Published in final edited form as:

J Math Psychol. 2009 April 1; 53(2): 86–91. doi:10.1016/j.jmp.2008.12.006.

Is perceptual space inherently non-Euclidean?

Julian Martin Fernandez^{*} and Bart Farell

Institute for Sensory Research, Syracuse University, Syracuse, NY, 13244, USA

Abstract

It is often assumed that the space we perceive is Euclidean, although this idea has been challenged by many authors. Here we show that, if spatial cues are combined as described by Maximum Likelihood Estimation, Bayesian, or equivalent models, as appears to be the case, then Euclidean geometry cannot describe our perceptual experience. Rather, our perceptual spatial structure would be better described as belonging to an arbitrarily curved Riemannian space.

Keywords

Euclidean; cue combination; Bayesian; maximum likelihood; depth perception; space perception

Introduction

Possible perceptual spaces

In his *Critique of Pure Reason*, Kant (1902) argued that the truths of geometry are synthetic *a priori* truths. That is, three-dimensional Euclidean space is a necessary, but not tautological, presupposed form underlying all human spatial experience. Kant seems to be referring to what we today call the “intrinsic” geometry of perceptual space, rather than its “extrinsic” geometry. Extrinsic geometry refers to the relationship between the structure of the observer's perception and the actual structure of physical space. It gives the geometrical transformations necessary to map physical space onto perceptual space. Let us give an example: assume that both physical space and our internal representation of space are Euclidean, and that our percepts are distorted so that perceived objects are veridical only up to a scaling factor in depth (e.g., a circle is perceived as an ellipse). If such were the case, then the extrinsic geometry would be affine. It has long been known that many geometric relations are distorted in perception. For example, perceived distance is compressed over a large range (Gilinsky, 1951), apparent parallel alleys and equidistance alleys are not physically parallel and equidistant (Blumenfeld, 1913; Indow & Watanabe, 1984), apparent frontoparallel planes are not physically frontoparallel (Helmholtz, 1962; Ogle, 1964), and lines perceived as curved might be straight in the physical environment (Todd *et al.*, 2001).

Intrinsic geometry, by contrast, 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 topic here is the structure of the intrinsic geometry of perception¹. That is,

^{*} Corresponding author. Address: Institute for Sensory Research, Syracuse University, 621 Skytop Rd, Syracuse, NY 13244, USA. E-mail: julian_fernandez@isr.syr.edu.

Commercial relationships: none.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

we are dealing with the internal geometry of a subject's perceptual space and not at all with the relationship between perceived and actual shape. Of course, we are assuming that there is a geometry of internal perceptual space—in other words, that perception is stable and consistent enough to support such a geometry.

There are fundamentally two different ways in which intrinsic geometry could depart from being Euclidean. The first way is for it to stay within the realm of more primitive geometries, this is, geometries that do not have an internal metric structure. An instance of this is affine geometry. The second way is to keep a metric structure but to define this metric in a different way than in Euclidean geometry. An instance of this is Riemannian geometry.

The fact that perceived structure could be non-Euclidean² is not trivial. The interpretation of many psychophysical studies relies on the assumption of a Euclidean percept. The conclusions from any psychophysical experiment in which a variable is measured by an indirect method can potentially be affected by erroneously assuming Euclidean geometry as valid. By “indirect method” we mean estimating the variable of interest by measuring a different variable and linking the two variables through standard Euclidean geometry. For instance, in studies of structure from motion (Domini & Caudek, 2003; Domini & Braunstein, 1998; Domini, Caudek & Richman, 1998) and motion-stereo depth cue combination (Domini, Caudek & Tassinari, 2006), it is often assumed that the relationship between the perceived slant of a surface and the perceived relative depth between two points on that surface satisfies Euclidean geometry. It is unclear at this time whether and how the conclusions drawn from such studies would be affected if the Euclidean assumption were relaxed.

Siding with more primitive geometries, Gibson (1979) suggested that Euclidean metric distances in 3-dimensional space are not a primary component of an observer's perceptual experience. This hypothesis has been developed and tested by others (Domini & Caudek, 2003; Domini & Braunstein, 1998; Domini, Caudek & Richman, 1998; Norman & Todd, 1992, 1993; Tittle *et al.*, 1995; Todd & Bressan, 1990; Todd & Norman, 1991; Todd & Reichel, 1989). Their findings indicate that observers are quite accurate and reliable at judging an object's topological, ordinal, or affine properties and that the perception of rigid motion occurs when these properties remain invariant over time. However, accuracy is low for judgments requiring veridical perception of Euclidean metric structure, such as judgments of lengths or angles. One proposal, discussed in detail by Norman and Todd (1992), is that perceptual space can be best described as a more abstract space, in which the concept of distance, as found in Euclidean space, is not defined. Norman and Todd give the hierarchy of spatial structures that might be available to our perceptual system, from the most concrete to the most abstract (i.e., to the most primitive), as Euclidean, affine, ordinal, topological, and nominal (or categorical). This hierarchy has some resemblance to the classification of spaces given by Klein (1957), who showed that Euclidean geometry could be seen within an evolutionary sequence of more and more complex geometries: topological, affine, similarity group, Euclidean. Another, more extreme, proposal suggests that the internal structure of an object, as recovered from structure from motion, is internally inconsistent (Domini & Caudek, 2003; Domini & Braunstein, 1998; Domini, Caudek & Richman, 1998). Thus, although observers can exhibit a conceptual understanding of Euclidean metric structure, the basis of this knowledge might be more cognitive than perceptual.

