

Shape equivalence under perspective and projective transformations

JOHAN WAGEMANS, CHRISTIAN LAMOTE, and LUC VAN GOOL
University of Leuven, Leuven, Belgium

When a planar shape is viewed obliquely, it is deformed by a perspective deformation. If the visual system were to pick up geometrical invariants from such projections, these would necessarily be invariant under the wider class of projective transformations. To what extent can the visual system tell the difference between perspective and nonperspective but still projective deformations of shapes? To investigate this, observers were asked to indicate which of two test patterns most resembled a standard pattern. The test patterns were related to the standard pattern by a perspective or projective transformation, or they were completely unrelated. Performance was slightly better in a matching task with perspective and unrelated test patterns (92.6%) than in a projective-random matching task (88.8%). In a direct comparison, participants had a small preference (58.5%) for the perspective-related patterns over the projectively related ones. Preferences were based on the values of the transformation parameters (slant and shear). Hence, perspective and projective transformations yielded perceptual differences, but they were not treated in a categorically different manner by the human visual system.

When do two shapes look alike? Obviously, when you take a shape and translate, reflect, or rotate it, it will still be perceived as the same shape (except perhaps for some highly familiar shapes, such as the outline of the United States when presented in a strange orientation; see Rock, 1973). Similarly, an object seen at different distances, although producing retinal images of different sizes, is perceived as having the same shape. Moreover, an object and its cast shadow or its projected image are usually considered shape equivalent.

It is widely believed that shape constancy is based on the visual system's sensitivity to the geometric congruence under an increasingly larger set of operations in the above cases: from Euclidean, to similarity, to projective transformations. This kind of innate geometry can take the form of unconscious or automatic recovery of shape constancy based on implicit knowledge of the laws of geometry (e.g., Helmholtz, 1868/1968; Rock, 1983), or it can take the form of sensitivity to the invariants under each group of transformations (e.g., Cutting, 1986; Gibson, 1950).

How large the set of operations can be for shapes before and after transformation to be considered perceptually equivalent is not clear yet. This may seem strange be-

cause psychological experiments could test specific predictions derived from a strong mathematical framework: Felix Klein's (1872/1893) Erlangen program of geometry. Within this framework, geometry is defined as a system of definitions and theorems that remain invariant under a group of transformations (see Cutting, 1986, chap. 5, for further discussion). Also, different groups of transformations are ordered in a hierarchy of progressively "weaker" geometries in the sense that they have fewer and broader invariants, which give rise to progressively wider equivalence classes: from Euclidean, to similarity and affine, to projective geometry and topology (see Michaels & Carello, 1981, pp. 30–37, for an intuitive and pictorial presentation).

On the one hand, there is evidence suggesting that perceptual shape equivalence is not based on projective congruence. Niall and Macnamara (1990) presented convex pentagons with variable three-dimensional orientations and asked observers to select projectively equivalent matching patterns. The five-point cross ratios of the selected comparison figures differed significantly from those of the standard figures, which made the authors conclude that sensitivity to projective equivalence seems an implausible basis for shape constancy (later confirmed, by Niall, 1992, with shapes continuously rotating in depth). More recently, we demonstrated that human observers are able to determine shape equivalence under affine transformations (which are simpler than projective transformations), even when only minimal information is provided (i.e., four points; see Van Gool, Moons, Pauwels, & Wagemans, 1994, for more mathematical background) and no depth cues are available (Wagemans, De Troy, Van Gool, Foster, & Wood, 1994; see also Kukkonen, Foster, Wood, Wagemans, & Van Gool, 1996). In combination, these studies

The work presented in this paper was supported by an Esprit Basic Research Action ("VIVA," or "Viewpoint-Invariant Visual Acquisition"), Grant F.K.F.O. No. 2.0065.90 from the Belgian Fund for Basic Research, and Grant N.F.W.O. No. G.0210.97N from the Belgian National Fund for Scientific Research. The results have been presented in preliminary form at the 18th European Conference of Visual Perception in Tübingen, Germany (August 1995). The quality of the presentation has been improved thanks to helpful suggestions by James Cutting, Michael Kubovy, and two anonymous reviewers. Correspondence should be addressed to J. Wagemans, Department of Psychology, University of Leuven, Tiensestraat 102, B-3000 Leuven, Belgium (e-mail: johan.wagemans@kuleuven.ac.be).

might be taken to imply that the class of projective transformations is too wide.

