# On the relative detectability of configural properties

**James T. Todd** The Ohio State University, Columbus, OH, USA

**Eric Weismantel** The Ohio State University, Columbus, OH, USA

**Christopher S. Kallie** The Ohio State University, Columbus, OH, USA

**A match-to-sample shape-discrimination task was employed to measure the detectability of different types of transformations. To create the foils for this task, the standard object could be altered by adding a small hole (a topological property), adding small bumps to straight edges (a projective property), changing the relative orientations of parallel contours (an affine property), or stretching the standard object to alter its aspect ratio (a Euclidean property). The results revealed that the relative perceptual salience of different types of shape change is consistent with the Klein hierarchy of geometries. That is to say observers were most sensitive to changes in topological structure, followed by changes in projective, affine, and Euclidean structure, respectively. The predicted patterns of performance among the different conditions were computed using a wide variety of commonly used shape-difference metrics, but none of them had a significant positive correlation with the observers' thresholds.**

## Introduction

It has long been recognized that some configurations of contours are easier to detect than others. A classic demonstration of this is shown in Figure 1, which was first reported by Pomerantz, Sager, and Stoever (1977). Observers in their experiment were required to indicate the odd man out in an array of four figures by moving a joystick in the appropriate direction as quickly as possible. Note in Figure 1 that all four quadrants of each array are identical except for the orientation of a single line. Although the other contextual contours by themselves provided no useful information, they sometimes had a dramatic effect on performance. Pomerantz et al. argued that emergent features can be formed when contextual contours are combined with the target lines and that some of these features are more discriminable than their component parts presented in isolation (see also Pomerantz & Portillo, 2011). They described this phenomenon as a configural superiority effect.

Why should some differences in contour structure be easier to detect than others? Lin Chen (1982, 1983, 2005) has proposed an interesting theoretical hypothesis that the relative perceptual salience of different possible attributes of an object may be systematically related to their structural stability under change in a manner that is similar to the Klein hierarchy of geometries. According to this hypothesis, observers should be most sensitive to those aspects of an object's structure that remain invariant over the largest number of possible transformations. Thus, the easiest type of shape change to detect should be one that alters the connections between contours or the number of holes in an object, which remain invariant over all possible topological transformations. The next most salient shape changes should be those that alter a property that is invariant over all projective transformations, such as whether a contour is straight or curved. These, in turn, should be easier to detect than changes in properties that are invariant under affine transformations, such as whether a pair of contours is parallel or nonparallel. The least stable properties are those that are only invariant under Euclidean rigid body motions, such as the relative lengths of contours with different orientations (see also Todd & Bressan, 1990; Todd, Chen, & Norman, 1998). Note that this hypothesis provides an elegant account of the relative detectability of the odd men out in Figure 1. The target in Figure 1B is topologically distinct from the other items in the array whereas the targets in Figure 1A and C differ from the foils by the affine property of parallelism. It is important to point out in this context that there are several different topological properties that could be used to identify the target in Figure 1B. These include the number of terminators in the figures, the types of vertices that are present (i.e., Ls versus arrows), and the

Citation: Todd, J. T., Weismantel, E., & Kallie, C. S. (2014). On the relative detectability of configural properties. *Journal of Vision*, *14*(1):18, 1–8, http://www.journalofvision.org/content/14/1/18, doi:10.1167/14.1.18.

doi: 10.1167/14.1.18 Received October 28, 2013; published January 22, 2014 ISSN 1534-7362 © 2014 ARVO

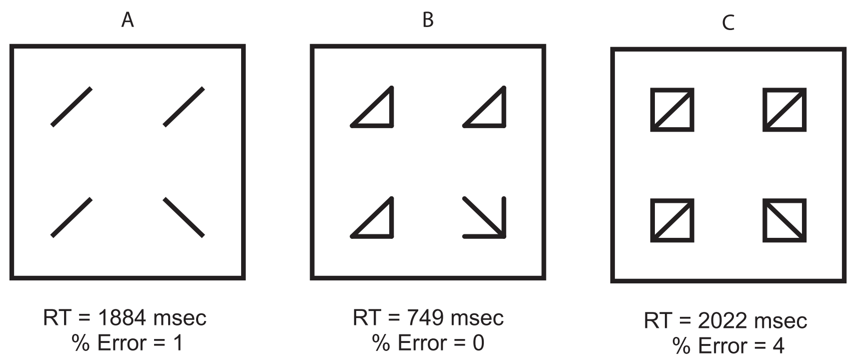


Figure 1. The configural superiority effect. Some example stimuli used by Pomerantz, Sager, and Stoever (1977) to demonstrate that some types of configural changes are easier to detect than others. Observers were required to indicate the odd man out in an array of four figures by moving a joystick in the appropriate direction as quickly as possible. Note that the reaction times for the stimuli in panels A and C were over 2.5 times slower than the one in panel B.

property of closure. Note that all of these properties are topological because they remain invariant under all possible topological transformations.