¹Notice that, as used in psychophysics, the words “intrinsic” and “extrinsic” have different meanings than those given in differential geometry.

²We will use the term non-Euclidean to include any geometry that is not Euclidean. In some branches of geometry, the term non-Euclidean is usually restricted as meaning Riemannian geometries other than Euclidean.

However, in all the cases described above the psychophysical evidence is also consistent with an alternative interpretation, one in which the recovered structure of the perceived object (that is, its intrinsic structure) is still essentially Euclidean but physically inaccurate. That is, the perceived shape of an object differs from the actual or simulated shape, so these experiments might reflect the properties of extrinsic rather than intrinsic geometry.

By a second alternative, mentioned above, the departure from Euclidean geometry is through a change in the metric of the space, so that the concept of distance is still defined but differently from the way it is defined in Euclidean geometry. This generalization of Euclidean space is Riemannian space. Riemannian geometry was first put forward in a general form by Bernhard Riemann in the nineteenth century (Petersen, 1998). It deals with a broad range of geometries whose metric properties vary from point to point, as well as two standard types of non-Euclidean geometry, spherical geometry and hyperbolic geometry, plus Euclidean geometry itself. Euclidean space, whose curvature is zero, is the simplest case of Riemannian space. The essential difference between Euclidean and Riemannian geometry is the nature of parallel lines. In Euclidean geometry, if we start with a line l and a point A not on l , then we can draw only one line through A that is parallel to l . In hyperbolic geometry, by contrast, there are infinitely many lines through A parallel to l , and in elliptic geometry, parallel lines do not exist. An example often used to make Riemannian space more intuitive is to look at what happens in 2D space. Euclidean geometry is the geometry perceived by an observer living within a flat surface (for example, a plane). A non-Euclidean geometry is the geometry perceived by an observer living within a curved surface (for example, a sphere), the surface being curved into a third spatial dimension. Riemannian geometry, as a description of perceptual space, is clearly a different alternative than those of the first of our categories, which contains more abstract (primitive) geometries—e.g., affine, topological, and similarity spaces—in which distance is not defined.

Luneburg (1947) introduced a Riemannian space of constant curvature as a description for visual space, a model that was further developed by Blank (1958, 1978). Recent experiments have shown that the assumption of a constant curvature is generally not valid (Cuijpers, Kappers & Koenderink, 2001; Koenderink & van Doorn, 2000). For instance, Koenderink and van Doorn (2000) used classic geometry in a method that positioned the subject at the barycenter of equilateral triangles of various sizes. By remote control, the subject rotated a horizontal arrow located at one vertex of the triangle so that it appeared to point to a sphere located at another vertex. From the angle subtended by the visual direction to the arrow and the exocentric pointing direction—the veridical value was 30° —Koenderink and van Doorn obtained the perceived angle subtended by a vertex of the equilateral triangle. Then they derived the curvature of Riemannian geometry from the departure of the sum of angles of the three vertexes from 180° . Koenderink and van Doorn found that the curvature changed from elliptic in near space to hyperbolic in far space, to parabolic at very large distances. Though Howard and Rogers (2002) note the possibility of measurement bias distorting subjects' settings, the results of this elegant experiment are intriguing and important.

Multiple cues and the geometry of perceptual experience

Regardless of what geometry best describes our perceptual experience, shape must be obtained by combining information from the multiple cues available in the sensory input. One of the major problems in vision (and in perception in general) is to understand how the brain integrates the information provided by multiple cues. For example, by combining information from several depth cues, the visual system can estimate 3D layout with greater precision across a wider variety of viewing conditions than it could by relying on any one cue alone (Clark & Yuille, 1990). To realize this advantage, the reliability of each depth cue must be factored into the combination rule.