On the other hand, there is evidence suggesting that the perceptual system can deal even with nonprojective, nonrigid shape deformations, such as cardioidal strain in the context of perceiving the relative age of faces and even Volkswagen Beetles (e.g., Mark & Todd, 1985; Mark, Todd, & Shaw, 1981; Pittenger & Shaw, 1975; Pittenger, Shaw, & Mark, 1979).

In this paper, we will have a closer look at shape equivalence under projective transformations. More specifically, we will ask whether all projective transformations play the same role for the visual system or, alternatively, whether one subset of these transformations has a special perceptual status. We need to provide some background in geometry and optics to further clarify this issue. Under normal viewing conditions, the geometry involved in projecting pencils of light from an object onto an eye generates a perspective distortion. Although the transformations involved in perspective projection are usually described as *projective transformations* or *projectivities* (e.g., Carlbohm & Paciorek, 1978), they constitute, in fact, only a subset of these, which would be better denoted as *perspective transformations* or *perspectivities*. The reason these are usually not distinguished is that the perspectivities do not form a group of their own (and, hence, do not have their own invariants): In general, when two perspectivities are combined, one no longer gets a perspectivity but gets a projectivity instead; in other words, the set of perspectivities does not obey the law of closure, which is one of the conditions to have a group in the mathematical sense.

We do not yet know whether this mathematical distinction has any perceptual consequences. On the one hand, a painting still looks alright when viewed from the incorrect projection center (e.g., Cutting, 1987b, 1988; Goldstein, 1987; Kubovy, 1986), a phenomenon called *the robustness of perspective* (see Ittelson, 1996, for a recent discussion). This suggests that the visual system can easily deal with combinations of two perspectivities and thus projective shape equivalence. On the other hand, Pizlo (1994, Figure 2) presented an example of a perspective transformed shape that still looked like the original shape and a more generally projectively transformed variant that did not. Pizlo and Rosenfeld (1992) proposed that perspective transformations have "pseudo-" or "quasi-invariants": properties that may vary, but do so only within a small range of values over a large range of transformations. Hence, they are not perfectly invariant, but they might be useful for object recognition as demonstrated experimentally (Pizlo, 1994; Pizlo & Salach-Golyska, 1995).

In sum, according to one view, perspective transformations may be special because the visual system may have evolved to deal primarily with those transformations that are involved in viewing objects from different viewpoints, not with more general projective transformations. A good reason might be that perspective shape equivalence can always be related to a single planar shape undergo-

ing rigid motion in space, whereas this is not the case with projective shape equivalence. According to another view, perspective transformations are not special; they are only a subset of the group of projective transformations, and the invariants that they have in common (i.e., the projective invariants) form the basis of perceived shape equivalence. In the former view, the optics of the vertebrate eye is taken more seriously than are the mathematical principles of group theory defining the hierarchy of geometries. In the latter view, the opposite is true.

In an experiment designed to tackle these issues, participants saw three patterns presented simultaneously on a computer screen, one on top and two below it. The one on top was to be taken as the standard pattern viewed orthogonally, and observers had to determine which of the two patterns below best matched the standard pattern. We used three matching tasks. In the first, called *perspective versus random*, one of the two choice patterns was a perspective transformed version of the standard pattern and the other was a perspective transformed version of another random pattern. In the second, called *projective versus random*, one of the two choice patterns was related to the standard pattern through a projective transformation and the other was unrelated (random). In the third matching task, called *perspective versus projective*, the two choice patterns were a perspective transformed and a projectively transformed version of the standard pattern.

The predictions are straightforward (see Table 1). When perspectivities have a special status for the visual system, the perspective equivalent pattern in the first task should be easy to select (i.e., close to 100% correct), but the projectively equivalent pattern in the second task should be difficult to select (i.e., close to chance level, which is 50% correct). When both of these transformations are compared directly in the third task, participants should select the perspective one in an overwhelmingly large portion of the trials (i.e., close to 100%). In contrast, when perspective transformations are not categorically distinct from other projective transformations, projectively transformed versions of the standard patterns in the second task should be selected as often as the perspective transformed ones in the first task (i.e., both close to 100%), and there should be no difference in the third task (i.e., both close to 50%).

