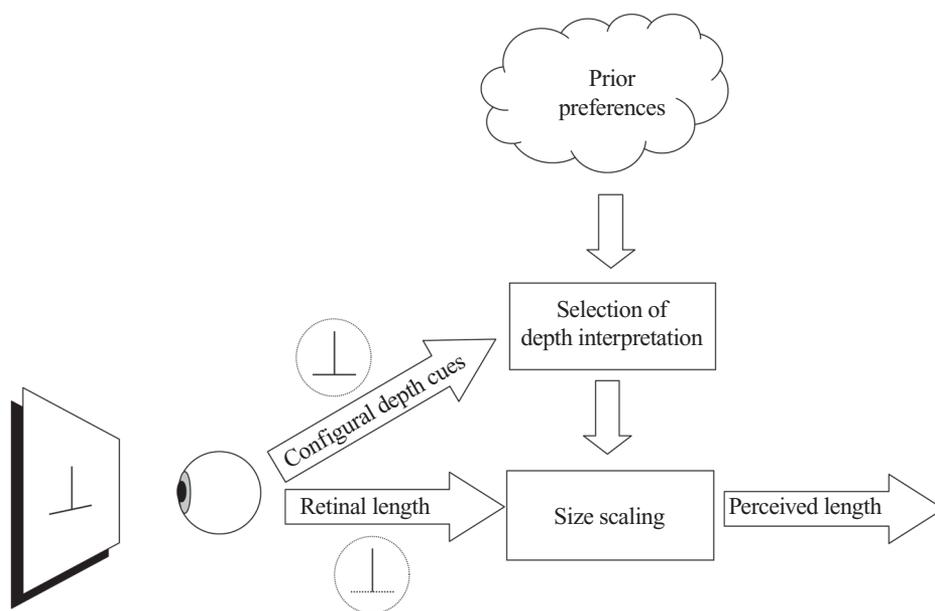
Figure 6. Information flow in a revised perspective theory. It is identical to that depicted in Figure 2, except that the selection of depth interpretation is influenced by prior preferences of certain 3-D interpretations—for example, orthogonal line intersections.

ter of the scene shown. Vertical lines that do not originate at point C (such as those starting at points L and R) do not represent lines receding into the distance, but are interpreted as lying in the frontoparallel plane, orthogonal to the horizontal ground line. In order for a line starting at those points to be interpreted as both receding and orthogonal to the horizontal, it has to form an angle other than 90° with the horizontal in the 2-D projection.

Other researchers have found evidence indicating a prior preference for orthogonal line intersections in 3-D (e.g., Griffiths & Zaidi, 2000). One of the oldest examples is the Ames trapezoid (Ames, 1951). In viewing the Ames configuration, observers show a preference to interpret an angle as 90°, and in doing so, override other depth cues. The angles are not 90°, and the result is perceptual error. Several authors have considered using the formalism of Bayesian decision theory as a way to incorporate such prior assumptions into visual scene interpretation (Knill & Richards, 1996; Maloney, 2002; Mamassian, Landy, & Maloney, 2002).

To explain the data from Experiment 1, we hypothesize that only vertical lines that coincide with the axis of symmetry of a configuration (as does the line in the inverted T) are interpreted as originating at the visual scene's center point and are therefore consistent with the interpretation that they are both orthogonal to the horizontal and receding into the distance. Vertical lines in asymmetrical configurations are seen as originating at points to the left or right of the center, and the only way they can be orthogonal to the horizontal is if they are not receding but instead lying in the frontoparallel plane.

Consistent with the results of Experiment 1, the lines in the asymmetrical configurations would be less likely to trigger size-constancy scaling and would lead to an illusion that is smaller than that triggered by the symmetrical configuration. The prior bias toward orthogonality counteracts the prior bias for a slanted stimulus plane and the effect of elevation, thus decreasing the illusion for all but the configuration of Figure 1A, the inverted T of the bisection illusion.

