On the Affine Structure of Perceptual Space

James T. Todd, Augustinus H. J. Oomes
The Ohio State University

Jan J. Koenderink, Astrid M. L. Kappers
Universiteit Utrecht

Affine geometry is a generalization of Euclidean geometry, in which distance can be scaled along parallel directions, though relative distances in different directions may be incommensurable. A new procedure is presented for testing the intrinsic affine structure of a psychological space by having subjects perform bisection judgments over multiple directions. If those judgments are internally consistent with one another, they must satisfy a theorem first proved by Pierre Varignon around 300 years ago. An experiment is reported in which this procedure was employed to measure the perceived structure of a visual ground surface. The results revealed that observers' judgments were systematically distorted relative to the physical environment, but that the judged bisections in different directions had an internally consistent affine structure. Implications of these findings for other possible response tasks are considered.

Introduction

One of the most popular ways of modeling perceptual (or cognitive) phenomena is to represent percepts (or concepts) as points within some psychological space. Often it is assumed, moreover, that a psychological space has an underlying geometric structure that constrains how different judgments are related to one another. There are many different geometries that could potentially be used for modeling psychological phenomena (e.g., see Suppes, 1977; Suppes, et al, 1989). In the present article we will consider one called affine geometry that may be unfamiliar to many researchers in psychology. We will first motivate why this particular geometry might be useful, and we will then introduce a new technique for testing whether the affine properties of a psychological space are internally consistent.

Of all the possible aspects of geometric structure, the one that has attracted the most attention in the psychological literature is called a distance metric. The metric of a space is a mathematical function that relates distances in different directions. For example, the metric of Euclidean space is defined by the Pythagorean theorem. Euclidean geometry was developed over 2000 years ago as an abstract model of the physical environment, but it does not always provide an adequate description of the distance relations within a psychological space. Other possible metrics that have been proposed in this context include Minkowski metrics (e.g., see Attneave, 1950) and Riemannian metrics (e.g., see Luneberg, 1947), both of which are based on alternative conceptions of space that are non-Euclidean.

For any space to be considered as metric, there must be a distance function δ by which every pair of points can be assigned a nonnegative distance value that conforms with three basic axioms:

Minimality:

$$\delta(a,b) \geq \delta(a,a) = 0$$

Symmetry:

$$\delta(a,b) = \delta(b,a)$$

The Triangle Inequality:

$$\delta(a,b) + \delta(b,c) \geq \delta(a,c)$$

It is interesting to consider, therefore, the extent to which these axioms are satisfied for various types of psychological judgments. One early attempt to address this issue was performed by Shepard (1964). He asked subjects to make similarity judgments for line drawings of a circle with a single radial line, in which the size of the circle and the orientation of the line could be systematically varied. His results revealed that the observers' responses could not be adequately characterized by any static metric, because the relative weighting of the different dimensions in the overall similarity judgment could change with a subject's state of attention. When responses obtained with different states of attention were combined, the resulting data contained clear violations of the triangle inequality.

To provide a more intuitive example of how this could occur, Shepard also offered the following Gedanken experiment: Suppose that subjects were asked to rate the dissimilarity (i.e. distance) between pairs of words on a scale of 1 to 100. If given the pair “table” and “fable”, they would likely notice that the words sound alike, and assign a low dissimilarity rating, such as 15. If given the words “table” and “chair”, the subjects would probably switch their attention to the frequent co-occurrence of these objects in the natural environment, and would again provide a low dissimilarity rating, such as 10. However, the pair “chair” and “fable” have no apparent resemblance at all, and would likely be given a high dissimilarity rating, such as 80. This pattern of results would be a violation of the triangle inequality, and would therefore indicate that the psychological representation of these words is nonmetric.

Let us now consider how affine geometry might provide a more appropriate model of these phenomena. Affine geometry is a generalization of Euclidean geometry, with a more limited set of assumptions. Both geometries share the axioms of incidence and the parallel postulate, but affine geometry does not require the axioms of congruence. (See Blumenthal, 1961, for a more detailed discussion of the axiomatic bases of alternative geometries.) Within an affine space it is possible to compare the relative lengths of all parallel line segments[1], though the relative lengths of nonparallel line segments may be incommensurable.

