Perception, 2000, volume 29, pages 69 – 79

DOI:10.1068/p2921

Direct measurement of the curvature of visual space

Jan J Koenderink, Andrea J van Doorn
Department of Physics and Astronomy, Universiteit Utrecht, Princetonplein 5, 3584 CC Utrecht, The Netherlands; e-mail: j.j.koenderink@phys.uu.nl

Joseph S Lappin
Vanderbilt Vision Research Center, Vanderbilt University, Nashville, TN 37240, USA; e-mail: joe.lappin@vanderbilt.edu
Received 25 March 1999, in revised form 23 August 1999

Abstract. We consider the horizontal plane at eye height, that is all objects seen at the horizon. Although this plane visually degenerates into a line in the visual field, the 'depth' dimension nevertheless gives it a two-dimensional structure. We address the problem of intrinsic curvature of this plane. The classical geometric method is based on Gauss's original definition: The angular excess in a triangle equals the integral curvature over the area of the triangle. Angles were directly measured by a novel method of exocentric pointing. Experiments were performed outside, in the natural environment, under natural viewing conditions. The observers were instructed not to move from a set location and to maintain eye height, but were otherwise free to perform eye, head, and body movements. We measured the angular excess for equilateral triangles with sides of 2 – 20 m, the vantage position at the barycenter. We found angular excesses and deficits of up to 30°. From these data we constructed the metric. The curvature changes from elliptic in near space to hyperbolic in far space. At very large distances the plane becomes parabolic.

1 Introduction

People routinely commit appreciable and typically systematic errors when asked to estimate geometrical properties of their environment from a fixed vantage point (Battro et al 1976, 1978; Wagner 1985). This is not surprising in view of the fact that these geometrical properties are only implicitly specified by the retinal image(s). Conventionally one distinguishes between the visual field and the visual world (Pirenne 1970). The geometry of the visual field is that of the simultaneous presence of variously colored patches. The visual world involves an additional 'depth' dimension. Depth properties are specified by binocular disparities and optic flow, and by a varied bouquet of 'depth cues'. The former involve problems of 'correspondence' (simultaneously between the retinal images in the case of binocular disparities, successively in the case of optic flow), and various prior assumptions, but no interpretations of the retinal illuminance structures. The latter involve such interpretations.

Our experiments were performed in the natural environment, an open field under bright daylight conditions with everything in plain sight. No doubt the observers use large parts of the visual field in order to estimate the depth of objects that appear on their horizon. When we mention 'points' or 'lines', these have to be understood as mere formal abstractions, not realities. Since we are primarily interested in the 'depth' dimension we limit our investigations to the horizontal plane through the eyes. This plane appears as a line in the visual field—the so called 'horizon' of the observer. Of course, we have in mind geometry here; in reality the horizon need not be specified optically, nor do we confine our research to points. The azimuth is measured along the horizon and is a dimension that is measured in the visual field. Depth is a 'quality' of points on the horizon.

Visually perceived distance, that is 'depth', usually has been studied under tightly controlled conditions in the laboratory. Such conditions typically involve extreme stimulus

reduction, say a few luminous points in the dark (nothing else visible, intensities carefully balanced); extreme response reduction, say a forced binary choice; and extreme reduction of the viewing conditions, say fixed head (probably a biteboard, at least a chinrest) and possibly fixation instructions (Blank 1953, 1959; Indow et al 1962a, 1962b, 1963). Under natural conditions most of these constraints are violated. 'Natural conditions' are indeed the opposite of stimulus reduction; restricting viewing conditions is also at odds with the 'natural setting'. In our case we constrained location and eye height, but left observers quite free to make eye, head, and (torsional) body movements; even turning on the spot was allowed. We also designed a task that felt much more 'natural' to the observers than the forced binary choice.

From an applied perspective, such studies under natural conditions are obviously desirable. The scientific value, however, might be less obvious because such studies are less controlled (or controllable) than the laboratory studies. The potential value of studies of depth perception in natural environments arises from the likelihood that observers may change tactics when granted additional freedom and optical structure: Strictly controlled laboratory studies cover only part of the story.