We then predict the following: In asymmetrical configurations, such as the L stimuli, the illusion extent should increase if the stimuli are modified in a way that permits an interpretation of the nonhorizontal line as both orthogonal to the horizontal and receding into the distance. This could be achieved by varying the angle between the two lines away from 90°. For example, the nonhorizontal lines of the two L configurations in Figure 5 with obtuse and acute angles represent receding lines that are orthogonal to the horizontal. We would predict configurations like these to trigger misapplied size-constancy scaling, therefore leading to an increase in illusion extent relative to the right-angle L stimulus used in Experiment 1. On the other hand, since the inverted T stimulus is already consistent with an interpretation of the vertical as both orthogonal to the horizontal and receding into the distance, we would predict no further increases and possibly a decrease in illusion extent for angles other than 90° in this symmetric configuration.

In Figure 6, we present a revised diagram of the flow of information in perspective theories. Other nonsensory sources of information now play a role in depth interpretation.

Method

Observers. Six observers (4 female and 2 male), all of whom had taken part in Experiment 1, were paid to participate in the experiment. All had normal or corrected-to-normal vision and were

naive about of the purpose of the experiment. Each required 1 h 45 min to 2 h to complete eight sessions with rest breaks.

Stimuli and Apparatus. We tested our prediction by varying the angle between the horizontal and the nonhorizontal line in the L and inverted T configurations between values of 30º and 150º in 30º steps. The 10 different stimuli can be seen on the abscissas of Figures 7A, 7B, and 7C and Figure 8. The apparatus used was the same as in Experiment 1.

Procedure. The procedure was generally the same as in Experiment 1. On each trial, the observer judged whether the nonhorizontal line was longer or shorter than the horizontal line and recorded his or her judgment by pressing one of two keys. The observer judged each of the 10 stimulus configurations 160 times (a total of 1,600 trials). The 1,600 trials were broken into eight sessions of 200 trials each (20 judgments of each stimulus configuration) with rest breaks between sessions.

Analysis. The analysis was the same as in Experiment 1. In addition, since the 6 subjects in Experiment 2 had participated in Experiment 1, we were able to test whether illusion magnitude changes

with repeated viewing of the stimuli. A comparison of the illusion magnitude for the T and L configurations in the two experiments yielded no significant differences (paired, two-tailed t test; $p > .05$).