The psychological distinction between parallel and nonparallel distance intervals was nicely demonstrated in another experiment reported by Shepard (1964). He asked observers to perform two tasks, one that compared intervals along a single dimension, and another that compared intervals across different dimensions. On the intra-dimensional task, observers were required to adjust an angle between two lines to match the angle between two other lines at different orientations. On the cross-dimensional task, they were required to adjust an angle to match the size difference between two circles. All subjects agreed that the intra-dimensional task was quite trivial, but most of them insisted that they could find no intuitive basis for performing the cross-dimensional judgments. Consistent with these subjective reports, the variances obtained for the cross-dimensional task were 19 times larger than when the matching stimuli differed along a single dimension.

Shepard (1964) argued that cross-dimensional confusions are likely to be limited to artificially defined spaces, whose component dimensions have no natural relationship with one another. However, more recent evidence has revealed that similar distinctions between judgments of parallel and nonparallel distance intervals can also occur for other psychological spaces that are generally believed to be homogeneous. For example, when observers are asked to compare distance intervals in the physical environment, their judgments are more accurate and reliable when the comparison stimuli are parallel to one another (Norman, Todd, Perotti, & Tittle, 1996; Todd & Bressan, 1990). These findings provide strong evidence that affine properties of the physical environment may be more perceptually salient than its metrical properties.

It is important to keep in mind when considering these issues that there are many possible geometries involving different sets of underlying assumptions that could potentially be useful within mathematical psychology. Faced with this plethora of possibilities, how is one to decide which geometry is most appropriate in any given context? One could of course just assume a particular geometry and hope that its underlying assumptions are valid, but we believe that a better approach is to perform an independent check of those assumptions, such as Shepard’s (1964) test of the triangle inequality. In order to adopt this approach, however, it would be necessary to devise a set of formal procedures for assessing the internal consistency of observers’ judgments about the relevant properties of each geometry to be considered.

One way of testing the internal consistency of affine judgments is to exploit a theorem that was first proven by Pierre Varignon around 1700. Let P_1, P_2, P_3 and P_4 be arbitrarily selected points that define the vertices of a quadrilateral. Let Q_1, Q_2, Q_3 and Q_4 be the bisection points along each

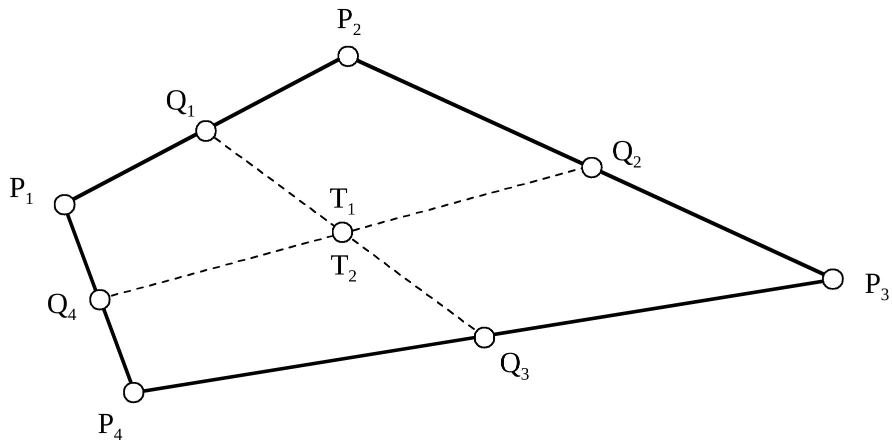

Figure 1 -- A Varignon configuration similar to those used in the present experiment to investigate the affine structure of perceptual space. Points P_1-P_4 mark the vertices of a quadrilateral. The points Q_1-Q_4 bisect the edges of the quadrilateral, and the points T_1 and T_2 bisect the intervals between the opposing edge bisections Q_1Q_3 and Q_2Q_4.

quadrilateral edge, and let T_1 and T_2 bisect the intervals between each opposing edge bisection (see Figure 1). In an affine space, the points Q_1-Q_4 will form a parallelogram, and the points T_1 and T_2 will be coincident with one another. It is important to recognize that intersecting line segments do not generically bisect one another. Thus, the coincidence of T_1 and T_2 in Varignon's theorem imposes a severe constraint on the structure of an affine space, by which interval bisections in different directions are formally related to one another.