Table 1
Predictions of Possible Outcomes in Three Comparisons Involving Perspective and Projective Shape Equivalence and the Results of an Experiment Testing Them

Comparison Task	Predictions Based on:		Results
	Optic Principles	Group-Theoretic Principles	
Perspective versus random	100	100	92.6
Projective versus random	50	100	88.8
Perspective versus projective	100	50	58.5

Note—All values are percentages. Predictions based on optic principles imply that perspectivities are special. Predictions based on group-theoretic principles imply that perspectivities are just projectivities.

METHOD

Participants

Sixteen undergraduates at the University of Leuven volunteered to participate for partial course credit. The participants were naive. They had normal or corrected-to-normal vision.

Apparatus

We used 16 personal computers with Intel 80486 processors and identical SVGA boards. The patterns were presented on computer screens with 56-Hz temporal resolution and 800×600 spatial resolution, viewed from an average distance of 50 cm.

Stimuli

The standard patterns were irregular pentagons, with the positions of the vertices chosen pseudorandomly (for more details, see Wagemans et al., 1994); they could be convex or concave. The pentagons were dark gray on a light gray background, except for one of the five sides (randomly selected), which was colored yellow to facilitate the correspondence finding. The original patterns were made to fit in an imaginary circle with a radius of 4.5 cm (5° visual angle), positioned in the center of the top half of the screen (i.e., the circle's midpoint was at $x = 400$, $y = 150$). The two patterns below the standard were, depending on the discrimination task, perspective or projective transformations of the standard or of a new pseudorandomly generated pattern. These two patterns were also included in two imaginary circles of the same size presented in the bottom half of the screen, one centered in the left half (at 200, 450) and one in the right half (at 600, 450). The three-pattern configuration remained on the screen until the participants had indicated their choice.

We will first describe the manner in which the perspective transformation was implemented. First, the standard pattern, with all z -coordinates arbitrarily set to zero (i.e., parallel to the projection plane), was rotated (ϕ), slanted (σ), and tilted (τ), in that order:

$$\begin{bmatrix} xt \\ yt \\ zt \end{bmatrix} = \begin{bmatrix} \cos \tau & -\sin \tau & 0 \\ \sin \tau & \cos \tau & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \sigma & -\sin \sigma \\ 0 & \sin \sigma & \cos \sigma \end{bmatrix} \cdot \begin{bmatrix} \cos \phi & -\sin \phi & 0 \\ \sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

This transformed pattern was then projected onto the image plane ($z = 0$), with the center of projection at $(0, 0, -300)$. This projection results in a rather large perspective distortion, whereas the size of the image is of the same order as the original pattern:

$$x_{\text{persp}} = f \cdot \frac{xt}{(zt + f)}$$

$$y_{\text{persp}} = f \cdot \frac{yt}{(zt + f)}$$

The overall transformation yields a rather general perspective transformation from the original (x, y) to the new $(x_{\text{persp}}, y_{\text{persp}})$ coordinates. Projective transformations were implemented in the same way, except for an additional (area-preserving) shear operation that was performed on the projected points:

$$x_{\text{proj}} = x_{\text{persp}} \cdot \cosh(sh) + y_{\text{persp}} \cdot \sinh(sh)$$

$$y_{\text{proj}} = x_{\text{persp}} \cdot \sinh(sh) + y_{\text{persp}} \cdot \cosh(sh).$$

Mathematically, every projectivity can be achieved by the concatenation of a perspectivity and an affine transformation. We preferred this implementation with shear (i.e., $a = \cosh$, $b = \sinh$) because the straightforward parameterization provided a good metric for the difference between perspective and projective transformation to compare performance against while keeping area constant (to avoid area being a cue). Examples of stimuli are shown in Figure 1.

Design

Three different matching tasks were used. To avoid a combinatorial explosion of trial numbers only a maximum of three parameters were varied in a systematic way. For each trial, rotation (ϕ) and tilt (τ) were chosen randomly out of nine possible angles (0° , 22.5° , 45° , 67.5° , 90° , 112.5° , 135° , 157.5° , and 180° ; sign was random).

In the first task, the two choice patterns were a perspective transformation of the standard pattern and a perspective transformation (with the same parameter values) of another random pattern. We presented 50 trials for each of four levels of slant angle ($\sigma = 30^\circ$, 45° , 60° , or 75° ; sign was random), resulting in 200 trials. In the second task, the two choice patterns were a projective transformation of the standard pattern and a projective transformation (with the same parameter values) of another random pattern. Here, we presented 10 trials for each combination of four slant angles (same as in Task 1) and five shear values ($sh = .1, .2, .3, .4, .5$; sign was random), also resulting in 200 trials. In the third task, both a perspective transformation and a projective transformation of the standard pattern were shown. Here, we used 10 trials for each combination of three perspective slant angles (30° , 52.5° , or 75° ; sign was random), three projective slants (idem), and five shear values (as in Task 2), yielding 450 trials in total.