Results and Discussion

The results for the L stimuli are shown in Figure 7 for all observers. Stimulus configuration is plotted along the horizontal axis with the five configurations depicted. From left to right along the axis, the angle between the horizontal and nonhorizontal lines of the L changes from acute to obtuse in 30º steps, with the center stimulus (Configuration 3) being the L with a right angle. The illusion extent corresponding to a given stimulus is plotted along the vertical axis. The vertical error bars represent one standard error. Figures 7B and 7C show the normalized and averaged data, respectively, and Figure 7D shows an illustration that, for each angle exam-
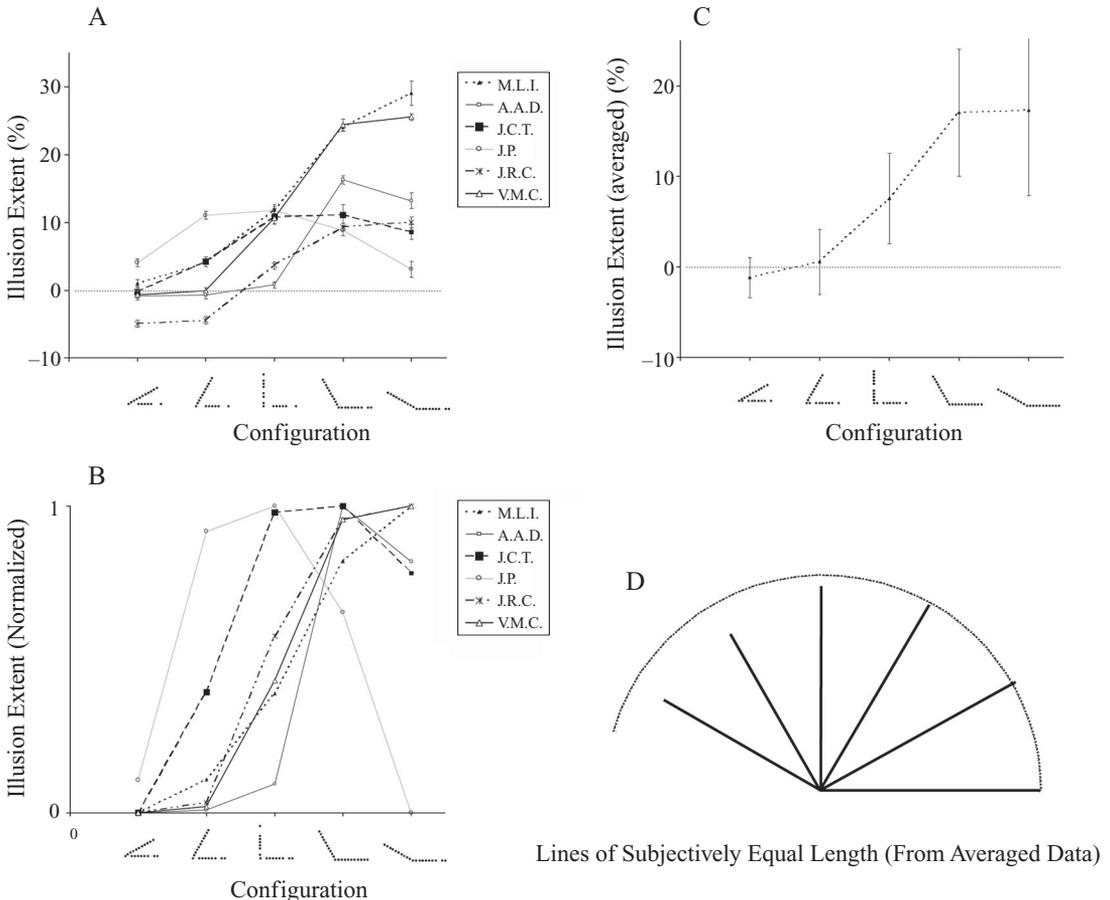

Figure 7. Experiment 2: Results for the L stimuli for the 6 observers. In panels A, B, and C, stimulus configuration is plotted and depicted along the horizontal axis. From left to right along the axis, the angle between the horizontal and nonhorizontal lines of the L changes from acute to obtuse in 30º steps, with the center stimulus (Configuration 3) being the L with a right angle. (A) The illusion extent corresponding to a given stimulus is plotted along the vertical axis. The vertical error bars represent one standard error. The illusion extent tends to increase as the angle between the two lines changes from acute to obtuse. (B) Illusion extent normalized as in Equation 3 versus stimulus configuration. (C) Illusion extent for each configuration averaged across observers. (D) Illustration that, for each angle examined, depicts the line length that, on average, was judged to be equal to the length of the horizontal line. The dotted arc marks the radial distance to the left endpoint of the horizontal that is equal to the length of the horizontal.
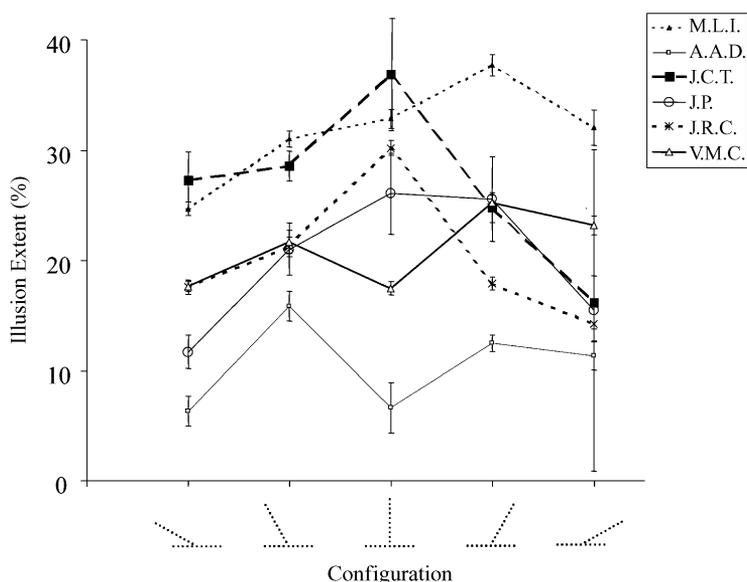
Figure 8. Experiment 2: Results for the T configuration with stimulus configuration plotted and depicted along the horizontal axis and corresponding illusion extent plotted on the vertical axis. The vertical error bars represent one standard error. The data show no consistent pattern of change in distortion magnitude with configuration for the group of observers.

ined, depicts the line length that, on average, was judged to be equal to the length of the horizontal line.