Several approaches have been proposed to understand how our perceptual system can solve the cue-combination problem in an effective manner. One approach to optimizing cue combination, the maximum-likelihood estimation model (MLE), is statistical and uses a cue-combination rule that results in an estimator that is often unbiased and has minimum variance (for a review, see Oruc, Maloney & Landy, 2003). An alternative approach is to apply Bayesian methods, in which the observer chooses an estimate that is the most probable given the image data (for a review, see Kersten, Mamassian & Yuille, 2004). The two approaches give very similar estimates under many circumstances. A number of studies have tested and confirmed the quantitative predictions of the MLE model for various cues (Alais & Burr, 2004; Ernst & Banks, 2002; Gepshtein & Banks, 2003; Knill & Saunders, 2003; Landy & Kojima, 2001; Hillis, Watt, Landy & Banks, 2004). Here we show that if the MLE, Bayesian, or equivalent models are valid characterizations of cue combination (as they seem to be), then Euclidean geometry cannot describe our perceptual experience. Rather, our perceptual spatial structure would be better described, as best, as belonging to an arbitrarily curved Riemannian space.

Summing up, there is ample evidence that perceptual space is not Euclidean, though there is still no consensus in the scientific community about this. As previously mentioned, many authors still treat or make the assumption that perceptual space is Euclidean. The intention of this article is to take a different approach from previous work by showing in a more explicit way why perceptual space cannot be Euclidean; the direct empirical evidence on this question is not our point. Instead, we assume first that MLE is a valid approach, based on empirical evidence, and then show that if this is so, then perceptual space is not Euclidean. In short, we will show that a Euclidean perceptual space is not compatible with MLE.

We can summarize the logic of this article as: 1) Mathematically, the shape obtained when two Euclidean shapes are combined by means of an average—an average that can be taken not just over depth values, but over any observable quantity, such as an angle—will only be Euclidean if the weights have very specific values. 2) Maximum Likelihood Estimated (MLE) weights are determined by reliabilities (specifically, the reliability of a given cue in predicting a variable), not by geometrical constraints; one proof of this is that perceptual learning can change the weights of a given depth cue (see, e.g., Jacobs & Fine [1999]). 3) Conclusion: MLE is not Euclidean (except by chance).

Analysis of the inconsistency between MLE and Euclidean perceptual space

Cue combination as described by MLE and Bayesian models seems to be in close agreement with whatever procedure our perceptual system uses. They provide a rule to obtain the final estimate of a perceptual variable as the weighted average of the estimates of the individual cues. For instance, our perception of depth might combine information from motion parallax, binocular disparity, shading, texture, perspective, and so forth.

Let us assume that we have a series of variables x_i (n in total), each of which represents a measurable perceptual property (distance between two points, angle, etc). Individual cues, represented by the index j (m in total), each provide an estimate, $\hat{x}_i^j$, for each variable x_i . The estimates given by each cue are, in general, different from one another. A linear depth cue combination model states that the observed value, $\hat{x}_i^{obs}$, will be:

$$\hat{x}_i^{obs} = \sum_{j=1}^m \omega_i^j \hat{x}_i^j \quad (1)$$

where the weights ω_i^j , satisfy

$$\sum_{j=1}^m \omega_i^j = 1. \quad (2)$$

(Note that MLE estimates are not necessarily minimum variance—but they are minimum variance estimates whenever the noise in the cues involved is uncorrelated (Cochran, 1963)—and not necessarily unbiased. In what follows we will assume for simplicity that these estimates are unbiased, so Eq. 1 is valid. This will not be a limitation to the generality of our results: If a different strategy other than weighted averages were used for depth-cue combination, we would still have a constraint to be satisfied, albeit this will differ from Eq. 1. The results that follow are due to the existence of a constraint in itself, regardless of its specific format, that is, whether it is given by Eq. 1 or by some other cue-combination rule. For instance, the weighted sum model is equivalent to a Bayesian model with a flat prior and flat cost function.)

In principle, there could be any number of variables and any number of cues; each additional one will increase the number of weights and thus the complexity of the problem. Thus, let us limit our analysis to the simplest case, the two-variable, two-cue scenario³. To show that nothing essential is lost by this constraint, notice that there could be at most $n-1$ independent geometric relations among the variables x_i . (A simple example of a geometric relation would be $x_1 = 2\pi x_2$, where x_1 is the circumference of a circle and x_2 is the radius.) Let us write them as:

$$\begin{aligned} x_1 &= f_1(x_2, \dots, x_n), \\ x_2 &= f_2(x_3, \dots, x_n), \\ &\dots \\ x_i &= f_i(x_{i+1}, \dots, x_n), \\ &\dots \\ x_{n-2} &= f_{n-2}(x_{n-1}, x_n), \\ x_{n-1} &= f_{n-1}(x_n), \end{aligned} \quad (3)$$