This also provides a straightforward procedure for measuring the internal consistency of affine structure for any psychological space in which it is possible to make bisection judgments. The procedure involves two separate phases. Subjects would first make bisection judgments for each edge of an arbitrary quadrilateral to obtain Q_1-Q_4. Next they would bisect the intervals between the judged bisections of the opposing edges to obtain T_1 and T_2. If the space is affine, then the final two judgments must be statistically equivalent.

Recently we have employed this procedure to measure the intrinsic affine structure of a perceived ground surface. This particular type of psychological space is of special interest, because prior research has shown that it is systematically distorted relative to the actual physical environment (e.g., see Battro, Netto, & Rozestraten, 1976; Koenderink, van Doorn, & Lappin, 2000a; Norman et al., 1996). That is to say, physically straight lines can appear perceptually to be curved, and equal colinear length intervals can appear perceptually to be unequal. These findings demonstrate that observers' judgments of affine properties in the environment can be physically inaccurate, but are they internally consistent? Our research was designed to address this question.

Methods

Apparatus. The stimuli were created and displayed on a Mac G3 computer with a 21 inch monitor. The displays were viewed through LCD (liquid crystal display) shuttered glasses that were synchronized with the monitor's refresh rate. The different views of a stereo pair were displayed at the same position on the monitor screen, but they were temporally offset. The left and right lenses of the LCD glasses shuttered synchronously with the display at an alternation rate of 60 Hz, so that each of the two stereo images could only be seen by the appropriate eye. When operating in stereo mode, the spatial resolution of the monitor was 1280 X 484 pixels. The displays were viewed from a distance of 57.3 cm so that each image subtended 38.5 X 25.8 degrees of visual angle. Head movements were restricted using a chin rest.

Stimuli. Each display depicted a blue and black textured ground surface with three vertical red posts (see Figure 2). The simulated ground surface was constructed from a 6.0 X 6.0 m rectangle located 15.0 cm below the point of observation. The right and left edges of this rectangle were located just outside the viewing frustum of the display window so that they would not be visible. The surface texture was created from a 256 X 256 random check pattern of blue squares, whose individual intensities were selected at random from a range of 256 possible values. This pattern was blurred by a Gaussian function with a standard deviation equal to the width of one check, and it was then thresholded to produce a binary blue and black texture map. Eight different textures were created, one of which was selected at random on each trial.

The depicted surface had a random pattern of hills and valleys so that the depth of a post could not be determined by its point of intersection with the ground. These bumps were created by subdividing the rectangular surface into a 64 X 64 square grid. The vertices of this grid were assigned random heights over a range of –2.0 to 2.0 cm, and the regions between vertices were smoothly interpolated to create the appearance of a continuous surface. A different random pattern of bumps was generated for each trial.

The heights of the red posts were also varied, over a range of 11.3 to 18.8 cm. This was intended to prevent observers from judging the depth of a post from its height in the visual field. Two of the posts presented in each display had fixed positions on the ground that varied across trials. The position of the third post could be adjusted both horizontally and in depth by manipulating a hand held mouse. Twelve different Varignon configurations were created in which the depths of the probe points (P_1-