For our group of observers, the illusion extent ranges from −5.0% to 5.0% for the most acute angle (average: −0.4%), from 0.7% to 11.9% for the right angle (average: 7.0%), and from 3.0% to 29.0% for the most obtuse angle (average: 13.3%). The general trend, seen in the figures and exhibited by all observers but one (JP), is for the illusion extent to increase as the angle between the horizontal and nonhorizontal changes from acute to obtuse. That means that, contrary to our hypothesis, making the angle in the L configurations more acute did not lead to an increase in illusion extent. Instead, the illusion extent increased only when the angle was made more obtuse but *decreased* when it was made more acute. This trend is opposite to the prediction of our hypothesis. We do find that a vertical line does not create the greatest possible distortion, and that, depending on the configuration, equal deviations from vertical can increase or decrease the illusion extent.

One reviewer noted that for the stimuli with the most acute angles, the free endpoints of the lines of the L configuration come closer together and the lines become closer to parallel. He suggested that this could make a judgment of relative length appear more like a vernier alignment task and in itself lead to a decrease in the illusion extent. We disagree. The issue is whether the observer interprets line length through a perspective transformation and what perspective transformation he or she chooses for different angles. There is no obvious reason why the convergence of the two lines should lead to a

choice of perspective transformation that renders the two lines equal in retinal length. The hypothesis under test, for example, predicts otherwise.

Consistent with our prediction, we did not see an increase in the illusion extent in the symmetric T stimulus as the angle between the two lines was changed away from 90°. Results for the T configuration are shown in Figure 8, with stimulus configurations plotted and depicted along the horizontal axis and the corresponding illusion extent plotted on the vertical axis. The data show no consistent pattern of change in distortion magnitude with configuration for the group of observers. The illusion extent ranged from 6.3% to 27.2% for Configuration 1 (average: 17.5%), from 15.8% to 28.5% for Configuration 3 (average: 25%), and from 11.34% to 32.0% in Configuration 5 (average: 18.7%).

EXPERIMENT 3

The trend in illusion extent that we observed in Experiment 2 was the opposite of the trend that we expected. However, in the L configurations in Experiment 2, the angular deviations (30° and 60°) from the right angle were large. The nonhorizontal line can be interpreted as orthogonal to the horizontal only if the slanted plane containing the stimuli is markedly slanted away from the observer (as in the lower part of Figure 5). Other depth cues (such as binocular disparity) are available to the observer that signal that the plane containing the stimuli is orthogonal to the line of sight. These cues conflict with the interpretation of the stimuli as simultaneously receding

in depth and with the nonvertical line being orthogonal to the horizontal. This conflict may preclude the interpretation of the two lines as orthogonal.

With smaller angular deviations, the reduction in cue conflict might not override the preference for orthogonality. We would still have no explanation for the large-angle results of Experiment 2, but by testing with smaller angular deviations, we can determine whether a preference for orthogonal interpretations plays a role in the observed distortions of length. We tested this possibility in 2 observers by presenting L stimuli with angles deviating from a right angle by only 5º and 10º in both directions. We also included conditions with T stimuli, as in Experiment 2.

Method

Observers. Two female observers participated in the experiment. Both were authors of the present article.[10] Each required about 2 h to complete eight sessions with rest breaks. Both had normal or corrected-to-normal vision.

Stimuli and Apparatus. The stimuli were similar to those employed in Experiment 2, except that we now varied the angle between the horizontal and the nonhorizontal line in the L and the inverted T configurations between values of 80º and 100º in 5º steps. The 10 different stimuli can be seen on the abscissas of Figures 9A and 9B. The apparatus used was the same as in Experiments 1 and 2.

