A statistical explanation of visual space

Zhiyong Yang & Dale Purves

The subjective visual space perceived by humans does not reflect a simple transformation of objective physical space; rather, perceived space has an idiosyncratic relationship with the real world. To date, there is no consensus about either the genesis of perceived visual space or the implications of its peculiar characteristics for visually guided behavior. Here we used laser range scanning to measure the actual distances from the image plane of all unoccluded points in a series of natural scenes. We then asked whether the differences between real and apparent distances could be explained by the statistical relationship of scene geometry and the observer. We were able to predict perceived distances in a variety of circumstances from the probability distribution of physical distances. This finding lends support to the idea that the characteristics of human visual space are determined probabilistically.

Visual space is characterized by perceived geometrical properties such as distance, linearity and parallelism. Intuitively, it seems that these properties are the result of a direct transformation of the Euclidean characteristics of physical space¹. This assumption is, however, inconsistent with a variety of puzzling and often subtle discrepancies between the predicted consequences of any direct mapping of physical space and what people actually see.

A number of examples in perceived distance, the simplest aspect of visual space, show that the apparent distance of objects bears no simple relation to their physical distance from the observer¹⁻⁴. Thus, when subjects are asked to make judgments with little or no contextual information, the distances reported differ in several ways from the corresponding physical distances. First, objects in these circumstances are typically perceived to be at a distance of 2-4 m, a phenomenon referred to as the 'specific distance tendency'5,6 (Fig. 1a). Second, the distance of an object from the observer appears to be about the same as that of neighboring objects in the retinal image, a phenomenon called the 'equidistance tendency'⁵ (Fig. 1b). Third, when presented at or near eye level, the distance of an object relatively near to the observer tends to be overestimated, whereas the distance of an object that is farther away tends to be underestimated^{7–11} (Fig. 1c). Fourth, the apparent distance of objects on the ground varies with the angle of declination of the line of sight¹²; objects on the ground that are at least several meters away appear closer than they really are and progressively more elevated than warranted by their physical position¹³ (Fig. 1d). Finally, under realistic outdoor conditions, the perceived distance of objects on the ground is influenced by the intervening structure of the ground surface^{14,15} (Fig. 1e). Although a variety of explanations have been proposed, there has been little or no agreement about the basis of this phenomen-ology $^{1-3}$.

Here we explore the hypothesis that these anomalies of perceived distance are all manifestations of a probabilistic strategy for generating visual percepts in response to inevitably ambiguous visual stimuli^{16,17}. A straightforward way to examine this idea in the case of visual space is to analyze the statistical relationship between geomet-

rical features (for example, points, lines and surfaces) in the retinal image plane and the corresponding physical geometry in representative visual scenes. Accordingly, we used a database of natural scene geometry acquired with a laser range scanner to test whether the otherwise puzzling phenomenology of perceived distance can be explained in statistical terms. Our results show that perceived distance is always biased toward the most probable physical distance underlying the stimulus, consistent with the general idea that the structure of visual space is determined statistically according to the probability distributions of the possible stimulus sources.

RESULTS

A probabilistic concept of visual space

The inevitable ambiguity of visual stimuli presents a challenge for generating perceptions of distance (and spatial relationships more generally). When a point in space is projected onto the retina, the corresponding point in the retinal projection could have been generated by an infinite number of different locations in the physical world. Similarly, an array of points in the retinal image could have arisen from an infinite number of physical configurations. Therefore, the relationship between any projected image and its source is inherently ambiguous. Nevertheless, the distribution of the distances of unoccluded object surfaces from the observer and their spatial relationships in normal viewing must have a potentially informative statistical structure. Given this inevitable ambiguity in vision, it seems likely that visual systems have evolved to take advantage of such statistical structure, or probabilistic information, in generating perceptions of physical space.

Any probabilistic strategy of this sort can be formalized in terms of Bayesian 'optimal observer theory'^{16,18–23}. In this framework, the probability distribution of physical sources underlying a visual stimulus, P(S|I) can be expressed as

$$P(S|I) = P(I|S)P(S)/P(I)$$
(1)

Department of Neurobiology, Box 3209, Duke University Medical Center, Durham, North Carolina 27710, USA. Correspondence should be addressed to Z.Y. (yang@neuro.duke.edu).

ARTICLES



d Perceived distance of objects on the ground



e Effects of terrain on distance perception



where S represents the parameters of physical scene geometry and I represents the visual image. P(S), is the probability distribution of scene geometry in typical visual environments (the prior), P(I|S) the probability distribution of stimulus I generated by the scene geometry S (the likelihood function), and P(I) is a normalization constant.