One possible reason why the human visual system might be biased toward object properties that are invariant under change is to facilitate shape constancy. There are a number of different problems in natural vision for which this approach might be particularly useful. Perhaps the most obvious of these is the inevitable distortion that occurs when a 3-D structure is optically projected onto a 2-D visual image. Lowe (1984, 1987) and Biederman (1987) have argued that properties that are invariant under optical projection, such as the connectivity patterns of contours, and the distinctions between straight versus curved or parallel versus nonparallel are important sources of information for object recognition. Although it is mathematically possible to adopt a viewpoint from which curved edges appear straight, nonparallel contours appear parallel, or unconnected edges appear to coterminate, the probability of encountering such a viewpoint by accident in a natural environment is vanishingly small. Thus, they are often referred to as nonaccidental properties. The advantage of Chen's approach, in our opinion, is that it is more firmly grounded in mathematics and that it highlights the fact that different nonaccidental properties are not all equally stable. For example, parallelism is not generally preserved in visual images unless objects are observed with small visual angles (see Todd et al., 2007). Under most natural viewing conditions, the optical projections of parallel lines will converge toward a vanishing point as is demonstrated quite clearly by linear perspective.

The research described in the present article was designed to evaluate Chen's hypothesis using a match-to-sample shape-discrimination task. To create the foils for this task, the standard object could be altered by adding a small hole (a topological property), adding small bumps to straight edges (a projective property), changing the relative orientations of parallel contours (an affine property), or stretching the standard object to alter its aspect ratio (a Euclidean property).

There is, however, a fundamental methodological problem that needed to be addressed in order to compare results from these different conditions. For any given type of transformation, it will always be the case that small amounts of change will be harder to detect than larger amounts of change unless performance has reached ceiling. How then do we compare two qualitatively different types of transformation? This can only be achieved if there is some sort of common currency for measuring the magnitude of each one (e.g., see Biederman & Bar, 1999; Vogels, Biederman, Bar, & Lorincz, 2001).

There are many possible metrics that could potentially be employed for this purpose. Although some of them are highly correlated with one another, others are not. Consider, for example, the three pairs of shapes depicted in Figure 2. The shape change shown in Figure

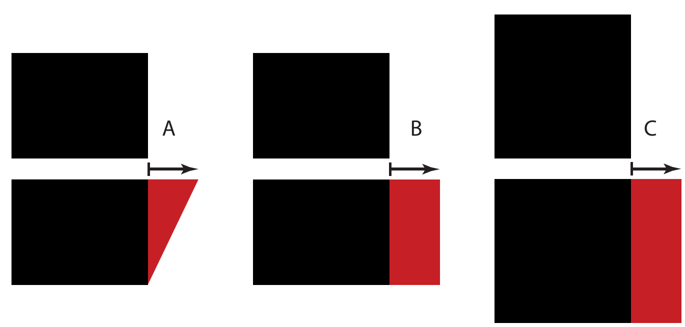


Figure 2. Three different types of shape change similar to those used in the present experiment. The one in panel A was created by displacing one vertex on the object in a horizontal direction whereas the ones in panels B and C were created by displacing an entire edge. The maximum displacement for each type of change is depicted with a black arrow, and the area of total displacement is highlighted in red.

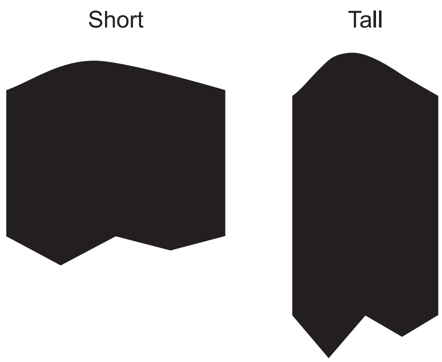


Figure 3. Two of the standard objects that were used in the present experiment. The same two objects could also be flipped horizontally or vertically in any given trial.