Procedure. The procedure was the same as in Experiment 2. On each trial, the observer judged whether the nonhorizontal line was longer or shorter than the horizontal line, and recorded her judgment by pressing one of two keys. The observer judged each of the 10 stimulus configurations 160 times (a total of 1,600 trials). The 1,600 trials were broken into eight sessions of 200 trials each (20 judgments of each stimulus configuration) with rest breaks between sessions.

Analysis. The analysis was the same as in Experiment 2.

Results and Discussion

The results for the L configuration are shown in Figure 9A. We tested for a linear trend across angle by applying the contrast $[-2\ -1\ 0\ 1\ 2]$ to the five means arranged in order of increasing angle and computing the standard deviation of the contrast from the bootstrap estimates of standard error for each mean. For both observers, there was a significant increase in the apparent length of the oblique line with increasing angle (M.T., $z = 5.435, p < .0001$; U.W., $z = 6.7414, p < .0001$), the same pattern found in the results for Experiment 2 with large-angle differences, and the opposite trend to that expected if the orthogonal preference hypothesis under test were true. The results for the T configuration have the same V-shaped pattern of results as was found with large-angle deviations.

CONCLUSIONS

In this article, we have concentrated on perspective theory explanations of distortions of perceived length. We did so partly because such theories are highly influential and partly because, if they proved to be valid, they would provide remarkably parsimonious predictions of perceived length distortions. The evident weakness of such theories as they are commonly employed is that there are no rules for deciding *in advance* what depth interpretations will control size-constancy scaling. The upper pathway in Figures 2 and 6 is left unspecified. Perspective theories are typically invoked post hoc, to explain previously observed distortions of perceived length and thus have little predictive power.
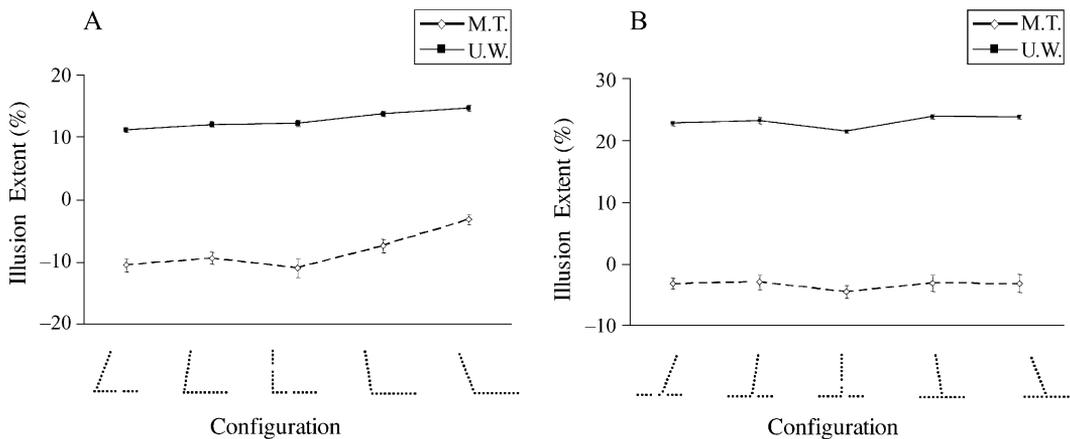

Figure 9. Experiment 3: Results of a replication of Experiment 2 with smaller angular deviations from the vertical. (A) Results for the L stimuli for 2 observers. From left to right along the horizontal axis, the angle between the horizontal and nonhorizontal lines of the L changes from acute to obtuse in 5º steps, with the center stimulus (Configuration 3) being the L with a right angle. The illusion extent corresponding to a given stimulus is plotted along the vertical axis. The vertical error bars represent one standard error. The apparent length of the oblique line tends to increase as the angle between the two lines changes from acute to obtuse. (B) Results for the T stimuli for the same 2 observers. From left to right along the horizontal axis, the angle between the horizontal and nonhorizontal lines of the L changes from acute to obtuse in 5º steps, with the center stimulus (Configuration 3) being the T with a right angle. The illusion extent corresponding to a given stimulus is plotted along the vertical axis. The vertical error bars represent one standard error. The apparent length of the oblique line varies little with angle.