By substituting the equation at the bottom into the one above it, and then repeating this procedure until we reach the equation located at the top, this can be reduced to the $n-1$ relations:

$$\begin{aligned} x_1 &= g_1(x_n), \\ &\dots \\ x_i &= g_i(x_n), \\ &\dots \\ x_{n-1} &= g_{n-1}(x_n), \end{aligned} \quad (4)$$

Here $g_{n-1}(x_n) = f_{n-1}(x_n)$, $g_{n-2}(x_n) = f_{n-2}(f_{n-1}(x_n), x_n)$, and so on. This shows that the variables in Eqs. 4 are related only in pairs, and thus we can restrict ourselves, without loss of generality, to the analysis of the two-variable scenario⁴. Let us call the two variables x and y . Then $n=2$, and the set of $n-1$ relations shown in Eqs. 4 become the single equation

³The introduction of additional variables into a single relation introduces additional weights but adds nothing new to the analysis or the conclusions that will be obtained here, which in the end are the result of having variables subject to sets of independent constraints (that is, MLE and geometric constraints).

$$x=g(y). \quad (5)$$

Let us assume the best-case scenario, in which the structure recovered from each individual cue, j , is Euclidean, and thus all cues satisfy the geometric relation given by Eq. 5:

$$\widehat{x}^j = g(\widehat{y}^j) \quad (6)$$

The question we ask is: Does the geometric relation given by Eq. 5 still hold for the final percept? In other words, does the equation

$$\widehat{x}^{obs} = g(\widehat{y}^{obs}) \quad (7)$$

hold?

As mentioned, we will restrict our analysis to the simplest, two-cue scenario. Call the two cues M and S (e.g., motion and stereo). Let us assume that Eq. 7 is valid. Then substituting Eq. 1 into Eq. 7 we get:

$$\begin{aligned} \omega_x^M \widehat{x}^M + (1 - \omega_x^M) \widehat{x}^S \\ = g[\omega_y^M \widehat{y}^M + (1 - \omega_y^M) \widehat{y}^S] \end{aligned} \quad (8)$$

from which we obtain:

$$\begin{aligned} \omega_x^M &= \frac{g[\omega_y^M \widehat{y}^M + (1 - \omega_y^M) \widehat{y}^S] - \widehat{x}^S}{\widehat{x}^M - \widehat{x}^S} \\ &= \frac{\widetilde{x} - \widehat{x}^S}{\widehat{x}^M - \widehat{x}^S} \end{aligned} \quad (9)$$

$$\widetilde{x} = g[\omega_y^M \widehat{y}^M + (1 - \omega_y^M) \widehat{y}^S] = g(\widehat{y}^{obs}).$$

where

Now, for Eq. 7 to be true, the value of ω_x^M given by Eq. 9 must lie in the interval $[0,1]$ for any value of $\omega_y^M \in [0,1]$. (Notice that, when $\omega_x^M = \omega_y^M$, Eq. 8 reduces to Jensen's inequality [see e.g., Hardy, Littlewood & Polya, 1934]. In such a case, the left side of Eq. 8 is usually different than the right side; they are equal only if the function g satisfies certain conditions). It is easy to see that if $\omega_y^M = 0$, then $\omega_x^M = 0$, and if $\omega_y^M = 1$, then $\omega_x^M = 1$. For intermediate values of ω_y^M , this condition is equivalent to:

⁴We will not analyze the case of fewer than $n-1$ relations, which is not reducible to be a function of a single variable, but instead will be a function of more than one. Again, the introduction of additional variables into a single relation introduces additional weights but adds nothing new to the analysis or the conclusions that will be obtained here.

$$\begin{aligned}
 \hat{x}^S < \tilde{x} < \hat{x}^M & \quad \text{if } \hat{x}^S < \hat{x}^M \\
 \hat{x}^M < \tilde{x} < \hat{x}^S & \quad \text{if } \hat{x}^M < \hat{x}^S
 \end{aligned}
 \tag{10}$$

If $x=g(y)$ is monotonic (see Fig. 1a), the inequality given in Eq. 10 implies that there is a single solution to Eq. 9 and thus, given $\hat{x}^S, \hat{x}^M, \hat{y}^S$ and $\hat{y}^M$, Eq. 7 is valid only for a single pair of weights ω_x^M and ω_y^M . If $x=g(y)$ is not monotonic, then there are values of $\hat{x}^S, \hat{x}^M, \hat{y}^S$ and $\hat{y}^M$ for which the inequality given in Eq. 10 is not satisfied (see Fig. 1b) and thus Eq. 7 is not true. In this later case the percept cannot be Euclidean.