2A was created by displacing one vertex on the object in a horizontal direction. The ones in Figure 2B and C were created by displacing an entire edge. The only difference between them is that the displaced edge in Figure 2B is shorter than the one in Figure 2C. One possible measure, called the Hausdorf metric, scales shape changes based on the maximum displacement among all the different points on an object's boundary. According to that measure, all three of the shape changes in Figure 2 have exactly the same magnitude. Another way of scaling these changes might be to measure the total displacement summed over all the different points on the boundary. Note that this is quite different from the Hausdorf metric. When scaled based on the total displacement, the shape change in Figure 2A is only half the magnitude of the one in Figure 2B, and the displacement of the short edge in Figure 2B has a smaller magnitude than the displacement of the longer edge in Figure 2C. There are few empirical studies that have attempted to assess the relative psychological validity of possible shape-difference metrics (see, however, Wilbraham, Martinez, Christensen, & Todd, 2008). Thus, in the present investigation, we evaluated a wide variety of metrics in an effort to determine if any of them can predict the relative detectability the different types of shape change we employed.

# Methods

### *Observers*

Four observers participated in the experiment, including two of the authors and two others who were naïve about purpose of the experiment. All had normal or corrected-to-normal vision.

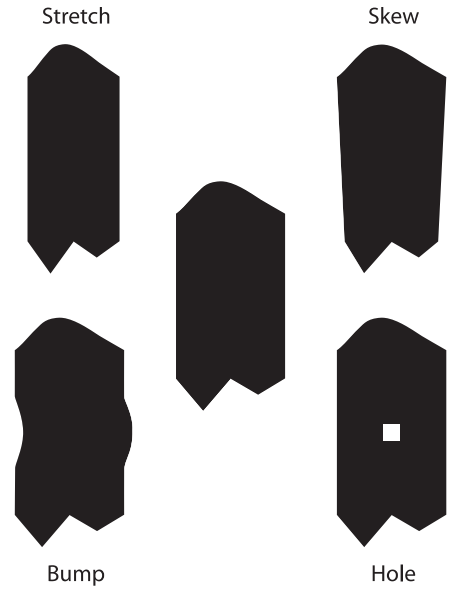


Figure 4. The four different types of shape change that were used in a match-to-sample discrimination task. All of the depicted changes have a maximum displacement of 15 pixels.

### *Apparatus*

The experiment was controlled using a Dell Dimension 8300 computer with a 21-in. monitor. The monitor was viewed at a distance of 65 in., and it had a spatial resolution of 1280 × 1024 pixels.

### *Stimuli*

Two of the possible standard objects are shown in Figure 3 and will be referred to hereafter as the short and tall object, respectively. The manipulation of object height was included because it affects some shape-difference metrics as shown in Figure 2. The same two objects could also be flipped horizontally or vertically in any given trial. Four different transformations were performed on these objects to create possible foils for a match-to-sample discrimination task as shown in Figure 4: A stretching transformation could be performed that expanded or compressed the objects along the horizontal axis. A skewing transformation could be performed that caused the parallel vertical edges to taper inward or outward. Small bumps could be added to both vertical edges. The heights of these bumps were one third the heights of the vertical edges, and their maximum displacements could be either rightward or leftward. Finally, a small hole could be added to the center of the object. The height of this hole was 19 pixels for the short object and 28 pixels for the tall object.

All these transformations were parameterized in terms of their maximum displacements (i.e., the Hausdorf metric). Based on the results of pilot

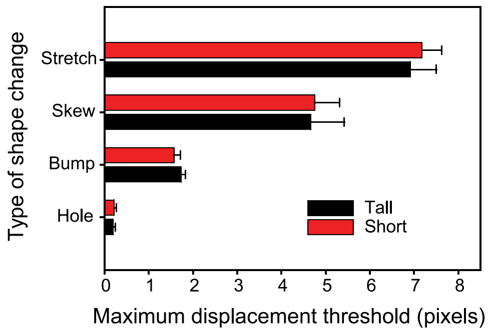


Figure 5. The mean discrimination thresholds for the short and tall standards with each type of shape change. These results are nearly identical to those obtained for individual observers.