Task and Procedure

The 16 participants performed the experiment simultaneously in a classroom equipped with identical personal computers. All participants worked individually and at their own pace on a personal computer that was separated from a neighboring one by at least 75 cm.

We asked the participants to indicate which of the two patterns presented on the lower half of the screen was most likely to be the same planar object, although viewed from a different angle, as the standard pattern presented in the top half of the screen. We told them explicitly that the top pattern was a planar object oriented perpendicular to the line of sight. The participants gave their responses by pressing one of two arrow keys on the keyboard (e.g., left arrow for pattern at the left). The three tasks were run in a random sequence. After each task, and in the middle of the perspective versus projective task, a small break was provided. The participants were not told that the experiment consisted of three tasks.

RESULTS

To avoid potential problems with ceiling effects on the raw choice frequencies, we have also analyzed transformed data as recommended by Tukey (1977, pp. 498–508). More specifically, for each choice frequency, we have calculated the *folded logarithms* or *flogs*: $.5 \ln [(\text{choice frequency for alternative 1} + 1/6) / (\text{choice frequency for alternative 2} + 1/6)]$. All of the results of the analysis of variance (ANOVA) were the same, except one (which will be noted below). For the reader's convenience, we will report the untransformed choice frequencies.

Perspective Versus Random

In 92.6% of the trials, the perspectively transformed pattern was chosen as being most closely equivalent to the standard pattern. For increasing slant angles, the percent correct responses decreased slightly: 94.6%, 94.6%, 92.0%, and 89.1%. An ANOVA showed that this effect was statistically significant [$F(3,45) = 5.40$, $MS_e = 0.102$, $p < .003$].

Projective Versus Random

In 88.8% of the trials, the projectively transformed pattern was chosen as being most closely equivalent to the standard pattern. A t test for dependent samples showed that this performance differed reliably from that in the previous task [$t(15) = 2.78$, $p < .015$]. Larger slant angles reduced the number of correct answers, as did higher shear values. Percent projective responses were 91.4%, 90.6%, 88.9%, and 84.2% for increasing