If visual space is indeed determined by the probability distribution of scene geometry underlying visual stimuli, then, under reduced-cue conditions, the prior probability distribution of distances to the observer in typical viewing environments should bias perceived distances. By the same token, the conditional probability distribution of the distances between locations in a scene should bias the apparent relative distances among them. Finally, when addiFigure 1 Anomalies in perceived distance. (a) Specific distance tendency. When a simple object is presented in an otherwise dark environment, observers usually judge it to be at a distance of 2-4 m, regardless of its actual distance. (In these diagrams, which are not to scale, 'Phy' indicates the physical position of the object and 'Per' the perceived position.) (b) Equidistance tendency. Under these same conditions, an object is usually judged to be at about the same distance from the observer as neighboring objects, even when their physical distances differ. (c) Perceived distance of objects at eye level. The distances of nearby objects presented at eye level tend to be overestimated, whereas the distances of farther objects tend to be underestimated. (d) Perceived distance of objects on the ground. An object on the ground a few meters away tends to appear closer and slightly elevated with respect to its physical position. Moreover, the perceived location becomes increasingly elevated and relatively closer to the observer as the angle of the line of sight approaches the horizontal plane at eye level. (e) Effects of terrain on distance perception. Under more realistic conditions, the distance of an object on a uniform ground-plane a few meters away from the observer is usually accurately perceived. When, however, the terrain is disrupted by a dip (upper panel), the same object appears to be farther away; conversely, when the ground-plane is disrupted by a hump (lower panel), the object tends to appear closer than it is.

tional information pertinent to distance is present, these 'default' biases will be reduced.

Below, we show that these predictions explain the phenomenology of apparent distance (Fig. 1).

Probability distributions of distances in natural scenes

The information at each pixel in the range image database is the distance, elevation and azimuth of the corresponding location in the physical scene relative to the laser scanner (Fig. 2). These data were used to compute the distribution of distances from the center of the scanner to locations in the physical scenes in the database.

Several statistical features were apparent in the analysis. First, the probability distribution of the radial distances from the scanner to physical locations in the scenes has a maximum at about 3 m, declining approximately exponentially over greater distances (Fig. 3a). This distribution is scale invariant, meaning that any scaled version of the geometry of a set of natural scenes will, in statistical terms, be much the same^{24,25}. This behavior is presumably due to the fact that the farther away an object is, the less area it spans in the image plane and the more likely it is to be occluded by other objects. Indeed, a simple model that incorporates this fact generates a scaling-invariant distribution of object distances nearly identical to that obtained from natural scenes (Fig. 3 legend).

A second statistical feature of the analysis concerns how different physical locations in natural scenes are typically related to each other with respect to distance from the observer. The distribution of the differences in the distance from the observer to any two physical locations is highly skewed, having a maximum near zero and a long tail (Fig. 3b). Even for angular separations as large as 30°, the most probable difference between the distances from the image plane of two locations is minimal.

A third statistical feature is that the probability distribution of horizontal distances from the scanner to physical locations changes relatively little with height in the scene (the height of the center of the scanner was always 1.65 m above the ground, thus approximating eye level of an average adult; Fig. 3c). The probability distribution of physical distances at eye level has a maximum at about 4.7 m and decays gradually as the distances increase (all distributions again being scale-invariant). The probability distributions of the



Intensity of laser return coded in gray scale



horizontal distances of physical locations at different heights above and below eye level also tend to have a maximum at about 3 m and are similar in shape.

Perceived distances in impoverished settings

How, then, do these scale-invariant distributions of distances from the image plane in natural scenes account for the anomalies of visual space summarized in Fig. 1?

When little or no other information is available in a scene, observers tend to perceive objects to be 2-4 m away^{5,6}. In the absence of any distance cues, the likelihood function is flat; the apparent distance of a point in physical space should therefore accord with the probability distribution of the distances of all points in typical visual scenes (equation 1). As indicated in Fig. 3a, this distribution has a maximum probability at about 3 m. The agreement between this distribution of distances in natural scenes and the relevant psychophysical evidence^{1,2,5,6} is thus consistent with a probabilistic explanation of the 'specific distance tendency'.