In two experiments, we attempted to develop a predictive theory of length distortion for a large class of two-line configurations that included two classic illusion configurations, the horizontal–vertical illusion and the bisection illusion. We chose not to treat these illusions as isolates but instead as representatives of the many ways two lines can be joined together in the visual field.

We first considered the hypothesis that the presence or absence of an occlusion cue was responsible for the observed difference in the magnitude of length distortion in the bisection illusion (Figure 1A) and horizontal–vertical illusion configurations (Figure 1B). We measured the illusion extent for 11 configurations, 9 of which had an occlusion cue at different locations in the configuration, and 2 of which were missing this cue. We measured the illusion extent for each configuration and observer separately and found marked individual differences. Our average results across observers for the bisection illusion and horizontal–vertical illusion configurations are in good agreement with the literature where typically only the average results across many observers are reported. Whereas different observers experienced different magnitudes of illusion, the overall patterns across observers were very similar and inconsistent with the predictions of the occlusion cue hypothesis. Even slight asymmetries in the inverted T junction configuration led to marked decreases in the illusion extent. We also found a slight asymmetry in the illusion extent as a function of the position of the vertical line in the left or right part of the visual field. The illusion extent in the left visual field was larger for many observers, which is a typical result in the literature.

We next developed a model that combined a preference for a 3-D interpretation of joint lines as orthogonal with a bias toward interpreting the stimulus configuration as slanted away from the line of sight, both of which are well-known prior preferences in visual interpretation of scenes (Mamassian & Landy, 2001). To the extent that the configuration is slanted away, the different elevations of the horizontal and vertical lines in the visual field would account, qualitatively, for both the bisection illusion and the horizontal–vertical illusion, but not for the relative magnitude of the illusions. Under the orthogonal preference hypothesis, the lesser magnitude of the horizontal–vertical illusion resulted from a conflict between prior preferences for scene interpretations with orthogonal lines and configurations slanted in depth.

We tested a prediction of the orthogonal preference hypothesis: The magnitude of illusion for the L configuration should be larger if the angle enclosed by the L were less than 90º (acute). Then the stimulus configuration could be interpreted as slanted away from the observer with the two lines intersecting at right angles. If the angle were obtuse, we would expect the same or a smaller illusion. For the T configuration of the bisection illusion, we would predict only a decrease in illusion for angles other than 90º.

We tested this hypothesis in two experiments. In Experiment 2, the angles between the lines in both L and T configurations were varied from 60º to 120º, in steps of 30º. In Experiment 3, the angles were varied only from 80º to 110º in steps of 5º. Contrary to the prediction based on the orthogonal preference hypothesis, we found, in both experiments, that deviations from a right angle in the L configuration led to an increase in illusion only if the angle was made obtuse. If the angle was made acute, a decrease in illusion, and sometimes a reversal, resulted. This is inconsistent with a simple model containing only assumptions about a preference for orthogonality and a bias in estimating the plane containing the stimulus.

In summary, we have created a much-needed gap in the literature. It is evident that currently no model can account for how the human observer will interpret two arbitrarily joined line segments in the frontoparallel plane. Despite the individual differences we found, there was enough of a common pattern in the data to suggest that such a model is possible. We simply do not know what it is.

REFERENCES

AMES, A. (1951). Visual perception and the rotating trapezoidal window. *Psychological Monographs*, **65**(7).

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. *Psychonomic Science*, **26**, 77-79.

BEGELMAN, D. A., & STEINFELD, J. (1971). Retinal vs. phenomenal determinants of the horizontal–vertical illusion. *Journal of General Psychology*, **85**, 77-85.

BLUMENFELD, W. (1913). Untersuchungen über die scheinbare Grösse im Sehraume. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, **65**, 241-404.

COHEN, P. R., & FEIGENBAUM, E. A. (1982). *The handbook of artificial intelligence* (Vol. 3, pp. 139-194). Los Altos, CA: Kaufmann.