Let us concentrate then on the cases in which the Euclidean constraint could still be satisfied by having adequate pair of weights: the monotonic case and the non-monotonic case with values of $\hat{x}^S, \hat{x}^M, \hat{y}^S$ and $\hat{y}^M$ for which the inequality given in Eq. 10 is satisfied. For these cases, Eq. 9 imposes a Euclidean constraint on perception by establishing the relationship that must hold between two parameters (the weights given to motion for two different variables, in this example). MLE as well as Bayesian and similar methods (and whatever method our visual system uses to combine cues) impose a constraint on the weights, too. These two constraints are independent: the Euclidean constraint is given by geometry, and the MLE constraint is related to the reliability of a given cue in predicting a variable. Let us look at the Cartesian space in which the two weights, ω_x^M and ω_y^M , represent the two coordinate axes. Cue combination constrains these two weights to take on a specific pair of values within the interval $[0,1]$, the values being a function of the reliabilities not only of the motion cue but also of the other cues for each of the two properties, x and y , being estimated. Such a pair of weight values is shown as point $(\omega_y^M_{MLE}, \omega_x^M_{MLE})$ in Cartesian space of Figure 2. The Euclidean constraint from Eq. 9 will result in a curve in this space, also shown in Figure 2. Let's assume that the actual value for ω_y^M agrees with the value prescribed by the depth cue combination rule. For the Euclidean constraint to be satisfied, ω_x^M has to take a specific value on the curve $(\omega_x^M_{Euclidean}$ in Fig. 2). It could happen that the Euclidean constraint curve includes the point $(\omega_y^M_{MLE}, \omega_x^M_{MLE})$ by chance; if so, then $\omega_x^M_{Euclidean} = \omega_x^M_{MLE}$. But this will not happen in general. In general, these two constraints will not give rise to the same weight values and will not be valid simultaneously. The reason is that, the probability that two independent constraints result in the same weights is zero. There still remains the question of whether the departure from Euclidean (as measured by the difference between the two weights) would be large enough to be detectable and thus leading to practical consequences. The answer to this question seems to be positive. A detailed example for the case of surface slant estimation is question seems to be positive. A detailed example for the case of surface slant estimation is analyzed in the *Appendix*.

Summing up, various studies have tested and confirmed the quantitative predictions of the MLE model for different cues, as indicated in the Introduction. The MLE constraint is independent of the Euclidean constraint given by geometry. The implication, therefore, is that perception, in general, is not Euclidean.

Implicit in this conclusion is the assumption that there is always more than one cue available. This seems like a valid assumption. There is the possibility that the visual system imposes a veto on all but one cue (for example, the most reliable one). In such an extreme case (robust estimation, see Landy *et al.*, 1995) there is no cue-combination constraint to violate the Euclidean geometry constraint. However, the single-cue case is suboptimal in general, inconsistent with most laboratory evidence, and therefore not likely to represent the way our perceptual system works in daily life.

Discussion

We have shown that even if individual perceptual cues have a Euclidean structure, the final percept that emerges from cue combination will not, in general, be Euclidean. As an intuitive example, let us assume that the two variables to be estimated are the radius, r , and the circumference, c , of a circle. Let us also assume that the recovered structure from each of the two cues, M and S , is also a circle, and that the estimates of r and c obtained from each cue differ. If both M and S provide an Euclidean percept, then $\hat{c}^M = 2\pi\hat{r}^M$ and $\hat{c}^S = 2\pi\hat{r}^S$. Assuming that the final percept is also be a circle, the perceived values for c and r will be (e.g., from MLE) $\hat{c}^{obs} = \omega_c^M \hat{c}^M + (1 - \omega_c^M) \hat{c}^S$ and $\hat{r}^{obs} = \omega_r^M \hat{r}^M + (1 - \omega_r^M) \hat{r}^S$. In general, however, these weights will not have values that satisfy the original relationship $c = 2\pi r$, and so $\hat{c}^{obs} \neq 2\pi\hat{r}^{obs}$. The circle will inhabit a curved Riemannian space whose local curvature will depend on the particular values obtained for $\hat{c}^{obs}$ and $\hat{r}^{obs}$.

Note that a percept that is non-Euclidean will not appear “strange” to the observer. An object will not pop-out just because its structure is non-Euclidean. A non-Euclidean slanted plane will still look like a Euclidean slanted plane, and a non-Euclidean circle will still look like a Euclidean circle. The only measurable effect will be that $\pi^{obs} \neq \pi$, but π in itself is not directly perceivable. However, in the case of the perception of surface slant (see *Appendix*), the magnitude of the departure from Euclidean structure seems to be large enough to be measured experimentally.