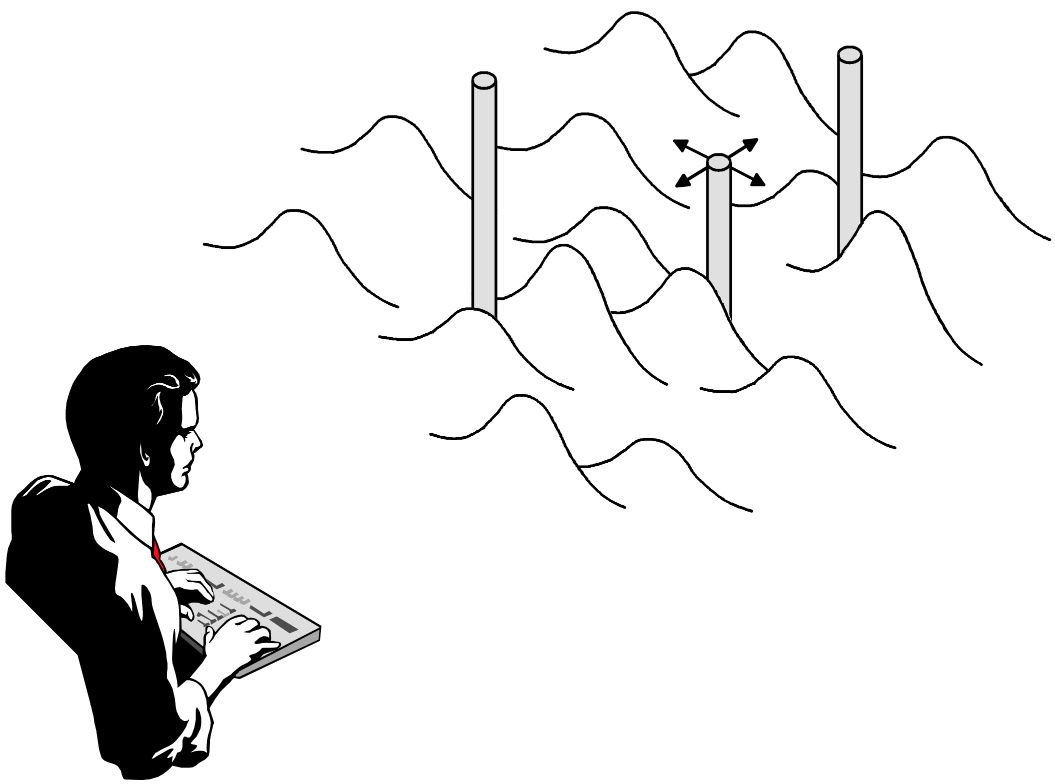

Figure 2 – A schematic view of a typical stimulus configuration used in the present experiment. Observers adjusted a vertical post on a bumpy ground surface until it appeared to bisect the interval between two other fixed posts.

P_4) ranged from 130 to 395 cm. These configurations were constrained such that no pair of posts could visually occlude one another.

Procedure -- The task on each trial was to adjust the moveable post so that it appeared to bisect an imaginary line between the two fixed posts (see Figure 2). The experiment was conducted in two separate phases. During the first phase, the positions of the fixed posts were selected from the points P_1, P_2, P_3 and P_4 of a predetermined Varignon configuration (see Figure 1), and each pair of points was repeated on 10 separate trials. In the second phase of the experiment, the fixed posts were positioned at the mean locations of the observers' judgments for Q_1Q_3 or Q_2Q_4 obtained during the first phase. Each pair of points was again repeated on 10 separate trials to estimate the apparent locations of T_1 and T_2. Three different configurations were interleaved within each of four experimental sessions.

Observers – Six naïve observers participated in the experiment and were paid $8 per hour for their services. All had normal or corrected to normal visual acuity. Each observer performed the sessions in a different random order.

Results

Two representative patterns of responses for different observers and different configurations are shown in Figure 3. The trapezoidal boundary in each figure shows the viewing frustum of the display window. The dotted lines show the actual Varignon configuration in physical space, with depth represented in the vertical direction. The solid circular arcs show the perceived configuration as revealed by the observer's judgments. The small ellipses mark the mean positions of the observer's settings, and the axis lengths of these ellipses in different directions show the standard deviations. The two ellipses in the center of the configuration show the judged positions of T_1 and T_2, respectively.