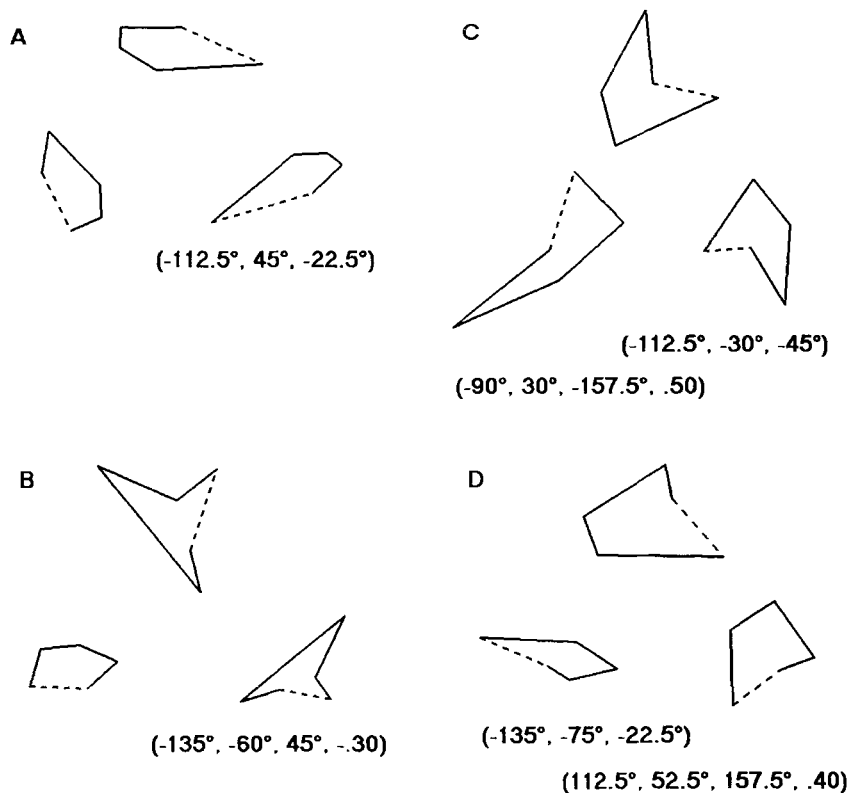


Figure 1. Examples of stimuli used in the experiment. The numbers underneath the test patterns indicate the parameter values when the test pattern is derived from the standard pattern. For the perspective transformations, the parameters indicate rotation, slant, and tilt, respectively; for the projective transformations, the parameters indicate rotation, slant, tilt, and shear, respectively. The dashed line represents the yellow line in the patterns used in the experiment. (A) Perspective versus random (note that this trial would not be easy because all three pentagons are convex). (B) Projective versus random (note that this trial would be easy because the distractor is completely convex, whereas each matching pattern has two concavities). (C) Perspective versus projective (note that shear probably determines the [perspective] preference here because perspective and projective slant happen to be equal). (D) Perspective versus projective (note that [projective] preference might be based on the smallest slant in this case).

slants and 90.5%, 91.2%, 89.0%, 87.5%, and 85.8% for increasing shear values. Both effects proved to be reliable for [slant, $F(3,45) = 11.68$, $MS_e = 0.070$, $p < .0001$; for shear, $F(4,60) = 4.35$, $MS_e = 0.072$, $p < .004$]. There was no significant interaction between these two factors ($p > .12$).

Perspective Versus Projective

In 58.5% of the trials, the participants selected the perspective transformed version of the standard pattern. A t test indicated that this was significantly better than chance [$t(15) = 10.77$, $p < .0001$]. A closer look at the data shows that this preference is modulated by the specific values of the transformation parameters. The preference for the perspective transformation decreased with increasing perspective slants [$F(2,30) = 161.87$, $MS_e = 0.535$, $p < .0001$] and decreasing projective slants [$F(2,30) = 43.26$, $MS_e = 0.719$, $p < .0001$]. In other words, there was a strong tendency to select the matching pattern with

the smallest slant, regardless of whether it was a perspective or a projectively equivalent version. Higher values of shear also resulted in an increased perspective preference [$F(4,60) = 30.17$, $MS_e = 0.271$, $p < .0001$].

Two interaction effects were also reliable. The interaction between perspective slant and projective slant¹ [$F(4,60) = 3.86$, $MS_e = 0.211$, $p < .0075$] means that the effect of projective slant became larger for larger perspective slant angles (see Figure 2A). For example, when the perspective slant was 75°, the percent perspective preference dropped to around 30% when the projective slant was smaller (i.e., 30° or 52.5°). It is also interesting to observe what happened for the three conditions in which both slant angles were equal: When both were 30°, there was still a pronounced perspective preference (67%); this dropped to 60% when both were 52.5° and to almost chance level (52%) when both slants were 75°. In fact, this means that the contribution of the second parameter, which can differ between the perspective and the

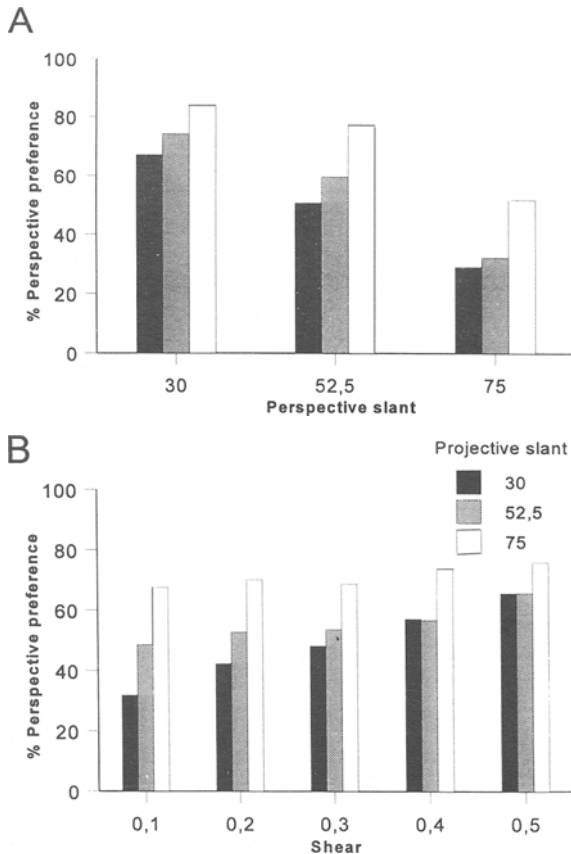


Figure 2. Parametric effects obtained in the third comparison task (perspective vs. projective). Percent perspective preference as a function of (A) perspective and projective slants and (B) projective slant and shear.