Our intention here was to show that perception in general cannot be Euclidean, and that “at best” it is Riemannian. This best-case scenario is that in which: 1) Perception is consistent and 2) perception has an internal metric, and 3) perception based on each individual cue is Euclidean. Thus, we did not to prove that perception is Riemannian, but only that it cannot be Euclidean. There are other scenarios, including those that might be equaled as “worst case”. For instance, if different judgments of space were inconsistent with one another then, not only they will not be Euclidean, but they will not even be Riemannian; there will be no consistent internal geometry in such a case.

It is worth noting, however, that the MLE models do not need to, and generally do not, assume a Euclidean perceptual representation. We've shown that they couldn't validly assume it. Most of the studies found in the literature do not make any assumptions about the internal geometry of perceptual space. Some of them claim otherwise, but in practice they do not test the assumption, because they usually measure only a single variable (e.g., perceived slant). They do not make any measurements that require, or test, a geometrical relationship between two different variables. Thus, most MLE results found in the literature are actually independent of the assumption of an internal Euclidean metric. Similarly, it is not necessary to assume a Euclidean metric to resolve issues of scale incompatibility between the cues. Landy et al. (1995) discussed the need to get estimates into the same units when combining cues that differ in scale. Landy et al. proposed ways of ‘promoting’ cues to metric scales, but cue promotion can be applied to any metrical space (that is, to any Riemannian space), whether the metric is Euclidean or not.

A few studies, however—already discussed in the previous paragraph—do assume a Euclidean relationship between variables, and thus their results should be interpreted with skepticism (but these studies are not tests of MLE, so the issue of the validity of MLE is not at stake in them).

All these cue-combination studies, and our demonstration of their implications for the geometry of internal perceptual space, seem to make an unsupported, and possibly implausible, assumption: that humans can readily estimate two or more geometric properties simultaneously. Perhaps the estimations can only be made successively, each constituting, in

effect, a separate task. We assumed here the best-case scenario, in which perceptual space remains stable over time, regardless of task. If this were not true, then perceptual space would be so unstable that it would make little sense to assume that it had an internal geometry at all; the statement that perceptual space is non-Euclidean would be true without need for further argument.

Finally, it is worth noting that our demonstration is very general. It applies not only to 3D shape in vision, but also to the structure of space in general, regardless of what kind of cues lay behind its recovery—whether they come from the same modality, e.g. vision, or from different modalities, e.g., vision and touch.

Acknowledgments

This research was supported by NEI Grants EY015673 (J.M.F.) and EY012286 (B.F.).

We thank the two anonymous reviewers for their valuable suggestions

Appendix

Example: surface slant estimation

As a concrete example, let us chose our variables to be the perceived slant of a planar surface, σ , and the depth difference (measured in a direction parallel to the line of sight) between two particular points on the surface, ΔZ . We will assume the surface has a tilt of zero; in other words, iso-distant lines on the surface are horizontal. In an Euclidean space, the two variables, perceived slant and depth difference, are related as:

$$\frac{\Delta Z}{Z_0} = \frac{\Delta y}{\sigma} \quad (\text{A1})$$

where Δy is the retinal vertical separation between the points, and Z_0 the distance between the observer and the surface. Eq. A1 is a concrete example of a geometrical relationship, as represented by Eq. 5 in the text. Let us assume that we have two depth cues, which could be, for example, texture, T , and binocular disparity, D . If the internal structure of the percepts from each of these cues is Euclidean, then they satisfy:

$$\Delta \hat{z}^T = \frac{\Delta y}{\sigma^T} \quad \text{and} \quad \Delta \hat{z}^D = \frac{\Delta y}{\sigma^D} \quad (\text{A2})$$

were distances were normalized as $z = Z/Z_0$.

Based on our previous analysis, we predict that the perceived object containing the two cues T and D will depart from the Euclidean expression:

$$\Delta \hat{z}^{obs} = \frac{\Delta y}{\sigma^{obs}} \quad \text{where} \quad \Delta \hat{z}^{obs} = \omega_z^T \Delta \hat{z}^T + (1 - \omega_z^T) \Delta \hat{z}^D \quad \text{and} \quad \hat{\sigma}^{obs} = \omega_\sigma^T \hat{\sigma}^T + (1 - \omega_\sigma^T) \hat{\sigma}^D \quad (\text{A3})$$

The ratio $r = \Delta \hat{z}^{obs} / (\frac{\Delta y}{\hat{\sigma}^{obs}})$ gives a measure of the departure from Euclidean perception. It is equal to 1.0 in the Euclidean case (see Eq. A3), and the more r departs from 1.0, the more non-Euclidean the percept will become.