The similar apparent distance of an object to the apparent distances of its near neighbors in the retinal image (the 'equidistance tendency'⁵) also accords with the probability distribution of the distances of locations in natural scenes. In the absence of additional information about differences in the distances of two nearby locations, the likelihood function is again more or less flat. As a result, the Figure 2 A representative range image taken from one of the wide-field images acquired by laser range scanning. (a) Image generated by the intensity of the laser return, indicated by the corresponding grayscale values. (b) Range image of the same scene; the distance of each pixel is indicated by color coding. Black areas are regions where the laser beam did not return a value.

probability distribution of the differences of the physical distances from the image plane to any two locations in natural scenes should strongly bias the perceived difference in their distances. As the distribution between two locations with relatively small angular separations (black line in Fig. 3b) has a maximum near zero, any two neighboring objects should be perceived to be at about the same distance from the observer. However, at larger angular separations (green line in Fig. 3b), the probability associated with small absolute differences in the distance to the two points is lower than the corresponding probabilities for smaller separations, and the distribution is relatively flatter. Accordingly, the tendency to see neighboring points at the same distance from the observer would be expected to decrease somewhat as a function of increasing angular separation. Finally, when more specific information about the distance difference is present, this tendency should decrease. Each of these tendencies has been observed in psychophysical studies of the 'equidistance tendency'5.

Perceived distances in more complex circumstances

160

The following explanations for the phenomena shown in Fig. 1c–e are somewhat more complex because, unlike the 'specific distance' and 'equidistance' tendencies, the relevant psychophysical observations were made under conditions that involved some degree of contextual visual information. Thus, the relevant likelihood functions are no longer flat. As their form is not known, we used a Gaussian to approximate the likelihood function in the following analyses (or, in determining the influence of the terrain, we obtained the posterior directly).

The probability distribution of physical distances at eye level (black line in Fig. 3c) accounts for the perceptual anomalies in response to stimuli generated by near and far objects presented at this height (Fig. 1c). As shown in Fig. 4a, the distance that should be perceived on this basis is approximately a linear function of physical distance, with near distances being overestimated and far distances underestimated; the physical distance at which overestimation changes to underestimation is about 5–6 m. The effect of these statistics accords both qualitatively and quantitatively with the distances reported under these experimental conditions¹¹.

To examine whether the perceptual observations summarized in Fig. 1d can also be explained in these terms (Fig. 4b), we computed the probability distribution of physical distances of points at different elevation angles of the laser beam relative to the horizontal



Figure 3 Probability distributions of the physical distances from the image plane of points in the range image database of natural scenes. (a) The scaleinvariant distribution of the distances from the center of the laser scanner to all the physical locations in the database (black line). The red line represents the distribution of distances derived from a simple model in which 1,000 planar rectangular surfaces were uniformly placed at distances of 2.5–300 m, from 150 m left to 150 m right, and from the ground to 25 m above the ground (which was 1.65 m below the image center). The sizes of these uniformly distributed surfaces ranged from 0.2 to 18 m. Five-hundred 512×512 images of this model made by a pinhole camera method showed statistical behavior similar to that derived from the range image database for a wide variety of specific values, although with different slopes and modes. The 2.5-m cut-off models the presumed tendency of observers to keep physical objects some distance away; even without this cut-off, however, zero is not the most probable distance but has a significant probability. The model also generated statistical behavior similar to that shown in panels **b** and **c** (not shown). (**b**) Probability distributions of the differences in the physical distances of two locations separated by three different angles in the horizontal plane (vertical separations, which are not shown, showed a similar result). (**c**) Probability distributions of the horizontal distances of physical locations at different heights with respect to eye level. Note that the probability distributions in this and following figures are presented as probability densities.

plane at eye level (that is, along different lines of sight) (Fig. 5). As shown in Fig. 5a, the probability distribution of distances is more dispersed when the line of sight is directed above rather than below eye level, showing a long tail that gradually approaches eye level. The distribution shifts toward nearer distances with increasing absolute elevation angle, a tendency that is more pronounced below than above eye level. A more detailed examination of the distribution within 30 m shows a single salient ridge below eye level (red), extending from ~3 m near the ground to ~10 m at an elevation of 10°; when the line of sight is above eye level, this bias is more dispersed (Fig. 5b). These statistical differences as a function of the elevation angle of the line of sight are even more apparent when the average distances to locations in the scene database are computed as a function of elevation angle (Fig. 5c). The distances of the average physical locations at different elevation angles of the scanning beam form a gentle curve. Below eye level, the height of this curve is rela-