The larger part of the theory of 'depth perception' is based upon a geometrical analysis of bicentric perspective (Helmholtz 1867; Luneburg 1947). This is only natural, in view of the fact that binocular disparity is about the only cue left to the observer in the severely controlled conditions (a few luminous points in the dark). There appears to be some consensus that the geometry of visual space (mainly that of the horizontal plane at eye height) is governed by a Riemannian metric (that is a global field of local metrics, ie the metric may change from point to point) and certainly is not Euclidean; possibly it is a Riemannian space of constant negative curvature (Blank 1953, 1959; Indow et al 1962a, 1962b, 1963). This is by no means the only theory, but its pervasiveness suggested to us that it might be of some interest to attempt to assess the validity of such a description in more natural conditions. We do not assume that the space is of constant curvature, since there are various reasons for doubt, even under controlled conditions. We do assume that the space is Riemannian, and that it is isotropic. The latter assumption is natural when the observer is allowed the freedom to turn on the feet and perform arbitrary head and body rotations about the vertical; indeed, one might say that isotropy is then enforced by the freedom—no visual direction is singled out. Under controlled conditions this is different: clearly the 'straight ahead' direction is special then. Under these assumptions we propose to determine the intrinsic curvature as a function of distance from the observer.

The notion of intrinsic curvature was introduced by Carl Friedrich Gauss in a seminal (in geometry) paper (see Gauss 1880) dating back to 1837. Although the theory is arcane to psychophysicists, some of its consequences are relatively easy to grasp. 'Intrinsic' curvature differs from 'extrinsic' curvature: a flat piece of paper remains intrinsically flat when you roll it into a cylinder or cone although its extrinsic curvature may change. Intrinsic curvature is evident from certain non-Euclidean properties of figures. One example treated by Gauss concerns the sum of interior angles of a triangle made up of geodesics (shortest arcs, the Euclidean straight lines). In the Euclidean plane the sum is 180°; any deviation (called 'angular excess') reveals an intrinsic curvature of the space. The excess is proportional to the area of the triangle; the excess per unit area ('specific angular excess') is a measure of the intrinsic curvature. Notice that its physical dimension is one over length squared, thus the inverse of the square root of the (absolute value of the) specific angular excess has the dimension of length. Gauss showed that it simply equals the radius in the simple case of a spherical surface.

Here is a very simple example: Consider the geographical triangle Amsterdam (latitude 52° 21′ north, longitude 4° 54′ east), Boston (42° 20′ north, 71° 5′ west), and Cape Town (35° 56′ south, 18° 22′ east). This triangle has an area of 38 700 000 km^2.

From Amsterdam, Boston is 5568 km at 20° 30′ north of west, Cape Town is 9914 km at 10° 52′ east of south, etc. The angle sum of the triangle is 234° 28.3′: This differs appreciably from the Euclidean 180°; thus the Earth is significantly curved. The specific spherical excess is 0.00507 s of arc km^{-2}: According to Gauss this equals the curvature averaged over the triangle. One over the square root of this curvature equals 6378 km which is the radius of the globe. Thus we obtain the radius of the Earth from the directions of the cities as seen from the others and from the area of the triangle. Notice that the specific curvature is small, for instance the excess in an equilateral triangle with sides of 1 km is only 0.002 s of arc: locally the Earth is quite flat.

Gauss was asked to perform geodesic measurements for his government (probably to increase the accuracy of military maps) and he used the occasion to measure angular excess of large triangles between mountain tops. He found values not significantly different from zero and concluded that physical space is flat for such extents. We decided to follow his lead and determine the angular excess of triangles of various sizes in visual space. From such measurements it is a simple matter to find the intrinsic curvature of visual space.

We implemented a simple, direct method to measure angles in visual space. To the best of our knowledge this has not been attempted before [closest is perhaps Ellis et al (1991)], although some investigators have used judgments of perpendicularity. Perpendicularity is special, however, because it assumes an angular metric. (For instance, perpendicularity is undefined in the projective or affine plane.) We only required our observers to indicate the direction in visual space from one point to another, neither point coincident with the egocenter. Pointing from the egocenter is generally known as 'aiming' as in pointing a rifle or a telescope. Taking aim, or '*egocentric pointing*', is a task that can be fully performed in the visual field and involves no depth. *Exocentric* pointing involves aiming from a location distinct from that of the observer. A common case of judging exocentric pointing is in perceiving who looks at whom at a cocktail party. We simply constructed a pointer that could be rotated by remote control (a radio link) from the observer's position. The observer sees both a target and the pointer and is instructed to rotate the pointer such that—from the position of the pointer—the pointer is aimed at the target. People find this a natural enough task to perform. By pointing from A to B we essentially determine the (visual) direction AB as seen from C, where A, B, and C are all distinct. By this method we may determine angular excess of any triangle in visual space.