experiments, we selected displacements for each transformation to obtain a range of accuracy from just above 50% to just below 100%. For the stretching transformation, the possible displacements were 4, 7, 10, and 13 pixels. For the skew transformation, they were 1, 4, 7, and 10 pixels, and for the bump condition, they were 1, 2, 3, and 4 pixels. A problem arose in selecting displacements for the hole condition because observers were 100% accurate when the hole had a width of only 1 pixel. In order to get psychometric functions for this condition, we needed to use subpixel displacements. The basic idea behind this technique is to mimic the optical effect of a higher-resolution display by manipulating the intensities of individual pixels. For example, 1 pixel at one fourth the maximum intensity would be optically equivalent to 1 pixel at full intensity on a display with four times the display resolution. The values selected for the hole condition to obtain an appropriate range of performance were 0.16, 0.18, 0.20, and 0.22 pixels. To summarize the overall experimental design, there were 32 distinct conditions, including four possible magnitudes of the four possible types of shape change for both the tall and short standard object.

### *Procedure*

Observers' discrimination thresholds were determined using a sequential match-to-sample task. Each trial began with a fixation point presented at the center of the display screen. The sample shape was then presented at a small random displacement from the center of the display screen for 300 ms, followed by a pattern mask for 300 ms. Next, the first test shape was presented on the left side of the display screen for 300 ms, followed by another pattern mask for 300 ms. The second test shape was then presented on the right side of the display screen for 300 ms, followed by a final pattern mask. All of the shapes presented in each trial were rotated about their centers by a random amount between $\pm 10°$ and $\pm 35°$. Observers were asked to indicate which test object had the same shape as the standard by pressing an appropriate response key, and the word "correct" or "incorrect" was presented in the center of the display screen to provide immediate feedback. A new trial could then be initiated by pressing the space bar. During their initial instructions, observers were shown examples of the different types of shape change that could occur at random in each trial. They participated in seven experimental sessions, each of which included eight presentations for each of the 32 possible conditions, resulting in a total of 256 trials per session. The first two of these sessions were considered as practice and excluded from subsequent analyses.

## Results

The four observers were remarkably consistent in their relative accuracy among the different conditions. Indeed, the correlations between each pair of observers were all above 0.95. For each type of shape change and each type of standard, we analyzed the results from the four possible magnitudes of shape change using logistic regression to estimate the 75% threshold. The averages obtained over all observers in each condition are shown in Figure 5. There are two important aspects of these results that deserve to be highlighted: First, there were no significant differences in the thresholds obtained for the short and tall standard objects. Second, there was a 20-fold difference in the thresholds obtained for the different types of shape change. This latter finding provides strong confirmation for Chen's hypothesis that the relative perceptual salience of object properties may be systematically related to their structural stability under change in a manner that is similar to the Klein hierarchy of geometries.

Is there some shape-difference metric that can account for this pattern of results? Suppose that observers' judgments are based on some hypothetical measure $\mu$, such that they can detect the difference between two shapes whenever this measure exceeds some critical threshold. When evaluated with respect to this metric, the thresholds obtained for different types of shape change should all be within measurement error of one another. Note that the Hausdorf metric does not satisfy this criterion, because the thresholds obtained for different types of shape change varied over a 20-fold range (see Figure 5).

That being said, the Hausdorf metric can still be used as a common currency for evaluating whether any other possible metric can account for the pattern of performance in this experiment. Consider the stretching transformation with the tall standard object. The maximum displacement threshold in this condition was approximately 7 pixels, and the total displacement at