COREN, S., & GIRGUS, J. S. (1978). *Seeing is deceiving: The psychology of visual illusions*. Hillsdale, NJ: Erlbaum.

EFRON, B., & TIBSHIRANI, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.

FILEHNE, W. (1898). Die geometrisch-optischen Täuschungen als Nachwirkungen der in körperlichen Sehen erworbenen Erfahrung. *Zeitschrift für Psychologie*, **17**, 15-61.

GILLAM, B. (1980). Geometrical illusions. *Scientific American*, **242**, 102-111.

GILLAM, B. (1998). Illusions at century's end. In J. Hochberg (Ed.), *Perception and cognition at century's end* (pp. 95-136). New York: Academic Press.

GIRGUS, J. S., & COREN, S. (1975). Depth cues and constancy scaling in the horizontal–vertical illusion: The bisection error. *Canadian Journal of Psychology*, **29**, 59-65.

GREENE, E. (1998). A test of the gravity lens theory. *Perception*, **27**, 1221-1228.

GREGORY, R. L. (1963). Distortion of visual space as inappropriate constancy scaling. *Nature*, **199**, 678-680.

GREGORY, R. L. (1970). *The intelligent eye*. New York: McGraw-Hill.

GREGORY, R. L. (1973). A discussion of G. H. Fisher's "Toward a new explanation for the geometrical illusions: Apparent depth or contour proximity?" and the inappropriate constancy-scaling theory. *British Journal of Psychology*, **64**, 623-626.

GREGORY, R. L. (1997). *Eye and brain: The psychology of seeing*. Princeton, NJ: Princeton University Press.

GRIFFITHS, A. F., & ZAIDI, Q. (2000). Perceptual assumption and projective distortions in a three-dimensional shape illusion. *Perception*, **29**, 171-200.

HOULARD, N., FRAISSE, P., & HECAEN, H. (1976). Effects of unilateral hemispheric lesions on two types of optico-geometric illusions. *Cortex*, **12**, 232-240.

HUFFMAN, D. A. (1971). Impossible objects as nonsense sentences. In R. Meltzer & D. Michie (Eds.), *Machine intelligence 6* (pp. 295-323). New York: Elsevier.

KNILL, D. C., & RICHARDS, W. (1996). *Perception as Bayesian inference.* Cambridge: Cambridge University Press.

KÜNNAPAS, T. M. (1955). An analysis of the vertical–horizontal illusion. *Journal of Experimental Psychology*, **49**, 134-140.

LEVINE, M. W. (2000). *Levine and Shefner's Fundamentals of sensation and perception* (3rd ed.). Oxford: Oxford University Press.

LUNEBURG, R. K. (1950). The metric of binocular visual space. *Journal of the Optical Society of America*, **40**, 627-642.

MALONEY, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world* (pp. 145-189). New York: Wiley.

MAMASSIAN, P., & LANDY, M. S. (2001). Interaction of visual prior constraints. *Vision Research*, **41**, 2653-2668.

MAMASSIAN, P., LANDY, M. S., & MALONEY, L. T. (2002). Bayesian modeling of visual perception. In R. Rao, B. Olshausen, & M. Lewicki (Eds.), *Probabilistic models of perception and brain function* (pp. 13-36). Cambridge, MA: MIT Press.

MCGRAW, P. V., & WHITAKER, D. (1999). Perceptual distortions in the neural representation of visual space. *Experimental Brain Research*, **125**, 122-128.

MILNER, A. D., & GOODALE, M. A. (1995). *The visual brain in action.* Oxford: Oxford University Press.

OPPEL, J. J. (1855). Über geometrisch-optische Täuschungen. *Jahresbericht des Frankfurter Vereins 1854-55*, **55**, 37-47.

PRINZMETAL, W., & GETTLEMAN, L. (1993). Vertical–horizontal illusion: One eye is better than two. *Perception & Psychophysics*, **53**, 81-88.