Figure 4 The perceived distances predicted for objects located at eye level, and for objects on the ground. (a) The perceived distances predicted from the probability distribution of physical distances measured at eye level. The solid line represents the local mass mean of the probability distribution obtained by multiplying the probability distribution in Fig. 3c (black line) by a Gaussian likelihood function of distances with a standard deviation of 1.4 m. The dashed line represents the equivalence of perceived and physical distances for comparison. When the standard deviation of the likelihood function was increased, the predicted distances showed greater deviation from the physical distances, and more closely approximated the known psychophysical data (1.4 m is therefore a conservative value). (b) The perceived distances of objects on the ground in the absence of other information predicted from the probability distribution in Fig. 5a. The likelihood function at an angular declination α was a Gaussian function: ~exp(– $(\alpha - \alpha_0)^2/2\Sigma^2$), where $\alpha_0 = \sin^{-1}(H/R)$, $\Sigma = 8^{\circ}$, R = radial distance and H =1.65 m). The prior was the distribution of distance at angular declinations within $[\alpha - 8^\circ, \alpha + 8^\circ]$. The ground in the diagram is a horizontal plane 1.65 m below eye level. The predicted perceptual locations of objects on the ground are indicated by the solid black line, which is slanted toward the observer.



Figure 5 Probability distribution of physical distances at different elevation angles. (a) Contour plot of the logarithm of the probability distribution of distances at elevation angles indicated by color coding. (b) Blowup of **a** showing the probability distribution of distances within 30 m in greater detail. (c) The average distance as a function of elevation angle, based on the data in **a**. The vertical axis is the height relative to eye level; the horizontal axis is the horizontal distance from the image plane. The curve below eye level, if modeled as a piece-wise plane, would have a slant of about 1.5° from a distance of 3–15 m, and about 5° from 15–24 m away.

tively near the ground for closer distances, but increases slowly as the horizontal distance from the observer increases. For elevations of the line of sight above eye level, the height of this average physical location is greatest at the highest elevation angle examined and decreases as the horizontal distance from the observer increases (the sky was automatically excluded). If the portion of the curve at heights below eye level in Fig. 5c is taken as an index of the average ground, it is apparent that the average ground is neither a horizontal plane nor a plane with constant slant, but a curved surface that is increasingly inclined toward the observer as a function of horizontal distance.

These statistical characteristics of distance as a function of the elevation of the line of sight can thus account for the otherwise puzzling perceptual effects shown in **Fig. 1d**. The perceived location of an object on the ground without much additional information about its actual distance varies according to the declination of the line of sight: objects appear closer and higher than they really are as a function of this angle^{12,13}. The apparent location of an object predicted by the probability distributions in **Fig. 5** is increasingly higher and closer to the observer as the declination of the line of sight decreases, in agreement with the relevant psychophysical data¹³ (**Fig. 4b**).

The effects of terrain

To understand the effect of the terrain (Fig. 1e), we examined the correlation of local variations in the terrain with the structure of the rest of the ground (Fig. 6). To this end, we computed the probability distributions of the distances to physical locations at all elevation angles when all the physical locations below eye level that fell within a restricted range of elevation angles (-30.8 to -26.8° in Fig. 6) were either more than 0.15 m below the ideal ground (a dip in the horizontal plane 1.65 m below eye level) or more than 0.15 m above the ideal ground (a hump).

Compared with the probability distribution in Fig. 5b, which includes all possible terrain profiles, the probability distribution of horizontal distances shifts toward greater distances when the ground dips locally. Conversely, when there is a local hump, the probability distribution shifts toward lesser distances (Fig. 6b). Moreover, the average ground surface is farther below eye level when a local dip is present, whereas the opposite is true when a local hump is present (Fig. 6c).

These observations show that local variations in the terrain exert a global influence on the statistical configuration of the rest of the ground. If visual space is generated probabilistically, then this robust relationship should have corresponding perceptual consequences. In fact, these changes in the probability distribution of distances from the image plane arising from particular local configurations of ground accord qualitatively with overestimation of physical distance when the ground is disrupted by a dip, and underestimation when the ground exhibits a local hump (Fig. 1e).

To examine the influence of the terrain on distance perception quantitatively, we determined the statistical relationship between the horizontal distances to locations on the ground at particular elevation angles and the terrain intervening between the locations and the observer (Fig. 7a). The surface formed by all the locations below eye level along each vertical scan was defined as the ground. By sampling all ground surfaces that showed a particular undulation, we obtained the probability distribution of horizontal distances to points on the ground at given elevations when either a dip or hump intervened (Fig. 7b).