In analyzing the errors in observers' judgments, it is useful to distinguish two separate components, which we will refer to as *intrinsic* and *extrinsic*. The extrinsic errors are revealed by systematic differences between the judged test points and their actual locations in physical space. Note in Figure 3, for example, that the judged locations of T_1 and T_2 were much closer in depth than their actual locations, which provides strong evidence that the extrinsic geometry of observers' judgments was systematically distorted. This finding is consistent with many previous reports in the literature that physically straight lines can appear perceptually to be curved, and that equal length intervals can appear perceptually to be unequal (e.g., see Battro et al., 1976; Koenderink et al., 2000a; Norman et al., 1996).

Although these extrinsic distortions are of considerable interest, our primary goal in the present experiment was to measure the intrinsic structure of perceptual space. Intrinsic errors are revealed in this context by violations of Varignon's theorem -- that is to say, when the two test points of a given configuration are significantly different from one another. For both of the examples provided in Figure 3 the response distributions for the two test points are almost identical, thus indicating that the observers' judgments had an internally consistent affine structure.

Figure 4 provides a summary of the intrinsic and extrinsic errors for all six observers in both a horizontal direction and in depth. It is clear from this figure that there was a strong anisotropy in the observers' judgments. Although there were large extrinsic errors along the depth dimension, the judged horizontal positions of the test points were relatively accurate. It is also important to note in this figure that the extrinsic errors were over four times larger than the intrinsic errors.

To analyze this pattern of results quantitatively, a series of Hotelling's T^2-tests were performed for each observer's judgments of each configuration. First we examined the extrinsic distortions of their judgments relative to the actual locations of T_1 and T_2 on the simulated ground surface. Out of 144 possible comparisons (6 observers X 12 configurations X 2 test points), 133 were significantly different.

Next we examined the internal consistency of the observers' perceptions by comparing their settings for the two test points in each configuration. Out of 72 possible comparisons (6 observers X 12 configurations), only 15 of the judged test pairs were significantly different from one another. Although this might seem to be higher than the expected outcome due to chance, our analysis did not take into account the propagation