RASMJOU, S., HAUSMANN, M., & GÜNTÜRKÜN, O. (1999). Hemispheric dominance and gender in the perception of an illusion. *Neuropsychologia*, **37**, 1041-1047.

ROTHWELL, B., & ZAIDEL, E. (1990). Visual field differences in the magnitude of the Oppel–Kundt illusion vary with processing time. *Perception & Psychophysics*, **47**, 180-190.

SEDGWICK, H. A., & NICHOLLS, A. L. (1993). Cross talk between the pictured surface and the pictured scene: Effects on perceived shape. *Perception*, **22**(Suppl. 109).

SEDGWICK, H. A., NICHOLLS, A. L., & BREHAUT, J. (1995). Perceptual interaction of surface and depth in optically minified pictures. *Investigative Ophthalmology & Visual Science*, **36**, S667.

THIÉRY, A. (1896). Über geometrisch-optische Täuschungen. *Philosophische Studien*, **12**, 67-126.

THOMPSON, J., & SCHIFFMAN, H. R. (1974). The effect on the magnitude of the horizontal–vertical illusion of horizontal retinal eccentricity. *Vision Research*, **14**, 1463-1465.

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.

WILLIAMS, P. A., & ENNS, J. T. (1996). Pictorial depth and framing have independent effects on the horizontal–vertical illusion. *Perception*, **25**, 921-926.

WOODWORTH, R. S. (1938). *Experimental psychology*. New York: Holt.

NOTES

1. Although these theories are intended to predict distortions of perceived length in the frontoparallel plane, they are evidently relevant to findings of fixed distortions of the geometry of binocular visual space (Blumenfeld, 1913; Luneburg, 1950). A complete theory of the geometry of visual space must include the retinal and encoding distortions just discussed.

2. One finding of this article is that large individual differences in perception of the illusion configurations are under study. Some readers may not perceive the illusion described in Figure 1A or 1B or with the configuration of Figure 1A rotated 90º.

3. Theories of this sort were proposed by researchers well before Woodworth. See, in particular, Thiéry (1896), Filehne (1898), and the useful reviews by Gillam (1980, 1998).

4. Of course, candidate cues to scene layout present in the L junction are not present in the inverted T configuration. For example, the L junction could be interpreted as the corner of a planar surface or the edge of a hole. In Experiment 1, we focus on the presence or absence of the inverted T junction using the L junction only as a point of comparison. The L junction, as a cue to scene layout, could certainly be studied in its own right.

5. The 11th configuration is the mirror reflection of Figure 1B.

6. We employed binocular viewing for the following reasons: (1) Any illusion due to erroneous depth interpretation is all the more impressive when binocular disparity contradicts the erroneous interpretation; (2) illusions are reported to be stronger under binocular viewing than under monocular (Prinzmetal & Gettleman, 1993); and (3) most past work on illusions has made use of binocular viewing.

7. The length of the line was computed and stored as a high-precision floating point number, which was adjusted according to the outcome of each staircase trial. Whenever the line was drawn on the screen as a dotted line, the location of each of the dots was computed to machine precision and then rounded so as to conform to the pixel grid of the monitor. Every dot was within 0.05 cm of its computed location.

8. One reader of an earlier version of this article thought that the illusion strength for the extreme inverted T configurations might be even less than that for the L configurations (see, e.g., Observer V.M.C. in Figure 4A). We formed a linear contrast for each observer by adding the PSEs for Configurations 1 and 11 and subtracting those for Configurations 2 and 10. The expected value of this contrast should be positive if the reviewer's conjecture is correct. We tested the hypothesis that the contrast value was 0 against the alternative that it was positive (one-tailed), basing the test on the bootstrap estimates of the standard deviation of each PSE. The null hypothesis was rejected only for Observer V.M.C. ($p < .001$).

9. To be precise: We are assuming that the stimulus plane appears to the observer to be rotated around a horizontal axis embedded in the stimulus plane, perpendicular to the observer's cyclopean line of sight. The top half appears to be farther away than the bottom half.

10. Both observers also completed pilot versions of Experiments 1 and 2. Their results in those experiments (not reported here) were similar to those of the observers who were naive about the purpose of the experiment.