As already described in qualitative terms, when there is a dip in the ideal ground plane, the probability distribution shifts toward larger horizontal distances. Fig. 7c shows more specifically the distances that should be perceived based on the probability distribution of the horizontal distances to locations in physical space. The expected overestimation when the terrain deviates negatively from a



Figure 6 Probability distributions of physical distances below eye level when the terrain has a local dip or a hump. (a) Contour plot of the logarithm of the probability distribution of distances when all the physical locations at elevations within $[-30.8^{\circ}, -26.8^{\circ}]$ were at least 0.15 m below the ideal ground (defined as 1.65 m below eye level), thus forming a dip. (b) Similar plot when all the physical locations at elevations within $[-30.8^{\circ}, -26.8^{\circ}]$ were at least 0.15 m below the ideal ground (defined as 1.65 m below eye level), thus forming a dip. (b) Similar plot when all the physical locations at elevations within $[-30.8^{\circ}, -26.8^{\circ}]$ were at least 0.15 m above the ideal ground, thus forming a hump. (c) Average profile of the ground obtained from the probability distributions in a (green line) and b (blue line), respectively. For comparison, the black line is the average ground derived from the probability distribution of all the range measurements below eye level.

more or less flat ground plane is about 0.2–0.6 m and increases slightly with increasing physical distance. These predicted values accord with psychophysical observations¹⁴.

The same approach was used to explore the predicted consequences of a hump in the terrain. As indicated in **Fig. 8a**, when there is positive deviation in the ground plane, the probability distribution shifts toward smaller horizontal distances. These data indicate that the perceived distance in the presence of a local hump should be slightly but consistently underestimated (by 0.1–0.5 m) in comparison with the physical distance, as shown in **Fig. 8b**. Again this prediction is in agreement with judgments of distance under the relevant circumstances (B. Wu, Z.J. He, & T.L. Ooi, *J. Vision Abstr.*, 2, 513a, 2002; Y.L. Yarbrough *et al.*, op. cit., 625a).

DISCUSSION

When projected onto the retina, three-dimensional (3D) spatial relationships in the physical world are necessarily transformed into two dimensions in the image plane. As a result, the physical sources underlying any geometrical configuration in the retinal image are uncertain: a multitude of different scene geometries could underlie any particular configuration in the image. This uncertain link between retinal stimuli and physical sources presents a biological dilemma, as an observer's visually guided behavior must accord with real-world physical sources.

Given this quandary, we used the phenomenology of visual space to test the idea that the uncertain relationship between images and sources could be addressed by a probabilistic strategy. If physical and perceptual space are indeed related in this way, then the characteristics of human visual space (using perceived distance as the simplest and most general index) should accord with the probability distributions of natural scene geometry. Observers would be expected to perceive objects in positions substantially and systema-tically different from their physical locations when countervailing empirical information is not available. When other contextual information is available, the perceived locations would be predicted by the altered probability distributions of the possible sources of the stimuli. Using a database of range images, we show that the phenomena illustrated in Fig. 1 can all be rationalized in this framework.

The fact that these otherwise puzzling features of visual space can be understood in terms of the probabilistic relationship between images and their possible sources accords with the successful explanation of many other visual phenomena in this way^{16,17}. Although the variety of discrepancies between physical measurements and the corresponding percepts seems 'maladaptive' on the face of it, given the problem of stimulus uncertainty that vision must inevitably contend with, this probabilistic strategy ensures routinely successful behavior in typical visual environments. The anomalies of perceived distance noted in numerous studies over the last century are evidently manifestations of this process.

Other approaches to rationalizing visual space

A number of studies have proposed that visual space is isomorphic with Euclidean space, a Riemann space with constant curvature or an affine space^{26–29}. Others have suggested that visual space is computed using information derived from perspective, texture gradients, binocular disparity or motion parallax by more or less independent visual processing modules^{1–3}. Perhaps the most influential theory of visual space has been put forward by Gibson³⁰, who argued that since human beings are terrestrial, the ground is the key factor in determining the perception of space. In this conception, a two-dimensional (2D) frame of reference built on the terrestrial surface is taken as the basis of visual space.

If visual space is indeed generated by a probabilistic strategy, then explaining the relevant perceptual phenomenology will require knowledge of the statistical properties of natural visual environments with respect to observers. Without this empirical information, any theoretical explanation of apparent distance is likely to be inadequate. For example, since the relationship of images and their sources is necessarily probabilistic, the assumption that visual space corresponds to a Riemann space of constant curvature or an affine space is unlikely to