In this experiment we propose to find the integral curvature (Gauss's "Curvature Integra"; it equals the spherical excess) over the areas of equilateral triangles with the observer situated at the barycenter. Processing of such pointing data should be trivial since we merely need to follow Gauss's example. The results do not imply a check on the Riemannian property of visual space: *Any* result is interpreted as a space variant intrinsic curvature. However, we directly address the problem of whether visual space—granted its Riemannian character—is a space of constant curvature. This is of some interest, since it remains unresolved in the literature (Indow 1991). The point is of conceptual interest because only spaces of constant curvature allow congruence, or the free movement of rigid bodies. This is indeed the heuristic behind Luneburg's hypothesis of constant curvature (Luneburg 1947).

2 Methods

We used a radiographically controlled pointer mounted on a tripod (see figure 1). The pointer was a white cube of 25 cm edge length, pierced with an arrow, shaft thickness 1.5 cm, total length 125 cm. The arrow stuck out 50 cm at both sides. The arrow head was 6 cm long and 3 cm in diameter. The cube was painted white, the arrow bright orange.

Figure 1. Photograph of the equipment. Notice that this is a 'mock-up': In the actual experiment distances were generally larger, a smaller sphere was used, the bottom halves of the tripods were lost in knee-high weeds and the horizon was more distant. The photograph provides a good idea of the target (sphere at the right), the pointing device (the cube pierced with an arrow in the middle), and the task of the observer (at the left). The observer held a transmitter and thus controlled the pointing device directly while looking back and forth between target and pointer. Notice the absence of any head-constraining device (biteboard or headrest): The observer was even free to turn on the spot. Also notice that the experiment was not conducted in the dark and that the target and pointer were not isolated luminous points but offered a number of monocular cues next to binocular disparity. In the Luneburg – Blank tradition the paradigm is insufficiently reduced and controlled. However, it is perhaps somewhat closer to natural conditions.

The transmitter was a standard Futuba F16 and yielded precise control over the orientation of the pointer. Targets were bright-orange spheres mounted on thin steel rods, themselves mounted on tripods. Sphere size was different for the various distances: the spheres subtended 20 ± 4 min of arc. Both pointer and targets were always at the eye height of the observer.

Experiments were performed in an unkept meadow with weeds growing up to knee and even waist height. Thus, it was typically not possible to see the base of the tripods meet the ground plane. The meadow was an open area with some distant trees and buildings visible, but no obvious landmarks. During the experiments we met with bright weather, mostly blue sky with less than $1/8$ cloud cover. In the course of the experiments the sun's direction changed appreciably.

The observer was instructed to keep the feet on a small platform and to stand upright at all times. The observer was permitted eye movements, head rotations, and torsions at the waist, even changes in placement of the feet (turning on the feet). The explicit instruction was to point the arrow to the target, ie to point it in such a way that someone at the pointer and using it as an aiming device would agree with the setting.

›

Since the observers were not at the pointer (this is exocentric pointing, not aiming), this was not easy, though most observers found the task a quite 'natural' one. At no time were the observers informed of their settings. Three observers participated in the experiment (the authors). All had normal acuity and binocular stereo vision.

Layout of positions was done with conventional geodesic aids (measuring tape and theodolite). Locations were accurate to 10 cm, pointer orientations to about 1°. This turned out to be sufficient, given the repeatability of the observer's settings.

3 Experiment

The experiment itself was straightforward. The observer was positioned at the barycenter of equilateral triangles of various edge lengths. The pointer was placed upon one vertex, the targets on another vertex. We always averaged over settings in which target and pointer were interchanged, thus assuming left/right symmetry. This technique obviates the need for an absolute calibration of pointer orientation, thus increasing accuracy. We also had the observer point at himself or herself: this effectively calibrates the pointer orientation absolutely. Given this absolute calibration we were in a position to check on left/right symmetry. We found no traces of left/right asymmetries throughout the experiments. We averaged over symmetrical cases in order to simplify later analysis. The data are then simply pointer orientation as a function of the size of the triangles.