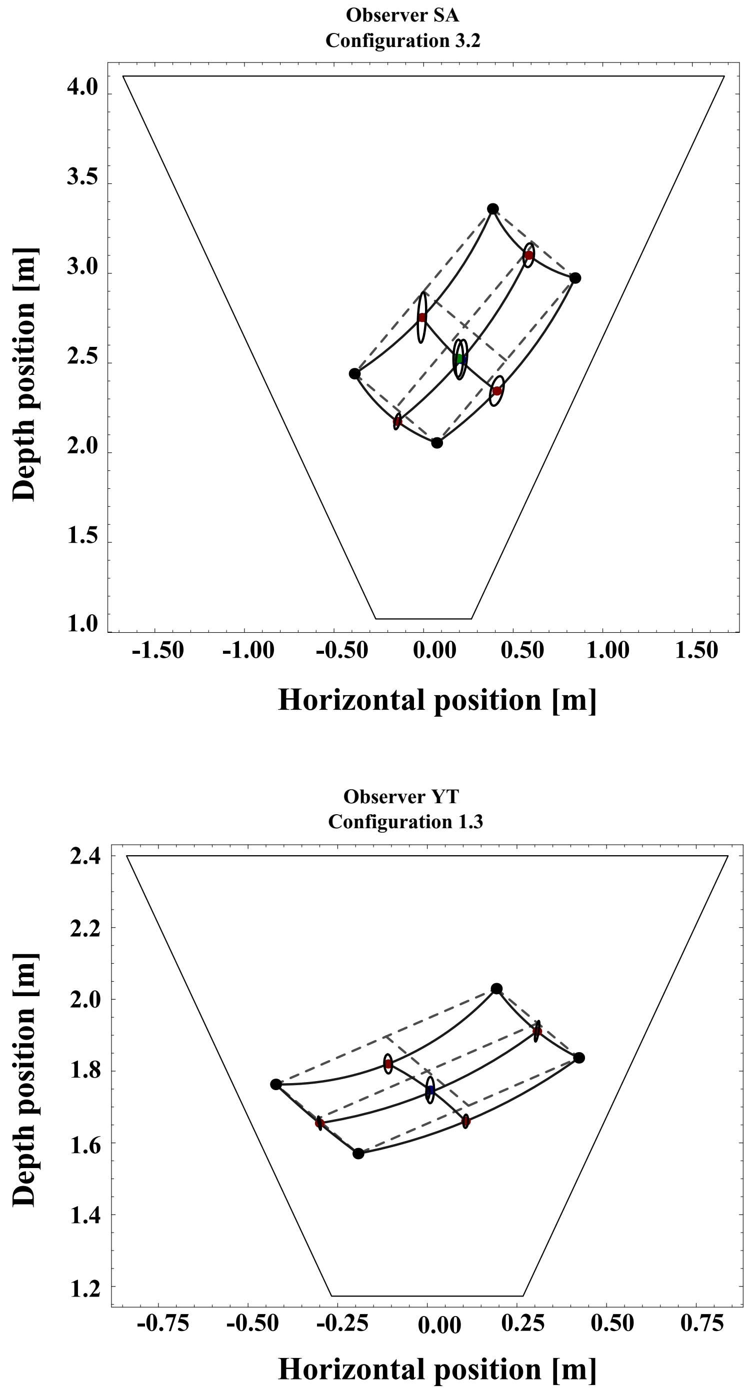

Figure 3 -- Two representative patterns of responses for different observers and different configurations. The trapizoidal boundary in each panel shows the viewing frustum of the display window. The dotted lines show the actual Varignon configuration in physical space, and the solid curves show the perceived configuration as revealed by the observer's judgments. The small ellipses mark the mean positions of the observer's settings, and the axis lengths of these ellipses in different directions show the standard deviations.

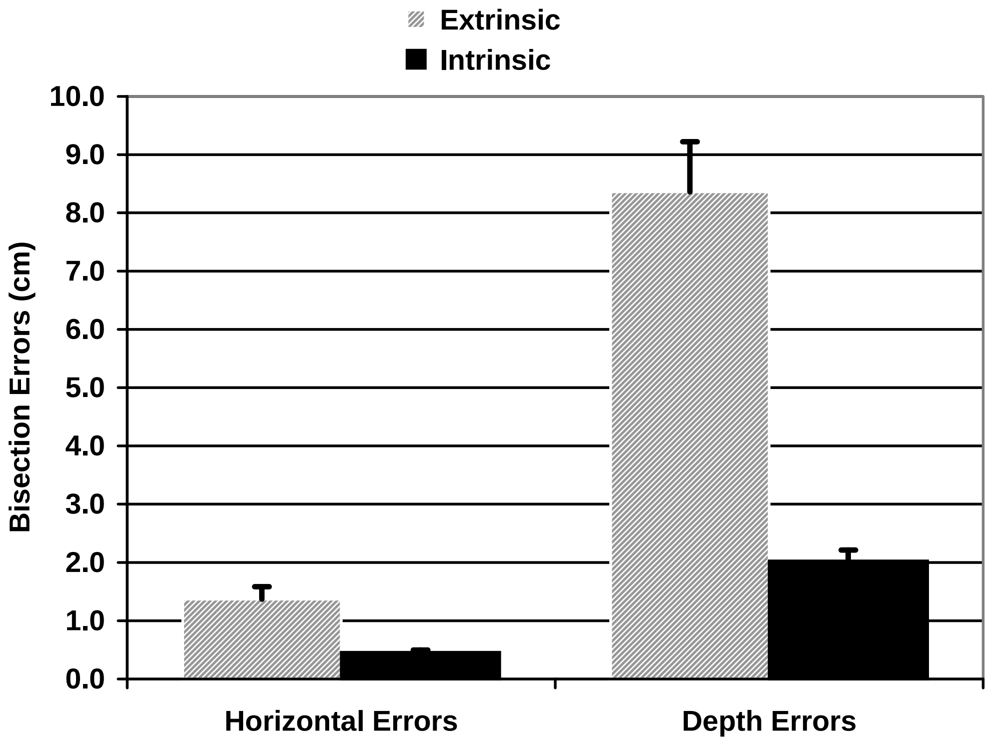

Figure 4 – The mean bisection errors in the horizontal direction and in depth averaged over all six observers. The intrinsic errors are the average distance between the two test points (T_1 and T_2) in each configuration. The extrinsic errors are the average distance between the judged test points and their actual locations in physical space.