ARTICLES



Figure 7 Statistical explanation of the effect of a dip in the ground-plane on perceived distance. (a) Diagram showing how the terrain in the range images was analyzed, and defining the relevant symbols. In this example, there is a dip at elevations in $[\alpha - \Delta \alpha_1, \alpha + \Delta \alpha_1]$ in an otherwise more-or-less flat ground plane. (b) The graphs show the probability distributions of the horizontal distances of the physical locations at elevations on the ground within $[\alpha + \delta - \Delta \alpha_2, \alpha + \delta + \Delta \alpha_2]$, given a dip in the otherwise flat terrain intervening between the locations and the observer $(\Delta \alpha_1 = 2^\circ, \delta = 3.6^\circ \text{ and } \Delta \alpha_2 = 0.58^\circ)$. The dip is closer to the observer in the the left panel ($\alpha = -26^\circ$), and farther away in the right panel ($\alpha = -14.4^\circ$). Black line, distribution when the ground is flat; red line, distribution when the ground is disrupted by the dip. (c) Perceived distances predicted for an object on the ground in presence of a dip, computed on the basis of probability distributions like those in **b**. The dashed line indicates the equivalent relationship between the perceived and physical distance of objects on flat ground; the solid line represents the predicted relationship when a dip is present.

account for the space that humans actually see. Visual space generated probabilistically will necessarily be a space in which perceived distances are not a simple mapping of physical distances; on the contrary, apparent distance will always be determined by the way all the available information at that moment affects the probability distribution of the gamut of the possible sources of any physical point in the scene. Although Gibson's emphasis on the terrain is a step in this direction, it accounts for only a small fraction of the empirical information that is typically available to the human visual system, and can thus explain little of the phenomenology illustrated in Fig. 1.

Other studies of natural scene statistics

Whereas the statistics of natural environments have been studied in considerable detail³¹, the statistics of the physical world in relation



Figure 8 Statistical explanation of the effect of a hump in the ground plane on perceived distance. (a) Probability distribution of horizontal distances given a hump that is relatively closer to the observer in left panel ($\alpha = -26^\circ$), and relatively farther away in the right panel ($\alpha = -14.4^\circ$). Black line, distribution when the ground is flat; red line, distribution when the ground is disrupted by a hump. (b) Predicted distances of an object on the ground in the presence of a hump, based on probability distributions like those in **a**. Dashed line and solid line defined as in **Fig. 7c**.

to the observer have not. The general assumption underlying most studies of natural images is that visual systems must encode image features with optimal efficiency. Accordingly, the statistics examined in such studies have focused on the probabilistic relationship of elements in the image plane and the pertinence of these relationships to efficient coding strategies. Although this approach has been both sensible and fruitful^{23,32–34}, the statistical relations among elements in the image plane are not directly informative about the physical scene geometry in relation to an observer and thus are not immediately pertinent to rationalizing the characteristics of human visual space.

METHODS

Range image database. The database was acquired using a high-precision range scanner (LMS-Z210 3D Laser Scanner, Riegl USA Inc.) mounted on a tripod and leveled in the horizontal plane at a height of 1.65 m. Although we avoided placing the scanner directly in front of large objects less than 3 m away (which would have blocked the majority of the scene), the site of the scanner was otherwise unconstrained. This device detects surfaces at distances of 2–300 m with an accuracy of ± 25 mm at a resolution of 0.144°. Twenty-three images were taken in fully natural settings that included trees and landscapes in nearby Duke Forest, and 51 images were taken in outdoor settings that included both natural and constructed objects at Duke University campus in stable daytime conditions³⁵ (see also ref. 36).

The raw images, all of which were used in the present analysis, comprised $333^{\circ} \times 80^{\circ}$. The edges of the images were excluded from the analysis, leaving $\sim 326^{\circ}$ horizontally $\times 72^{\circ}$ vertically; any location that didn't have a laser return (such as the sky) or happened to contain a moving object (such as a car) was also excluded. Although the resolution of the scanner was relatively low compared to the human visual system, given the scale-invariance of natural scenes^{24,25,31}, these images provide a reasonable sample of the normal visual environment.

Obtaining the distance distributions. The distribution of distances in the images was obtained by counting the frequency of occurrence of all the measured ranges (that is, the radial distance to the center of the scanner). The bin size was 20 cm. The distribution of horizontal distances at eye level (defined as 1.65 m, the height of the center of the scanner) was obtained by counting the samples within $\pm 2^{\circ}$ relative to the horizontal plane at this height. The distribution of horizontal distances within or at different heights above or below eye level was similarly obtained by counting samples taken from physical locations within the corresponding spaces centered 0.8, 1, 1.2, 1.4 or 1.6 m above or below eye level.

To obtain the distribution of the differences between the distances of any two locations, we randomly sampled pairs of locations that were separated either horizontally or vertically and counted the occurrences of the absolute difference of their distances from the image plane, using a bin size of 10 cm. To examine how distances are distributed at any elevation angle with respect to eye level, we tallied the distances of all the physical locations in the scenes that spanned a particular elevation angle relative to the horizontal plane at the level of the scanner.