Since the target and pointer were at an angular distance of 120 deg on the horizon, all observers found it necessary to make body and head turns in order to look back and forth between them.

The essential pointing geometry is shown in figure 2. Here the observer at V points the pointer P at the target T. The veridical pointing angle φ depends on the arc ϑ between the pointer and the target and on the ratio of their distances r_P/r_T:

$$\tan\varphi = \frac{\sin\vartheta}{r_P/r_T - \cos\vartheta}\,.$$

The pointer is illustrated in two orientations in figure 3. Notice that—apart from binocular stereo—various monocular cues are available that might reveal the orientation of the pointer. For instance, the ratios $\alpha/\bar{\alpha}$ or $\lambda/(1-\lambda)$ could be used, etc. Estimates suggest that the monocular cues probably dominate at all distances with the possible exception of the closest one. In this experiment $\vartheta = 120°$ and $r_P/r_T = 1$, thus the veridical pointing angle φ is 30°.

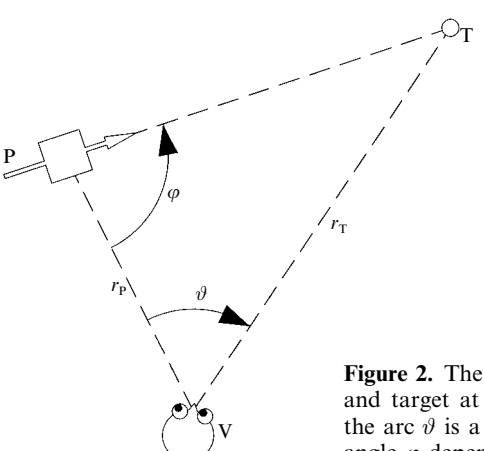

Figure 2. The pointing geometry. Observer at V, pointer at P, and target at T are in the horizontal plane at eye height. Thus the arc ϑ is a distance along the horizon. The veridical pointing angle φ depends on the ratio of distances and this arc.

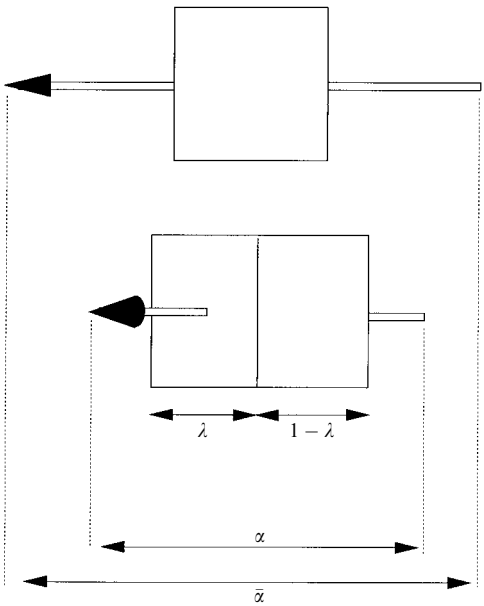

Figure 3. The pointer as it appears (schematically) to the observer.

4 Results

The data are expressed in terms of the angle subtended by the visual direction to the pointer and the exocentric pointing direction. Notice that the veridical value is 30°. We found that the actual values differ significantly and systematically from this value. In figure 4 we plot the actual values as a function of the size of the triangles. Notice that the results obtained for the three observers are quantitatively and qualitatively very similar, though different in detail: straight correlation (figures) yields R^2 values of 0.98 for AD – JK, 0.69 for JK – JL, and 0.74 for AD – JL. Linear regression reveals slopes significantly different from unity and offsets different from zero: deviations for AD are 75% those of JK with a 2.3° offset; those of JL are 48% those of AD with 13.8° offset (figure 5).