What we want to test here is not whether the cue combination percept departs from being Euclidean, which has been demonstrated mathematically, but rather whether the departure could be large enough to be measurable and thus lead to practical consequences. Unfortunately, there are no experiments in the literature that simultaneously measured perceptual estimates of $\Delta \hat{z}$ and $\hat{\sigma}$, for either a single cue or a combined-cue stimulus. Thus, it is not possible at this time to test how well psychophysical data agree with our predictions. But, even if the available data are not complete, we can still obtain some order-of-magnitude estimates using the data that are available. Specifically, we can use data from the literature to obtain estimates of r . Using the definition of r and Eqs. A2 and A3, we can rewrite r as:

$$r = (1 + \omega_z^T \delta_2)(1 - \omega_z^T \delta_1) \quad \text{where} \\ \delta_1 = \frac{\hat{\sigma}^T - \hat{\sigma}^D}{\hat{\sigma}^T} \quad \text{and} \quad \delta_2 = \frac{\hat{\sigma}^T - \hat{\sigma}^D}{\hat{\sigma}^D}. \quad (\text{A4})$$

For any pair of values of our choice for δ_1 and δ_2 , we can obtain reliable values for ω_z^T from the work of Hillis *et al.* (2004). Hillis *et al.* (2004) measured the reliability of slant estimation across a range of slant values for two cues, texture and stereo, and from these obtained predicted MLE weights for the cue combination condition. They showed that the actual weights used by human observers do agree with the MLE predictions for this slant estimation task. From Fig. 5 of Hillis *et al.* (2004) we can see that the JNDs (just noticeable differences) for slant discrimination are about the same for texture and stereo cues when the slants are about 30° for texture and 0° for stereo. This happens when the viewing distance is about 100 cm. Accurate values are not very important, as we will see shortly. The important issue here is that there are many instances for which the JND's (and thus the weights, ω_z^T and ω_z^D) of the two cues are the same (and equal to 0.5), but the slants provided by the two cues differ. In what follows we will assume that perceived and simulated slants for individual cues are about the same, although it is known that depth from both of these cues is not always veridically perceived. This approximation will allow us to use the simulated slant values instead of the perceived ones, $\hat{\sigma}^T$ and $\hat{\sigma}^D$, in Eq. A4.

The only parameter needed for estimating r that we lack is ω_z^T . We do not have data about this weight, so Fig. 3 shows r as a function of ω_z^T for the example mentioned above. Values used are $\hat{\sigma}^T = \tan(30^\circ)$, $\hat{\sigma}^D = \tan(0.01^\circ)$ and $\omega_z^T = .5$.

Fig. 3 shows that r departs very considerably from 1.0 for any reasonable value of ω_z^T : only for very small values of ω_z^T (equivalent to a veto of the texture cue for depth difference estimation when stereo is present) is r close to one. More dramatic departures of r from unity are found when using examples in which the slant difference between the texture and the stereo cues are larger, or in examples where $\hat{\sigma}^D$ is closer to zero. With the value of r estimated to differ by a large amount from 1.0, departures from Euclidean structure should in fact be measurable psychophysically.

References

- Alais D, Burr D. The ventriloquist effect results from near-optimal bimodal integration. *Current Biology* 2004;14:257–262. [PubMed: 14761661]
- Blank AA. Analysis of experiments in binocular space perception. *JOSA* 1958;48:911–925.