Finally, a scaling transform was constructed to examine how natural scene geometry changes as a function of scale. The 3D coordinates at the (i, j) pixel in the n^{th} scale were taken as the mean of 3D Euclidean coordinates at (2i, 2j), (2i - 1, 2j), (2i, 2j - 1) and (2i - 1, 2j - 1) in the $(n - 1)^{\text{th}}$ scale (the original range image is at the zero-order scale). Four scales, including the scale of the original images, were tested.

Measuring the influence of the terrain. We identified all possible ground surfaces in the database that had a particular vertical fluctuation (a dip or hump), and then determined the probability distribution of the horizontal distance of locations on the ground at particular elevation angles of scanning beam, given the vertical fluctuation intervening between the locations and the observer. A dip or hump was defined as a profile in which all the physical locations at elevations in $[\alpha - \Delta \alpha_1, \alpha + \Delta \alpha_1]$ were at least 0.1 m below or above the ideal ground (defined as 1.65 m below eye level; Fig. 7a).

The rest of the ground was accepted as flat if the mean of all relevant physical locations fell within 0.15 m above or below the ideal ground with a standard deviation less than 0.25 m. This constraint on the rest of the ground adjacent to a dip or hump necessarily reduced the number of samples obtained, but approximated the experimental conditions in which the observations illustrated in Fig. 1e were obtained. Less stringent criteria yielded similar results, whereas more stringent criteria greatly reduced the total number of samples.

For vertical scans below eye level that had dips/humps meeting these criteria, we tallied the frequency of occurrences of the horizontal distances of locations at elevations in $[\alpha + \delta - \Delta\alpha_2, \alpha + \delta + \Delta\alpha_2]$ on the ground, where α was $-14.4^\circ, -17.3^\circ, -20.2^\circ, -23^\circ, -26^\circ, -28.8^\circ, -31.7^\circ$ or -34.6° ($\Delta\alpha_1 = 2^\circ, \delta = 3.6^\circ, \Delta\alpha_2 = 0.58^\circ$). Thus, the extent of the deviation along the ground was less than 4°; moreover, the locations whose horizontal distances were tallied for later analysis were not within the dip or hump, but always at least 3.6° away from the nearest boundary of the deviation. The bin size for horizontal distances was 10 cm. The total number of samples obtained in this way ranged from ~2,500 to ~60,000.

Predicted percepts. The local mass mean of the relevant probability distribution was taken as the predicted distance that observers would be expected to perceive³⁷. The local loss function in this computation was a negative Gaussian function with a standard deviation of 0.2 m.

ACKNOWLEDGMENTS

We thank C. Howe, F. Long, S. Nundy, D. Schwartz and J. Voyvodic for useful comments, and M. Williams for help with the art. This project was supported by the National Institutes of Health and the Geller endowment.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Received 7 January; accepted 25 March 2003 Published online 18 May 2003; doi:10.1038/nn1059

- Hershenson, M. Visual Space Perception: a Primer (MIT Press, Cambridge, Massachusetts, 1999).
- Gillam, B. The perception of spatial layout from static optical information. in *Perception of Space and Motion* (eds. Epstein, W. & Rogers, S.) 23–67 (Academic, New York, 1995).
- Sedgwick, H.A. Space perception. in *Handbook of Perception and Human Performance* Vol. 1 (eds. Boff, K.R., Kaufman, L. & Thomas, J. P.) 21.1–21.57 (Wiley, Toronto, 1986).
- Loomis, J.M., Da Silva, J.A., Philbeck, J.W. & Fukusima, S.S. Visual perception of location and distance. *Curr. Dir. Psych. Sci.* 5, 72–77 (1996).
- Gogel, W.C. Equidistance tendency and its consequences. *Psychol. Bull.* 64, 153–163 (1965).
- Owens, D.A. & Leibowitz, H.W. Oculomotor adjustments in darkness and the specific distance tendency. *Percept. Psychophys.* 20, 2–9 (1976).
- Epstein, W. & Landauer, A.A. Size and distance judgments under reduced conditions of viewing. *Percept. Psychophys.* 6, 269–272 (1969).
- Gogel, W.C. & Tietz, J.D. A comparison of oculomotor and motion parallax cues of egocentric distance. Vis. Res. 19, 1161–1170 (1979).
- Morrison, J.D. & Whiteside, T.C.D. Binocular cues in the perception of distance to a point source of light. *Perception* 13, 555–566 (1984).
- Foley, J.M. Binocular distance perception: egocentric distance tasks. J. Exp. Psychol. Hum. Percept. Perform. 11, 133–149 (1985).
- Philbeck, J.W. & Loomis, J.M. Comparison of two indicators of perceived egocentric distance under full-cue and reduced-cue conditions. J. Exp. Psychol. Hum. Percept. Perform. 23, 72–85 (1997).
- Wallach, H. & O'Leary, A. Slope of regard as a distance cue. Percept. Psychophys. 31, 145–148 (1982).
- Ooi, T.L., Wu, B. & He, Z.J. Distance determined by the angular declination below the horizon. *Nature* **414**, 197–200 (2001).
- Sinai, M.J., Ooi, T.L. & He, Z.J. Terrain influences the accurate judgment of distance. *Nature* **395**, 497–500 (1998).
- Meng, J.C. & Sedgwick, H.A. Distance perception mediated through nested contact relations among surfaces. *Percept. Psychophys.* 63, 1–15 (2001).
- Knill, D.C. & Richards, W. Perception as Bayesian Inference (Cambridge Univ. Press, Cambridge, 1996).
- Purves, D. & Lotto, B. Why We See What We Do: an Empirical Theory of Vision (Sinauer, Sunderland, Massachusetts, 2003).
- Kersten, D. High-level vision as statistical inference. in *The New Cognitive Neurosciences*. 2nd edn. (ed. Gazzaniga, M.S.) 353–363 (MIT Press, Cambridge, Massachusetts, 1999).