An intuitively satisfactory way to represent the results is to use the fact that two points with directions attached to them define a unique circular arc. When we draw such arcs for the triangular configuration (figure 6) we notice that triangles in the near

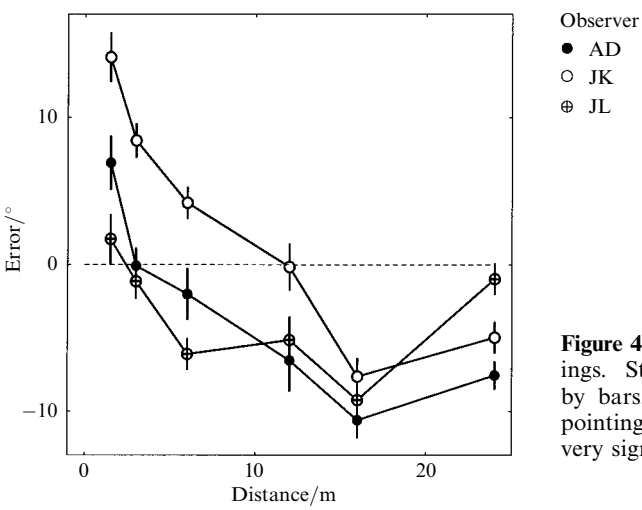

Figure 4. Systematic deviations of the pointings. Standard deviations are indicated by bars. The dotted line is the veridical pointing. Notice that all observers commit very significant pointing errors.

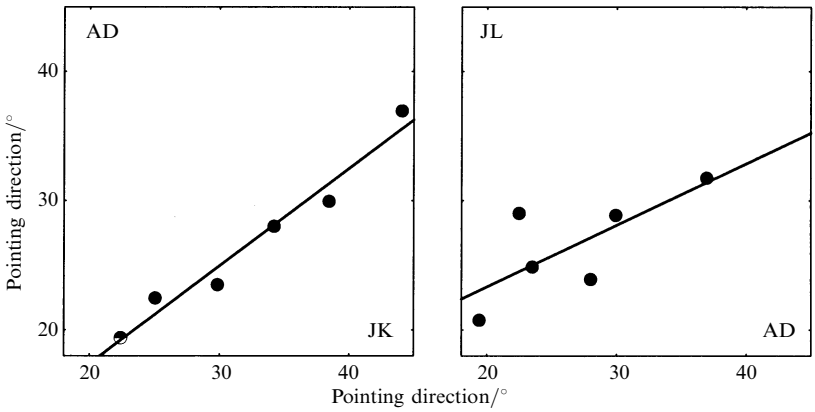

Figure 5. Scatterplots of the average pointings of individual observers.

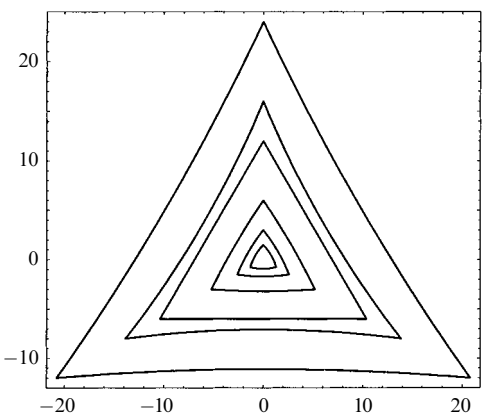

Figure 6. Raw settings of JK visualized in terms of circular arcs. The units shown are meters.

field look 'inflated' (thus there will be an angular excess) whereas those in the far field look 'deflated' (thus there will be an angular deficit). Without any more calculation, the raw data thus reveal that the optical space is elliptic (positively curved) in the near zone but hyperbolic (negatively curved) in the far zone.

Scatter in the settings is about 1.5° which appears to indicate remarkably accurate performance.

4.1 *Determination of the intrinsic curvature*

In order to determine the intrinsic curvature we proceed by finding the integral curvatures for the triangles (Gauss's "Curvatura Integra"). The integral curvatures simply equal the angular excesses, that is six times the actual pointing angles minus 180°. When we consider two triangles of different size, then the integral curvature of the larger triangle is the sum of the integral curvature of the smaller triangle (which is contained in it) and the integral curvature of the complement of the larger by the smaller triangle. Thus we obtain the integral curvature for triangular strips of various sizes by progressive subtraction of the integral curvatures. Since these strips contain only a limited range of distances, we obtain a measure of the specific excess, that is the intrinsic curvature, as a function of distance by dividing by the areas of the triangular strips.

Thus far the data processing is fairly trivial. The results are collected in table 1 and figure 7. Clearly the intrinsic curvature changes systematically with distance from the observer. Luneburg's hypothesis that human optical space has to be one of the classical spaces of constant curvature is not borne out in the present setting.