| | Empirical thresholds | Edged-based measures | | | Pixel-based measures | | | Gabor-based measures | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Max. disp. | Avg. disp. | Total disp. | Correl. | Distance | Angle | Correl. | Distance | Angle |
| Stretch | | | | | | | | | | |
| Tall | 6.92 | 7.00 | 7.00 | 7.00 | 7.00 | 7.00 | 7.00 | 7.00 | 7.00 | 7.00 |
| Short | 7.18 | 7.00 | 7.01 | 10.11 | 10.12 | 10.13 | 10.14 | 9.17 | 9.18 | 9.72 |
| Skew | | | | | | | | | | |
| Tall | 4.66 | 7.00 | 14.42 | 14.42 | 15.79 | 15.78 | 15.80 | 15.89 | 15.90 | 16.43 |
| Short | 4.76 | 7.00 | 14.33 | 21.57 | 23.25 | 23.22 | 23.23 | 20.27 | 20.28 | 22.22 |
| Bump | | | | | | | | | | |
| Tall | 1.73 | 7.00 | 30.59 | 30.59 | 31.08 | 30.90 | 30.28 | 23.00 | 23.00 | 23.91 |
| Short | 1.56 | 7.00 | 32.15 | 48.05 | 48.38 | 43.86 | 43.14 | 25.84 | 25.68 | 27.85 |
| Hole | | | | | | | | | | |
| Tall | 0.20 | 7.00 | 80.99 | 80.99 | 79.25 | 80.99 | 79.26 | 28.43 | 28.44 | 23.97 |
| Short | 0.22 | 7.00 | 85.09 | 124.21 | 121.55 | 122.15 | 119.84 | 34.79 | 34.81 | 33.01 |
| Correlation with empirical thresholds | | 0.000 | −0.87 | −0.83 | −0.84 | −0.83 | −0.83 | −0.96 | −0.96 | −0.91 |

Table 1. Predicted maximum displacement thresholds for several possible metrics for measuring the magnitude of shape change. *Notes*: Almost all of the metrics have a high negative correlation with the empirical results.

that threshold was 4256 pixels. If total displacement is the appropriate metric for detecting shape differences, then its threshold should be approximately 4256 pixels in all of the other conditions as well. Thus, the predicted Hausdorf threshold in any given condition can be determined by calculating the maximum displacement that produces a total displacement of 4256 pixels.

We have applied this procedure to compute the predicted Hausdorf thresholds for eight possible shape-difference metrics that have been employed in the literature. Two of these, the average displacement and total displacement, involved establishing a point-to-point correspondence along the boundaries of the two shapes to be compared. Three other measures were designed to assess differences in pixel intensities between any given pair of images. Within this context, it is useful to consider an image as a vector in a high-dimensional space in which each individual pixel defines a dimension, and the intensity of the pixel defines a specific position along that dimension. The difference between two images can be computed using the angle or correlation between their corresponding vectors or the Euclidean distance between their corresponding vector endpoints. Three other measures were designed to model how image data are processed by simple cells within the primary visual cortex. Each image was convolved with an array of Gabor filters (i.e., Gaussian-modulated sine gratings) with six different orientations and nine different scales that were separated from one another by 1.7 octaves (see Lades et al., 1993 for more explicit details). The Gaussian envelope for each filter had a standard deviation that was 0.45 times the period of its sine grating. As with the pixel measurements, we can again consider an image as a vector in a high-dimensional space in which each individual Gabor filter defines a dimension, and the output of the filter defines a specific position along that dimension. The difference between two images can be computed using the angle or correlation between their corresponding vectors or the Euclidean distance between their corresponding vector endpoints.

The results of these analyses are shown in Table 1. It is important to keep in mind that many of the metrics in this table are used in a wide variety of applications for scaling the differences between pairs of images. However, the results of the present experiment reveal that they have no psychological validity for comparing different types of shape change. Indeed, the predicted maximum displacement thresholds were almost all negatively correlated with the empirical thresholds obtained for actual human observers. The one exception was the Hausdorf metric, whose predicted thresholds had a zero correlation with human performance. Note that most of the metrics predict that the tall standard should produce lower thresholds than the short standard, but there were no significant differences between the standards for the empirical thresholds. Similarly, all of the metrics except the Hausdorf predict that the ordering of thresholds from lowest to highest should be produced for the stretch, skew, bump, and hole conditions, respectively, but the empirical thresholds were the opposite of this prediction.

## Discussion

The concept of invariance under change has been of central importance to the development of modern

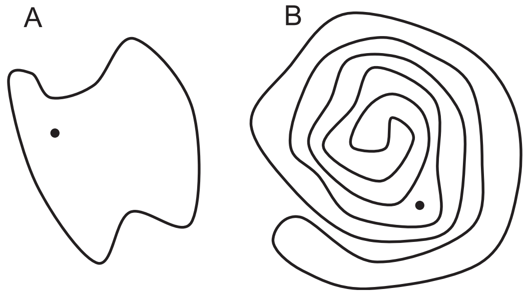


Figure 6. Two different patterns, which are topologically equivalent. For the one in panel A, it is trivial to determine that the dot is inside a closed boundary, but this judgment is quite difficult for the one in panel B. This phenomenon demonstrates that the "geometry" of visual perception may only be an approximation to the "geometry" of formal mathematics.