ARTICLES

- 19. Geisler, W.S. & Kersten, D. Illusions, perception and Bayes. *Nat. Neurosci.* 5, 508–510 (2002).
- 20. Belhumeur, P.N. A Bayesian approach to binocular stereopsis. Intl. J. Comp. Vision 19, 237–260 (1996).
- Bloj, M.G., Kersten, D. & Hurlbert, A.C. Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature* 402, 877–879 (1999).
- Weiss, Y., Simoncelli, E. & Adelson, E.H. Motion illusions as optimal percepts. *Nat. Neurosci.* 5, 598–604 (2002).
- Geisler, W.S., Perry, J.S., Super, B.J. & Gallogly, D.P. Edge co-occurrence in natural images predicts contour grouping performance. *Vis. Res.* 41, 711–724 (2001).
- 24. Mumford, D. & Gidas, B. Stochastic models for generic images. Q. J. Appl. Math. 59, 85-111 (2001).
- Lee, A.B., Mumford, D. & Huang, J. Occlusion models for natural images: a statistical study of a scale-invariant dead leaves model. *Int. J. Comput. Vision* 41, 35–59 (2001).
- 26. Luneberg, R.K. *Mathematical Analysis of Binocular Vision* (Princeton Univ. Press, Princeton, New Jersey, 1947).
- 27. Indow, T. A critical review of Luneburg's model with regard to global structure of visual space. *Psychol. Rev.* **98**, 430–453 (1991).
- 28. Wagner, M. The metric of visual space. *Percept. Psychophys.* **38**, 483–495 (1985).

- Todd, J.T., Oomes, A.H.J., Koenderink, J.J. & Kappers, A.M.L. On the affine structure of perceptual space. *Psychol. Sci.* 12, 191–196 (2001).
- 30. Gibson, J.J. The Perception of the Visual World (Houghton Mifflin, Boston, 1950).
- Simoncelli, E.P. & Olshausen, B.A. Natural image statistics and neural representation. Annu. Rev. Neurosci. 24, 1193–1216 (2001).
- 32. Olshausen, B.A. & Field, D.J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607–609 (1996).
- Vinje, W.E. & Gallant, J.L. Sparse coding and decorrelation in primary cortex during natural vision. *Science* **287**, 1273–1276 (2000).
 Sigman, M., Cecchi, G.A., Gilbert, C.D. & Magnasco, M.O. On a common circle:
- 34. Sigman, M., Cecchi, G.A., Gilbert, C.D. & Magnasco, M.O. On a common circle: natural scenes and Gestalt rules. *Proc. Natl. Acad. Sci. USA* 98, 1935–1940 (2001).
- 35. Howe, C.Q. & Purves, D. The statistics of range images can explain the anomalous perception of length. *Proc. Natl. Acad. Sci. USA* 99, 13184–13188 (2002).
- 36. Huang, J., Lee, A.B. & Mumford, D. Statistics of range images. Proc. IEEE Conf. CVPR 1, 324–331 (2000).
- 37. Brainard, D.H. & Freeman, W.T. Bayesian color constancy. J. Opt. Soc. Am. A 14, 1393–1411 (1997).