which increases the probability of falsely rejecting the null hypothesis. When we examined these apparent violations of Varignon's theorem, they appeared to be randomly distributed among the different observers and conditions. For each observer, the theorem was satisfied for at least 9 of the 12 possible configurations, and for each configuration, the theorem was satisfied by at least 4 of 6 observers. These results suggest, therefore, that the apparent violations of Varignon's theorem were most likely due to measurement error, and that the intrinsic geometry of the observers' perceptions had an internally consistent affine structure.

Discussion

It is important to recognize when evaluating the present experiment that the "geometry" of perceived space can be construed in two different ways. One possibility is to consider the extrinsic structure of observers' perceptions relative to the physical environment. From an extrinsic point of view, the structure of perceptual space (Ψ) is determined by its formal relation to physical space (Φ), such that $\Psi = f(\Phi)$. Within this context, the geometry of perceived space is defined by the particular set of properties that are invariant over the transformation $f(\Phi)$. In previous articles we have argued that the visual perception of 3D structure from motion or shading is affine in this extrinsic sense (Koenderink & van Doorn, 1991; Koenderink, van Doorn, Kappers, & Todd, 2000b; Todd & Bressan, 1990; Todd & Norman, 1991), but that is clearly not the case in the present experiment for the perception of 3D structure from binocular disparity. Note in Figure 3, for example, that straight lines in the physical environment can be perceived as curved, thus indicating that the transformation $f(\Phi)$ in this instance is neither affine nor projective.

In contrast to most previous studies on 3D form perception, our primary goal in the present experiment was to investigate the intrinsic geometry of perceptual space. Intrinsic geometry provides a global set of constraints by which the judgments of a given observer are formally related to one another, irrespective of their relation to the external environment. Our long term goal in this research is to develop a "tool box" of procedures for measuring the internal consistency of observers' judgments about various aspects of geometric structure. In this case our focus was on affine structure, and the tool for measuring its internal consistency was provided by Varignon's theorem.

Let us now consider how affine spaces are related to other types of geometric structure that have been proposed in the literature for modeling psychological phenomena. Affine geometry is based on several basic axioms, including the axioms of incidence and Euclid's parallel postulate. The axioms of incidence state that two points are connected by a single line, and that two lines are connected at a single point. The parallel postulate states that for a given line L and a given point P, there is a single line through P that is parallel to L. Euclidean geometry is a special case of affine geometry in which additional axioms are added. Thus, to demonstrate that a psychological space is affine does not preclude the possibility that it may also be Euclidean.

Other types of metrical geometry are based on assumptions that specifically contradict Euclid's parallel postulate. For example, in elliptic geometry there are no lines through P that are parallel to L, and in hyperbolic geometry there are an infinity of lines through P that are parallel to L. Euclidean, elliptic and hyperbolic geometries are all special cases of a more general framework called Riemannian geometry, which can describe the structure of any smooth manifold. Riemann spaces of constant curvature – the so-called homogeneous spaces – are the only ones that allow congruence. These can be subdivided into three distinct types based on the sign of the intrinsic curvature. The geometry of spaces with no intrinsic curvature (e.g., planes or cylinders) is Euclidean; the geometry of spaces with positive curvature (e.g., spheres) is elliptic; and the geometry of spaces with negative curvature (e.g., saddles) is hyperbolic.

It is interesting to note in this regard that there have been numerous experiments described in the literature that are purported to show that the intrinsic curvature of perceptual space is measurably non-Euclidean (e.g., Battro et al., 1976; Indow, 1991; Koenderink et al., 2000a; Norman et al., 1996). This would seem to contradict the results of the present study, since Euclidean space is the only Riemannian geometry that is also affine. There are several points to keep in mind when considering this issue. First, these experiments have all been based on an a priori assumption that perceptual space has a stable Riemannian distance metric, but there is no independent evidence to verify that assumption. To the extent that it may be true, moreover, there is strong evidence to suggest that the curvature of perceptual space varies with position (Indow, 1991; Koenderink et al., 2000a), and that there are large individual differences among observers. For example, in one series of experiments by Battro, et al. (1976) involving over 120 observers, the results obtained for 60% revealed a negative curvature, whereas those for the remaining 40% revealed a positive curvature.