Table 1. Table of Gaussian curvatures (m^{-2}) calculated from the angular excesses determined by the equilateral-triangles experiment. The final column contains the estimated standard errors in the curvature.

Range/m	Subject			
	AD	JK	JL	σ
0.0 – 1.04	$+2.2\times10^{-1}$	$+4.4\times10^{-1}$	$+5.4\times10^{-1}$	3.8×10^{-2}
1.04 – 2.07	-7.2×10^{-2}	-5.8×10^{-2}	-3.0×10^{-2}	1.3×10^{-2}
2.07 – 4.15	-5.0×10^{-3}	-1.1×10^{-2}	-1.3×10^{-2}	3.2×10^{-3}
4.15 – 8.29	-2.9×10^{-3}	-2.8×10^{-3}	-6.2×10^{-4}	8.0×10^{-4}
8.29 – 11.06	-2.6×10^{-3}	-4.7×10^{-3}	-2.6×10^{-3}	7.7×10^{-4}
11.06 – 16.50	$+6.7\times10^{-4}$	$+5.8\times10^{-4}$	$+1.8\times10^{-3}$	2.7×10^{-4}

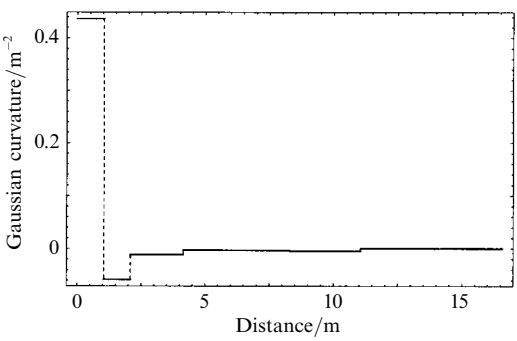

Figure 7. The intrinsic curvature for a number of distance ranges for observer JK.

From the scatter in the data we find that the curvatures are typically very significantly different from zero, except for a certain intermediate range that somewhat differs for our three observers.

5 The geodesic surface

We shall now perform one further step of data processing that is of a more complicated nature. Given the intrinsic curvature as a function of distance we proceed to construct a surface that can be depicted and has the same intrinsic curvature as the horizontal plane at eye height. In order to do this we add a dimension to the horizontal plane at eye height. This dimension is a purely formal one; it should not be confused with the vertical dimension of physical space, for instance! The only value of this formal embedding of the horizontal plane at eye height in a three-dimensional space is that it allows visualization of the curvature. We assume the surface to be a surface of revolution, reflecting the isotropic nature of optical space for an observer free to turn about the vertical. This is an approximation since the observers did not exhaust all possible head positions in any given task. However, it is clearly a reasonable first pass.

Because our data are rangewise constant, we may piece the 'geodesic surface' together from annular piece of surfaces of revolution of constant curvature (Spivak 1975; Struik 1950). One example would be a sphere, clearly a surface of revolution of constant curvature. However, it is important to appreciate that arbitrary bendings leave the intrinsic curvature invariant. Thus a cone has the same curvature as a plane (namely zero) because a cone can be obtained by bending a plane. Likewise, if you slit open a sphere by a meridian, you can bend it into a spindle-shaped body; and if you cut off two polar caps, then cut by a meridian, you can open up the sphere to look like a car-fender. Similarly, a variety of surfaces of negative constant curvature can be obtained from the familiar 'pseudosphere'. We apply the constraint that these annular pieces should join together smoothly. Even then, some freedom (of bending) is left. This can be removed

by applying a further constraint. We apply the constraint that the surface should be as close to horizontal at the origin as possible. It turns out that we cannot satisfy this latter constraint completely. The best solution still has a conical singularity at the origin.

A picture of the geodesic surface for observer JK is shown in figures 8 and 9. The geodesic surfaces for the other observers are comparable. 'Geodesics' (or shortest connections between pairs of points) can be realized by pulling a string taut between the points, constraining the string fully to the surface. The projections of such strings on the horizontal plane at eye height will be curved. By construction of the geodesic surface, the geodesic triangles on the geodesic surface will project exactly to the experimental results. Thus the geodesic surface is one (particularly intuitive) way to represent the data.