projective transformation (i.e., shear), had a smaller effect for the larger slants. This is corroborated by the reliable interaction between projective slant and shear [$F(8,120) = 5.88$, $MS_e = 0.219$, $p < .0001$]. Figure 2B demonstrates that the effect of shear became smaller when projective slant was larger.

DISCUSSION

Two major conclusions follow from these results. First, perceived shape equivalence under affine transformations as established earlier (Wagemans et al., 1994) generalizes to perspective and projective transformations. In fact, the average performance level in our first two tasks (around 90%) is even higher than what we found in our previous study with affine shape equivalence (around 75%) in which we always presented two patterns side by side on the screen (see Wagemans et al., 1994, for more details). The three-pattern configuration used here is easier because participants do not have to decide on a trial-by-trial basis whether the shapes are sufficiently equivalent to be given a *same* response; they can always look for the best possible match, which reduces the likelihood of false positives. Moreover, some trials (e.g., Figure 1B) were easy because they contained distractors that were qualitatively different from the other two patterns (e.g., with a concavity while the standard and match were both completely convex, or vice versa). With pentagons, this occurs more often than with the quadrilaterals that we have used before. In fact, post hoc computations showed that the size of the convex hull (i.e., the number of sides of the smallest possible convex polygon enclosing all vertices of the pen-

tagons) could have been a useful cue in 53% of the trials in the first two tasks. So, additional information must have been used to explain the obtained performance levels in these two tasks as well as the perspective preference in the third task. In any event, it will be clear that our task was easier than Niall and Macnamara's (1990) task in which participants had to rely on purely quantitative differences in cross ratios.

Second, comparing perspective with projective transformations, we must conclude that the results are much closer to the prediction (based on group-theoretic principles) that the visual system does not deal with perspectivities in a categorically different way from more general projectivities than to the prediction (derived from optic principles) that perspectivities have a special perceptual status (see Table 1). Nevertheless, we found that perspective-related shapes were easier to distinguish from unrelated shapes (in Task 1) than were projectively related shapes (in Task 2). Moreover, in a direct comparison task (perspective versus projective in Task 3), we found a significant perspective preference. Because the results were so strongly determined by specific parameter values, we are tempted to attribute these differences not to an inherent and fundamental (qualitative) difference between perspective and more general projective transformations but to specific (quantitative) parametric differences between them.

For example, it is clear that the degree of slant strongly affected the perceived shape equivalence: The test pattern with the smallest slant tended to be chosen. With equal slant angles, the extra shear parameter also affected the choice behavior, but to a much lesser extent with large slants. We would like to point out that, with negligible slants, the difference between a perspective transformation and a projective transformation, as we have implemented them, reduces to the difference between a Euclidean (rotation) transformation and an affine (shear) transformation, respectively. According to the group-theoretic principles, one would expect a larger effect of shear with smaller slants because the transformations then belong to different groups. The results supported that expectation.

These significant parametric influences are interesting because they temper the pretended explanatory power of any absolute approach to perception, such as optical versus geometrical principles in general or, more specifically, sensitivity to *only* perspectivities versus *all* projectivities alike. Clearly, with small enough differences between perspective and projective transformations, the difference would be indistinguishable for the perceptual system. For a mathematical framework such as invariants-under-transformations to be useful as a basis for a perceptual theory, at least this notion of a just noticeable difference would have to be incorporated (see Cutting, 1987a, for a general discussion; see Busey, Brady, & Cutting, 1990, for a discussion of indiscriminable distortions in the context of affine shape equivalence).

However, these parametric influences may also be a cause of concern. First, perhaps there is a simple Euclidean measure that covaries with perspective slant, projective slant, and projective shear and that completely determines performance in our comparison tasks. Indeed, Niall and Macnamara (1990, p. 657) have suggested, on the basis of their experimental results, that "observers are responsive to the Euclidean properties of perspective views rather than the invariant projective properties of those views." One such measure is the perimeter squared over the area, P^2/A , which is called *compactness* (see Zusne, 1970, pp. 206–223). Given a constant area, the shape with the longest perimeter will appear to have the least compact shape. Thus, with larger slants and shears, shape becomes less compact. Considering the sensitivity of human perceivers to this measure (see Cutting & Garvin, 1987, for a brief review), it might well be the case that our observers always selected the comparison shape that best resembled the standard shape in terms of compactness.