geometry, and there is a considerable amount of evidence to suggest that it is also of central importance to human perception. A fundamental fact of the natural environment is that all objects undergo change, and human observers possess the ability to identify objects over a wide range of possible transformations. It should not be surprising, therefore, that visual processes should focus primarily on those relatively stable properties that define the essence of an object's identity rather than those that change frequently under different viewing conditions. Consider, for example, the ability of normal individuals to identify their friends and loved ones. We are able to identify people with different facial expressions, hairstyles, makeup, hats, jewelry, glasses, patterns of illumination, or viewing directions. We can also identify people after they undergo growth or aging, gain or lose weight, or even when their faces are altered by cosmetic surgery. These observations suggest that the identity of an individual's face must be based on some remarkably abstract property that is somehow unaffected by all of the transformations that faces typically undergo in the natural environment.

There is another important reason why human observers might be most sensitive to attributes of objects that have a relatively high degree of stability. Although the natural environment is 3-D, the visual information by which objects are specified is confined to the 2-D projection surface of the retina. This projective mapping produces an inevitable loss of information, such that any given pattern of optical stimulation can have an infinite number of possible 3-D interpretations. Research has shown, however, that these ambiguities can be highly constrained in the sense that all of the possible interpretations are related by a limited class of transformations (see Koenderink, van Doorn, Kappers, & Todd, 2001; Todd & Bressan, 1990). Attributes of objects that are invariant over those transformations can be determined quite easily from the pattern of optical stimulation whereas other properties can only be estimated by incorporating additional ad hoc constraints. It makes sense, therefore, that the human visual system would concentrate its resources on object attributes that can be determined most reliably.

It is interesting to note that the present experiment is similar in some respects to a classic study by Biederman and Bar (1999). Observers in that study viewed pairs of objects presented in sequence and were asked to judge as quickly as possible whether their 3-D structures were the same or different. The "different" trials were divided into two conditions: In half of these trials, one object was created from the other by altering properties that are viewpoint invariant, including the coterminations of its contours (a topological property), whether its contours were straight or curved (a projective property), or whether they were parallel to one another (an affine property). In the remaining "different" trials, one object was created from the other by altering the lengths or curvatures of its contours without affecting their coterminations, linearity, or parallelism. Both conditions produced equal levels of performance when objects were presented at the same 3-D orientation, but when they were presented at different orientations, observers' responses were significantly more accurate and had shorter latencies for those pairs that differed in their viewpoint-invariant structure.

Unfortunately, Biederman and Bar (1999) did not analyze their data to see if there were also significant variations in performance among the different types of nonaccidental properties, so it is not possible to evaluate Chen's hypothesis about their relative detectability. The present experiment was designed in part to overcome this limitation, and the results demonstrate that there is a clear hierarchy in their relative perceptual salience. Our results agree with those of Biederman and Bar with respect to Euclidean metric properties. Changes in these properties are the hardest to detect, but there are also significant differences in the relative detectability of nonaccidental properties. Changes in topology are the easiest to detect, followed by changes in collinearity and parallelism, respectively. It should not be surprising that parallelism is the least salient of the nonaccidental properties. Although it is preserved under orthographic projection or small viewing angles under polar projection, it not should be considered a viewpoint-invariant property under most natural viewing conditions. In general, the optical projections of parallel lines in the environment will converge toward a vanishing point as is demonstrated quite clearly by linear perspective.

The Klein hierarchy of geometries provides a useful framework for assessing the invariance of object properties to various types of change, but it is probably best not to take it too literally as a model of human perception. Consider, for example, the topological property of whether a small probe dot is inside or outside a closed boundary. For the shape presented in Figure 6A, this type of judgment seems subjectively trivial and automatic. If, however, the boundary is made sufficiently convoluted as in Figure 6B, then the task becomes quite difficult—even though the two figures are topologically equivalent (see Minsky & Papert, 1969; Todd et al., 1998). This example is instructive because it suggests that that the "geometry" of visual perception may only be an approximation to the "geometry" of formal mathematics.