- Blank, AA. Metric geometry in human binocular perception: Theory and fact. In: Leeuwenberg, ELJ.; Buffart, HFJM., editors. *Formal theories of visual perception*. New York: Wiley; 1978. p. 83-102.
- Blumfeld W. Unsterchunger uber die scheinbare grosse im sehraume. *Zeitschrift fur Psychologie* 1913;65:241–404.
- Clark, JJ.; Yuille, AL. *Data Fusion for Sensory Information Processing Systems*. Kluwer; 1990.
- Cochran, WG. *Sampling Techniques*. 2nd. Wiley; 1963.
- Cuijpers RH, Kappers AML, Koenderink JJ. On the role of external reference frames on visual judgments of parallelity. *Acta Psychologica* 2001;108:283–302. [PubMed: 11693253]
- Domini F, Braunstein ML. Recovery of 3-D structure from motion is neither Euclidean nor affine. *J Exp Psychol : Human Perception and Performance* 1998;24(4):1273–1295.
- Domini F, Caudek C. 3-D structure perceived from dynamic information: a new theory. *Trends in Cognitive Sciences* 2003;10:444–449. [PubMed: 14550491]
- Domini F, Caudek C, Richman S. Distortions of depth-order relations and parallelism in structure from motion. *Perception & Psychophysics* 1998;60:1164–1174. [PubMed: 9821778]
- Domini F, Caudek C, Tassinari H. Stereo and motion information are not independently processed by the visual system. *Vision Research* 2006;46:1707–1723. [PubMed: 16412492]
- Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 2002;415:429–433. [PubMed: 11807554]
- Gepshtein S, Banks MS. Viewing geometry determines how vision and haptics combine in size perception. *Current Biology* 2003;13:483–488. [PubMed: 12646130]
- Gibson, JJ. *The ecological approach to visual perception*. Boston: Houghton Mifflin; 1979.
- Gilinski AS. Perceived size and distance in visual space. *Psychological Review* 1951;87:411–434.
- Hardy, GH.; Littlewood, JE.; Polya, G. *Inequalities*. Cambridge (England): The University Press; 1934.
- Helmholtz, H von. *Helmholtz's treatise on physiological optics*. Southall, JPC., translator. New York: Dover; 1962. original work published 1866
- Hillis JM, Watt SJ, Landy MS, Banks MS. Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision* 2004;4:967–992. [PubMed: 15669906]
- Howard, IP.; Rogers, BJ. *Seeing in depth*. Vol. 1. I. Porteous; Ontario, Canada: 2002.
- Indow T, Watanabe T. Parallel- and distance-alleys with moving points in the horizontal plane. *Perception and Psychophysics* 1984;35:144–154. [PubMed: 6718211]
- Kant, I. *Critique of pure reason*. New York: American Home Library Company; 1902.
- Kersten D, Mamassian P, Yuille A. Object perception as Bayesian inference. *Annual Rev Psychol* 2004;55:271–304. [PubMed: 14744217]
- Klein, F. *Vorlesungen uber hohere geometrie*. 3rd. [Lectures on higher geometry]. New York: Chelsea: 1957. Original work published 1893
- Knill DC, Saunders JA. Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research* 2003;43:2539–2558. [PubMed: 13129541]
- Koenderink JJ, van Doorn AJ. Direct measurement of the curvature of visual space. *Perception* 2000;29:69–79. [PubMed: 10820592]
- Landy MS, Maloney LT, Johnston EB, Young MJ. Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research* 1995;35:389–412. [PubMed: 7892735]
- Landy MS, Kojima H. Ideal cue combination for localizing texture-defined edges. *Journal of the Optical Society of America A* 2001;18:2307–2320.
- Luneburg, RK. *Mathematical analysis of binocular vision*. Princeton, NJ: Princeton University Press; 1947.
- Norman, JF.; Todd, JT. The visual perception of 3-dimensional form. In: Carpenter, GA.; Grossberg, S., editors. *Neural networks for vision and image processing*. Cambridge, MA: MIT press; 1992. p. 93-110.
- Norman JF, Todd JT. The Perceptual analysis of structure from motion for rotating objects undergoing affine stretching transformations. *Perception & Psychophysics* 1993;3:279–291. [PubMed: 8483692]
- Ogle, KN. *Researches in binocular vision*. New York: Hafner Publishing Company; 1964.

- Oruc I, Maloney LT, Landy MS. Weighted linear cue combination with possibly correlated error. *Vision Res* 2003;43:2451–2468. [PubMed: 12972395]
- Petersen, P. *Riemannian Geometry*. Berlin: Springer-Verlag; 1990.
- Tittle JS, Todd JT, Perotti VJ, Norman JF. The systematic distortion of perceived 3D structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception and Performance* 1995;21:663–678. [PubMed: 7790840]
- Todd JT, Bressan P. The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Perception & Psychophysics* 1990;48:419–430. [PubMed: 2247324]
- Todd JT, Norman JF. The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception & Psychophysics* 1991;50:509–523. [PubMed: 1780199]
- Todd JT, Oomes AHJ, Koenderink JJ, Kappers AML. On the affine structure of perceptual space. *Psychological Science* 2001;12:191–196. [PubMed: 11437300]
- Todd JT, Reichel FD. Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review* 1989;96:643–657. [PubMed: 2798652]

Figure 1.

a) If $x=g(y)$ is a monotonic function, then the condition expressed in Eq. 10 is fulfilled: $\tilde{x}$ always falls between $\hat{x}^M$ and $\hat{x}^S$. b) If $x=g(y)$ is a non-monotonic function, then the condition expressed in Eq. 10 is not valid: there are values of $\tilde{x}$ that do not fall between $\hat{x}^M$ and $\hat{x}^S$.

Figure 2.

The isolated dot represents the weights prescribed by MLE or any other cue-combination process. The solid curve represents the Euclidean constraint. In general, the MLE solution and the Euclidean curve will not intersect, except by chance. If they do intersect, the weights from both would agree, that is, $\omega_{x \text{ Euclidean}}^M = \omega_{x \text{ MLE}}^M$.

Figure 3. r as a function of ω_z^T for the parameter values discussed in the main text extracted from Hillis et al (2004). Any significant departure of r from unity implies a significant departure from an Euclidean structure.