One possible explanation for these large variations in the metric structure of perceptual space is that the underlying geometry of observers' judgments may be dependent on contextual factors (see Suppes, 1977), such as the presence of visible objects in the environment or how an observer interprets an experimenter's instructions. We believe it is best to be skeptical, however, about the generality of this phenomenon. Although context may be important for the metric structure of perceptual space, perhaps it is the case that there are other more primitive aspects of observers' perceptions that exhibit a higher degree of stability.

Within the hierarchy of possible geometries, affine structure is more primitive than Euclidean structure, because it is based on a smaller set of underlying assumptions, and is therefore invariant over a larger set of possible transformations. Alternative geometries can also be devised for which these assumptions are relaxed still further. For example, projective geometry is a generalization of affine geometry without any axioms about parallelism.

Just as it is possible to test the internal consistency of affine judgments using Varignon's theorem, one can also explore the projective structure of perceptual space using a theorem first proven by Pappus of Alexandria around 340 AD. The Pappus theorem provides a global constraint on how line segments in different directions must intersect one another. We have recently performed a new set of experiments to investigate whether this theorem is satisfied for observers' co-linearity judgments on a visible ground surface, and the results of those studies are described in (Koenderink, van Doorn, Kappers, & Todd, 2000c).

References

Attneave, F. (1950). Dimensions of similarity. *American Journal of Psychology, 63*, 515-556.

Battro, A. M., Netto, S. P., & 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.

Blumenthal, L. M. (1961). *A modern view of geometry*. New York: Dover.

Indow, T. (1991). A critical review of Luneburg's model with regard to global structure of visual space. *Psychol Rev, 98*(3), 430-53.

Koenderink, J. J., van Doorn, A., & Lappin, J. (2000a). Direct measurement of the curvature of visual space. *Perception, 29*, 69-80.

Koenderink, J. J., & van Doorn, A. J. (1991). Affine structure from motion. *Journal of the Optical Society of America A, 8*, 377-385.

Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2000b). Ambiguity and the "mental eye" in pictorial relief. *Perception, 29*.

Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2000c). Pappus in Optical Space. *Perception, 29*.

Luneberg, R. K. (1947). *Mathematical analysis of binocular vision*. Princeton, NJ: Princeton University Press.

Norman, J. F., Todd, J. T., Perotti, V. J., & Tittle, J. S. (1996). The visual perception of 3D length. *Journal of Experimental Psychology: Human Perception and Performance, 22*, 173-186.

Shepard, R. N. (1964). Attention and the metric structure of stimulus space. *Journal of Mathematical Psychology, 1*, 54-87.

Suppes, P. (1977). Is visual space Euclidean? *Synthese, 35*, 397-421.

Suppes, P., Krantz, D. H., Luce, R. D. & Tversky, A. (1989) *Foundations of Measurement: Volume 2.* Academic Press: San Diago, CA.

Todd, J. T., & Bressan, P. (1990). The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Perception & Psychophysics, 48*, 419-430.

Todd, J. T., & Norman, J. F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception & Psychophysics, 50*, 509-523.

Footnotes

This collaboration was supported in part by a NATO Scientific Exchange Grant (CRG 92065). In addition, James Todd and Stijn Oomes were supported by grants from NIH (R01-Ey12432-01) and NSF (SBR-9514522). Correspondence should be addressed to James T. Todd, Department of Psychology, 142 Townshend Hall, The Ohio State University, Columbus, OH 43210

[1] It is possible to define an affine geometry based only on the axioms of incidence and Euclid's parallel postulate that does not provide sufficient structure to establish an equivalence relation between parallel line segments. In order to ensure this property, it is necessary to include an additional axiom, which can take the form of either Desargues' theorem or Pappus' theorem (See Blumenthal, 1961).