### Are measures of shape change psychologically valid?

Perhaps the most surprising aspect of our results is that the measures of shape change we employed were almost all negatively correlated with perceptual performance. Most of these measures predict that changes to the tall objects should be easier to detect than changes to the short objects, but the empirical results revealed no significant differences in the observers' thresholds. Similarly, all of the measures except the Hausdorf predict that the changes in the stretch condition should be the easiest to detect, followed by the skew, bump, and hole conditions, respectively, but the empirical thresholds were the opposite of this prediction. It is important to keep in mind when evaluating these results that the measures we employed are used frequently for a wide variety of applications that require scaling the differences between pairs of images (e.g., see Biederman & Bar, 1999; Vogels et al., 2001). Thus, the present results should be considered as a warning sign that the psychological validity of these measures may be questionable (see also Wilbraham et al., 2008).

An illuminating discussion about the concept of shape was published over 50 years ago by the theoretical geographer William Bunge (1962). Bunge argued that an adequate measure of shape should satisfy four criteria: (a) It should be objective; (b) it should not include less than shape, such as a set of position coordinates; (c) it should not include more than shape, such as fitting it with a Fourier or Taylor series; and (d) it should not do violence to our intuitive notions of what constitutes shape. Note that existing measures for scaling the differences between images do not satisfy criteria b or d. A set of pixel intensities or Gabor wavelet responses is a far cry from what most people would consider as shape, and the results of the present experiment show clearly that these measures cannot predict the relative salience of different types of shape change. Discovering a more psychologically valid metric for scaling shape changes will remain as an important problem for future research.

*Keywords: configural properties, perceptual organization, shape, contour*

## Acknowledgments

This research was supported by a grant from the National Science Foundation (BCS-0962119).

Commercial relationships: none.
Corresponding author: James T. Todd.
Email: todd.44@osu.edu.
Address: Department of Psychology, The Ohio State University, Columbus, OH, USA.

## References

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–147.

Biederman, I., & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. *Vision Research, 39,* 2885–2889.

Bunge, W. (1962). *Theoretical geography. First edition. Lund studies in geography series C: General and mathematical geography*. Lund, Sweden: Gleerup.

Chen, L. (1982). Topological structure in visual perception. *Science, 218,* 699–700.

Chen, L. (1983). *What are the units of figure perceptual representation? (Studies in cognitive science no. 22)*. Irvine, CA: University of California, Irvine, School of Social Sciences.

Chen, L. (2005). The topological approach to perceptual organization. *Visual Cognition, 12,* 553–637.

Koenderink, J. J., van Doorn, A. J., Kappers, A. M., & Todd, J. T. (2001). Ambiguity and the "mental eye" in pictorial relief. *Perception, 30,* 431–448.

Lades, M., Vorbruggen, C., Buhmann, J., Lange, J., von der Malsburg, C., Wurtz, R., & Konen, W. (1993). Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers, 42,* 300–311.

Lowe, D. (1984). *Perceptual organization and visual recognition*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.

Lowe, D. (1987). The viewpoint consistency constraint. *International Journal of Computer Vision, 1,* 57–72.

Minsky, M. L., & Papert, S. (1969). *Perceptrons: An* *introduction to computational geometry*. Cambridge, MA: MIT Press.

Pomerantz, J. R., & Portillo, M. C. (2011). Grouping and emergent features in vision: Toward a theory of basic Gestalts. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(5), 1331–1349.

Pomerantz, J. R., Sager, L. C., & Stoever, R. J. (1977). Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 422–435.

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., Chen, L., & Norman, J. F. (1998). On the relative salience of Euclidean, affine and topological structure for 3D form discrimination. *Perception, 27,* 273–282.

Todd, J. T., Thaler, L., Dijkstra, T. M. H., Koenderink, J. J., & Kappers, A. M. L. (2007). The effects of viewing angle, camera angle and sign of surface curvature on the perception of 3D shape from texture. *Journal of Vision, 7*(12):9, 1–16, http://www.journalofvision.org/content/7/12/9, doi:10.1167/7.12.9. [PubMed] [Article]

Vogels, R., Biederman, I., Bar, M., & Lorincz, A. (2001). Inferior temporal neurons show greater sensitivity to nonaccidental than metric differences. *Journal of Cognitive Neuroscience, 134,* 444–453.

Wilbraham, D. A., Martinez, A. M., Christensen, J. C., & Todd, J. T. (2008). Can low level image differences account for the ability of human observers to discriminate facial identity? *Journal of Vision, 8*(15):5, 1–12, http://www.journalofvision.org/content/8/15/5, doi:10.1167/8.15.5. [PubMed] [Article]