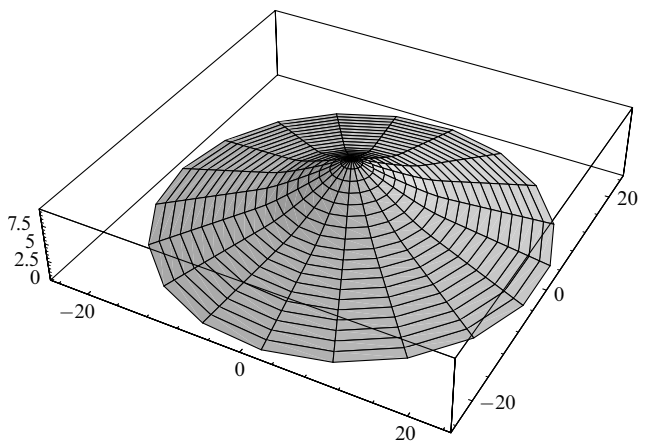

Figure 8. The geodesic surface for observer JK. Although this rendering is in 'true scale' (all numbers are meters, vertical and horizontal axes are identically scaled) this is not a surface of the space we move in. The vertical dimension is a purely formal one. The units shown are meters.

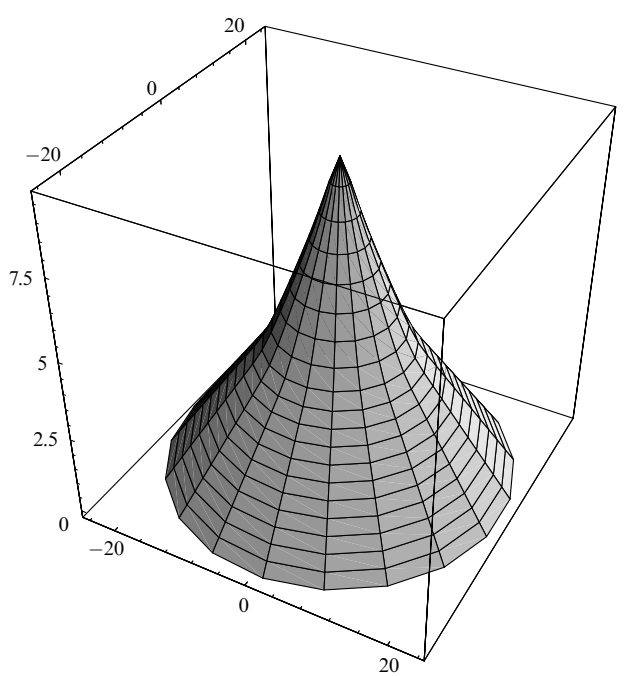

Figure 9. The geodesic surface for observer JK, vertical scale expanded. In this rendering the conical singularity at the origin is very evident. The units shown are meters.

The geodesic surface can be used to predict geodesics between any pair of points. Consequently this representation of the data allows the prediction of any exocentric pointing experiment on the basis of our present data.

The geodesic surfaces of the three observers are similar, though different in detail. They are compared in figure 10.

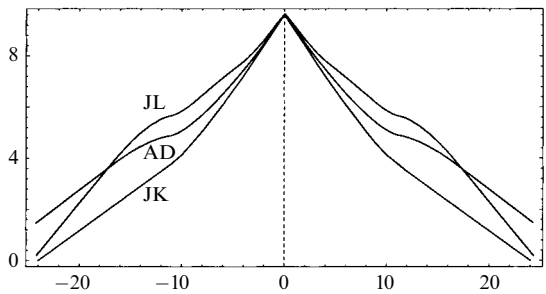

Figure 10. The geodesic surfaces for all observers compared. The graphs are sections of these surfaces by a vertical plane through the origin. The units shown are meters. (The geodesic surfaces are surfaces of revolution by construction.)

6 Discussion

We have introduced a novel method to probe human optical space, namely that of exocentric pointing. We have used this method to address the Luneburg hypothesis that human optical space should be of constant (negative) curvature. This hypothesis is clearly falsified by our data. The horizontal plane at eye height has an elliptic curvature in the near zone and a hyperbolic curvature in the far one. The neutral (parabolic) distance is about an eye height.