We have run a simulation study to test this possibility.² We have written an algorithm to calculate P^2/A for pentagons and a simple shape-matching algorithm that always selects a comparison shape that minimizes the compactness difference with a standard shape. We have then presented this computer program with the same number of stimuli and the same procedure to generate stimuli as those for our participants in order to create an artificial data set ($N = 49$) that could be subjected to the same ANOVAs. As expected, the same trends could be observed here. For example, the preference for a perspective transformed shape over a random one decreased with increasing slant (89.2%, 76.6%, 60.9%, and 51.1% for 30°, 45°, 60°, and 75° slant an-

gles, respectively), as did the preference for a projectively transformed shape over a random one with increasing slant (65.7%, 62.6%, 55.9%, and 50.7%, respectively) and shear (64.4%, 62.7%, 59.1%, 54.7%, and 52.7% for shear levels going from .1 to .5, respectively). However, when these simulated values are compared with the choice frequencies of our real observers, it is clear that compactness cannot explain everything. For Task 1, for example, the simulation yields a perspective preference of 89% at the smallest slant angle, but this drops off rapidly with larger slants, whereas the perspective preference obtained in the experiment remains high throughout (e.g., also 89% at the largest slant, compared with 51% for the simulation). Likewise, for Task 2, at the smallest levels of slant and shear, the simple preference for shapes with similar compactness would yield only a 65% preference for the truly equivalent shapes, whereas almost-chance performance would be obtained at the largest levels of slant and shear. In contrast, human perceivers have a very pronounced preference for truly equivalent shapes at all parameter levels (ranging from 90% to 85% in the most difficult conditions). Similar observations apply to the interactions and the data from the third task. In sum, although compactness clearly varies with our parameter values in the predicted way, our observers' choice behavior cannot be determined completely by a simple preference for shapes with similar compactness.

A second concern is that some people might reason that these parametric effects argue against the use of invariants: If invariants underly shape constancy, then performance in determining shape equivalence under a set of (perspective or projective) transformations should not vary with the size of the transformation. In contrast, we have argued, on both theoretical and empirical grounds, that the visual system's use of invariants need not itself be invariant (Wagemans, Van Gool, & Lamote, 1996). As was the case for affine invariants in the earlier study, the parametric effects of slant and shear in the present study might reflect the difficulty of the measurement of the invariants instead of a normalization operation as generally assumed (e.g., Tarr, 1995; Ullman, 1989).

REFERENCES

- BUSEY, T. A., BRADY, N. P., & CUTTING, J. E. (1990). Compensation is unnecessary for the perception of faces in slanted pictures. *Perception & Psychophysics*, **48**, 1-11.
- CARLBOM, I., & PACIOREK, J. (1978). Planar geometric projections and viewing transformations. *ACM Computing Surveys*, **10**, 465-502.
- CUTTING, J. E. (1986). *Perception with an eye for motion*. Cambridge, MA: MIT Press, Bradford Books.
- CUTTING, J. E. (1987a). Perception and information. *Annual Review of Psychology*, **38**, 61-90.
- CUTTING, J. E. (1987b). Rigidity in cinema seen from the front row, side aisle. *Journal of Experimental Psychology: Human Perception & Performance*, **13**, 323-334.
- CUTTING, J. E. (1988). Affine distortions of pictorial space: Some predictions for Goldstein (1987) that La Gournerie (1859) might have made. *Journal of Experimental Psychology: Human Perception & Performance*, **14**, 305-311.
- CUTTING, J. E., & GARVIN, J. J. (1987). Fractal curves and complexity. *Perception & Psychophysics*, **42**, 365-370.
- GIBSON, J. J. (1950). *The perception of the visual world*. Boston, MA: Houghton Mifflin.
- GOLDSTEIN, E. B. (1987). Spatial layout, orientation relative to the observer, and perceived projection in pictures viewed at an angle. *Journal of Experimental Psychology: Human Perception & Performance*, **13**, 256-266.
- HELMHOLTZ, H. (1968). The recent progress of the theory of vision. In R. M. Warren & R. P. Warren (Eds.), *Helmholtz on perception* (pp. 61-136). New York: Wiley. (Original work published 1868)
- ITTELSON, W. H. (1996). Visual perception of markings. *Psychonomic Bulletin & Review*, **3**, 171-187.
- KLEIN, F. (1893). Vergleichende Betrachtungen über neuere geometrische Forschungen [Comparative reflections on new geometric investigations]. *Mathematische Annalen*, **43**, 63-100. (Original work published 1872)
- KUBOVY, M. (1986). *The psychology of perspective and Renaissance art*. Cambridge: Cambridge University Press.
- KUKKONEN, H. T., FOSTER, D. H., WOOD, J. R., WAGEMANS, J., & VAN GOOL, L. (1996). Qualitative cues in the discrimination of affine-transformed minimal patterns. *Perception*, **25**, 195-206.
- MARK, L. S., & TODD, J. T. (1985). Describing perceptual information about human growth in terms of geometric invariants. *Perception & Psychophysics*, **37**, 249-256.
- MARK, L. S., TODD, J. T., & SHAW, R. E. (1981). Perception of growth: A geometric analysis of how different styles of change are distinguished. *Journal of Experimental Psychology: Human Perception & Performance*, **7**, 855-868.
- MICHAELS, C. F., & CARELLO, C. (1981). *Direct perception*. Englewood Cliffs, NJ: Prentice-Hall.
- NIALL, K. K. (1992). Projective invariance and the kinetic depth effect. *Acta Psychologica*, **81**, 127-168.
- NIALL, K. K., & MACNAMARA, J. (1990). Projective invariance and picture perception. *Perception*, **19**, 637-660.
- PITTINGER, J. B., & SHAW, R. E. (1975). Aging faces as viscal-elastic events: Implications for a theory of nonrigid shape perception. *Journal of Experimental Psychology: Human Perception & Performance*, **1**, 374-382.
- PITTINGER, J. B., SHAW, R. E., & MARK, L. S. (1979). Perceptual information for the age-level of faces as a higher-order invariant of growth. *Journal of Experimental Psychology: Human Perception & Performance*, **5**, 478-493.
- PIZLO, Z. (1994). A theory of shape constancy based on perspective invariants. *Vision Research*, **34**, 1637-1658.
- PIZLO, Z., & ROSENFELD, A. (1992). Recognition of planar shapes from perspective images using contour-based invariants. *Computer Vision, Graphics & Image Understanding*, **56**, 330-350.
- PIZLO, Z., & SALACH-GOLYSKA, M. (1995). 3-D shape perception. *Perception & Psychophysics*, **57**, 692-714.
- ROCK, I. (1973). *Orientation and form*. New York: Academic Press.
- ROCK, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
- TARR, M. J. (1995). Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bulletin & Review*, **2**, 55-82.
- TUKEY, J. W. (1977). *Exploratory data analysis*. Reading, MA: Addison-Wesley.
- ULLMAN, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, **32**, 193-254.
- VAN GOOL, L., MOONS, T., PAUWELS, E., & WAGEMANS, J. (1994). Invariance from the Euclidean geometer's perspective. *Perception*, **23**, 547-561.
- WAGEMANS, J., DE TROY, A., VAN GOOL, L., FOSTER, D. H., & WOOD, J. R. (1994). *Minimal information to determine affine shape equivalence*. Unpublished manuscript (Intern. Rep. No. 169). University of Leuven, Laboratory of Experimental Psychology.
- WAGEMANS, J., VAN GOOL, L., & LAMOTE, C. (1996). The visual system's measurement of invariants need not itself be invariant. *Psychological Science*, **7**, 232-236.
- ZUSNE, L. (1970). *Visual perception of form*. New York: Academic Press.

NOTES

1. This interaction was no longer significant when the flog-transformed data were analyzed [$F(4,60) < 1$]. This means that the specification of the effect and its subsequent discussion should be treated with caution.

2. It would have been slightly better to calculate P^2A over all shapes presented in the original experiment and regress the preference data against the difference values of both comparison shapes with the standard shape. However, our programs for stimulus presentation and response registration did not write the randomly generated coordinates of the pentagons into the output file used for record keeping. Instead of repeating the experiments with a modified program that would allow these post hoc statistical tests, we decided that a simulation study would provide the same insights more economically.

(Manuscript received October 2, 1995;

revision accepted for publication August 7, 1996.)