Existing data have predominantly been obtained under well-controlled conditions, typically with stimuli reduced to a few luminous points in the dark. Under such controlled conditions the only depth cue left is binocular disparity. In our case a far richer bouquet of depth cues was available to the observers. Moreover, our observers were free to make head movements (no biteboard), torsions at the waist, and were even permitted to turn on the feet. The task was clearly different too; nothing similar to exocentric pointing appears to have been used in the classical studies. There is thus little reason to expect any close connection between these data. After Luneburg's work many researchers attempted to determine 'the' curvature of human optical space (Blank 1953, 1959; Indow et al 1962a, 1962b, 1963). Eventually some concensus was formed (Indow 1991) that the curvature might in fact not be constant, and our data agree with this suggestion. However, as said above, our observers almost certainly used different cues to perform a different task in a very different setting.

As remarked in the introduction, our experiment is no check on the validity of the hypothesis that human optical space is of a fixed, Riemannian nature. On the face of it this seems unlikely (though invariably assumed in the classical work), but we have no independent check. In order to investigate this issue we should attempt verification of predictions from our data that go beyond our data. The geodesic surface allows us to predict results of arbitrary exocentric pointing tasks. One prediction that can immediately be derived is very surprising indeed. Owing to the fact that there is a conical singularity at the origin, we predict that points that are diametrically opposite to the observer can be joined by two distinct geodesics. Thus we predict that the observer would need to choose between two equally 'visually correct' exocentric pointing directions. From the semitop angle of the conical singularity we predict a difference of $5^\circ - 10^\circ$ between such equally acceptable pointings: Thus such an effect is in the measurable range. Although such a prediction appears prima facie rather unlikely, we have actual indications that such an effect may in fact occur. We hope to report on this in a subsequent communication.

References

Battro A M, Piero Netto S di, Rozestraten R J A, 1976 "Riemannian geometries of variable curvature in visual space: visual alleys, horopters, and triangles in big open fields" *Perception* **5** 9 – 23

Battro A M, Reggini H C, Karts C, 1978 "Perspectives in open spaces: a geometrical application of the Thouless index" *Perception* **7** 583 – 588

Beltrami E, 1865, see Struik (1950)

Blank A A, 1953 "The Luneburg theory of binocular visual space" *Journal of the Optical Society of America* **43** 717 – 727

Blank A A, 1959 "The Luneburg theory of binocular perception", in *Psychology: A Study of a Science* volume 1 *Sensory, Perceptual, and Physiological Formulations* Ed. S Koch (New York: McGraw-Hill) pp 395 – 426

Ellis S R, Smith S, Grunwald A, McGreevy M W, 1991 "Direction judgement error in computer generated displays and actual scenes", in *Pictorial Communication in Virtual and Real Environments* Ed. S R Ellis (London: Taylor and Francis) part IV, chapter 34, pp 504 – 526

Gauss K F, 1880 "Disquisitiones generales circa superificies curvas", in *Werke* (Göttingen: Königliche Gesellschaft der Wissenschaften zu Göttingen) volume 4, pp 217 – 258

Helmholtz H, 1867 *Handbuch der physiologischen Optik* 1st edition (Leipzig: Leopold Voss)

Indow T, Inoue E, Matsushima K, 1962a "An experimental study of the Luneburg theory of binocular space perception (1). The 3- and 4-point experiments" *Japanese Psychological Research* **4** 6 – 16

Indow T, Inoue E, Matsushima K, 1962b "An experimental study of the Luneburg theory of binocular space perception (2). The alley experiments" *Japanese Psychological Research* **4** 17 – 24

Indow T, Inoue E, Matsushima K, 1963 "An experimental study of the Luneburg theory of binocular space perception (3). The experiments in a spacious field" *Japanese Psychological Research* **5** 10 – 27

Indow T, 1991 "A critical review of Luneburg's model with regard to global structure of visual space" *Psychological Review* **98** 430 – 453

Luneburg R K, 1947 *Mathematical Analysis of Binocular Vision* (Princeton, NJ: Princeton University Press)

Pirenne M H, 1970 *Optics, Painting and Photography* (Cambridge: Cambridge University Press)

Spivak M, 1975 *A Comprehensive Introduction to Differential Geometry* volume III (Houston, TX: Publish or Perish)

Struik D J, 1950 *Lectures on Classical Differential Geometry* (New York: Dover)

Wagner M, 1985 "The metric of visual space" *Perception & Psychophysics* **38** 483 – 495