

Categorical perception of relative orientation in visual object recognition

LUKE J. ROSIELLE and ERIC E. COOPER
Iowa State University, Ames, Iowa

The purpose of the present investigation was to determine whether the orientation between an object's parts is coded categorically for object recognition and physical discrimination. In three experiments, line drawings of novel objects in which the relative orientation of object parts varied by steps of 30° were used. Participants performed either an object recognition task, in which they had to determine whether two objects were composed of the same set of parts, or a physical discrimination task, in which they had to determine whether two objects were physically identical. For object recognition, participants found it more difficult to compare the 0° and 30° versions and the 90° and 60° versions of an object than to compare the 30° and 60° versions, but only at an extended interstimulus interval (ISI). Categorical coding was also found in the physical discrimination task. These results suggest that relative orientation is coded categorically for both object recognition and physical discrimination, although metric information appears to be coded as well, especially at brief ISIs.

Virtually all sensory input varies along a continuum, yet people tend to perceive the world in terms of discrete categories. For example, although the visible spectrum of light varies continuously from a wavelength of approximately 370 to 730 nanometers, when viewed by an observer, this continuum appears to be composed of relatively discrete bands of color. In fact, the tendency to categorize the sensory world is so strong that a person's ability to physically discriminate among stimuli that are members of the same category (e.g., two shades of red) may suffer considerably. This phenomenon is known as categorical perception.

Harnad (1987) has defined categorical perception as an "analog to digital transformation that recodes a continuous region of physical variation as a discrete, labeled equivalence class" (p. 4). Importantly, Harnad identified two general criteria that must be met in order to claim that a given perceptual event is an instance of categorical perception:

1. First, two stimuli that are members of the same category should be harder to physically discriminate than two stimuli that are members of a different category, even though the physical distances between the two sets of stimuli are equal. Research has shown that two wavelengths of light both labeled "blue" by the participant are more difficult to physically discriminate than two wavelengths of light, one of which is labeled "blue" and the other la-

beled "green," even though the physical distances between the two "blue" wavelengths and the "blue" and "green" wavelengths are equal (Bornstein & Korda, 1984).

2. Second, stimuli that vary along a physical continuum should be classified into several discrete categories. For example, it is a well-established phenomenon that people perceive the visible spectrum of light as a few discrete bands of hues, each of which can be uniquely described as being qualitatively different from the others (Bornstein, 1987). A wide variety of experimental paradigms have been used to assess classification, including overt labeling (e.g., "What is the name of this stimulus?"), *same-different* identity judgments (e.g., "Do these two stimuli have the same name?"), and forced-choice (e.g., "Is this stimulus a member of Category X or Category Y?"; see, e.g., Cooper & Biederman, 1993; Etcoff & Magee, 1992; Liberman, Harris, Hoffman, & Griffith, 1957; Rosch, 1975).

What would be the computational advantages of parsing continuous sensory information into discrete categories? Numerous authors have suggested that categorical perception helps to significantly reduce the amount of information a sensory system has to process (Beale & Keil, 1995; Smith & Medin, 1981; Snowdon, 1987) by allowing the brain to act upon a relatively small number of categories instead of a much larger stimulus continuum. Although some degree of precision is sacrificed in the process, categorical perception would thus allow for quicker, easier processing of sensory stimuli. In addition, categorical perception allows a much greater toleration of perceptual noise because a range of input values will all activate the same category in memory. Thus, despite some variation in the input, the correct representation in memory is more likely to be activated (see Harnad, 1987).

Categorical perception has been found to occur in a number of sensory domains and for a variety of different

Portions of this research were presented at the 1998 Meeting of the Association for Research in Vision and Ophthalmology, Ft. Lauderdale, Florida. The authors thank Charles Leek, Veronica J. Dark, William S. Robinson, and Lori D. Lindley for their careful readings of an earlier version of this manuscript. We also thank Brandi King, Kris Ghosh, Tammy Sass, Eric Esser, and Shala Wyant for their assistance in conducting the research described in this paper. Correspondence should be addressed to L. J. Rosielle, W112 Lagomarcino Hall, Ames, IA 50011 (e-mail: zaphrod@iastate.edu).

stimuli. Initially, the phenomenon was thought to be limited to lower level stimuli such as phonemes and colors (Bornstein, 1987; Bornstein & Korda, 1984; Liberman et al., 1957). However, investigators have since found evidence that categorical perception also governs some higher level aspects of perception, most notably in the domain of face processing. For example, Etcoff and Magee (1992) have found categorical perception effects in the perception of facial expressions (see also Calder, Young, Perrett, Etcoff, & Rowland, 1996), whereas Beale and Keil (1995) have found categorical perception effects in the perception of facial identity.

If categorical perception affects aspects of face perception, then it seems plausible that categorical perception may also be involved in aspects of object (nonface) perception. However, virtually no direct evidence speaks to this proposition, and the relationship between categorical perception and object recognition has yet to be empirically explored. Thus, one of the goals of the present research is to seek evidence that speaks to whether categorical perception is involved in the perception of objects.

A good place to start might be to look for evidence of categorical perception in the coding of the spatial relations among an object's parts. In particular, a useful candidate for a type of spatial relation that might be perceived categorically is the relative orientation of an object's parts. The orientation of simple shapes, for example, is known to be discriminated categorically (see, e.g., Appelle, 1972; Essock, 1980), and this finding may extend to the perception of the orientation between an object's parts as well. Unfortunately, little empirical evidence speaks to the role relative orientation plays in recognition, let alone whether relative orientation constitutes an instance of categorical perception. Thus a second goal of the present research is to investigate how relative orientation is coded in the representation used for object recognition.

On computational grounds, some theories of object recognition have posited that relative orientation might be coded during object recognition. Specifically, some authors (e.g., Biederman, 1987; Hummel & Biederman, 1992) have suggested that relative orientation is coded categorically via the categories parallel, perpendicular, and oblique for the purpose of basic-level object recognition. For example, in these theories that posit categorical coding, the handle of a coffee cup might simply be coded as being "parallel to" the body of the cup, as opposed to its exact orientation being specified. Given that objects regularly vary in the orientation of their parts, it may be computationally useful to categorically code the orientation of an object's parts in order to make the representation more robust to changes in viewpoint. Like other types of categorical relations, coding relative orientation (as opposed to absolute orientation) allows the memory representation used to recognize that object to become invariant to many of the variations that characterize naturalistic viewing conditions, such as changes in

the object's size, position in the visual field, or orientation. For example, if a cup were to change its size, position, or orientation, the handle of the cup would still be "parallel to" the body of the cup.

It is important to note, however, that categorical coding of orientation, as defined by models of shape recognition that posit categorical coding (e.g., Hummel & Stankiewicz, 1998), only implies that an output function is nonlinear across categorical boundaries. Thus, according to these models, noncategorical (metric) information is not necessarily excluded from a categorical representation. Categorical coding, therefore, only makes the strong prediction that recognition time to appreciate the similarity of two stimuli that have a metric difference in the orientation of their parts should be nonlinear across categorical boundaries.

Why do theories of object recognition that posit categorical coding of orientation use the categories parallel, perpendicular, and oblique when, conceivably, any arbitrary range of orientations could serve as a category? For example, a possible categorical relation could be "between 0° and 45°," and a second possible categorical relation could be "between 45° and 90°." Whereas any arbitrary range of orientations that served as a category would not be disrupted by changes in the size or the position of the object, they might be disrupted by rotations in depth. As an object rotates in depth, the degree of obliqueness between the object's parts changes in the two-dimensional image projected by the object onto the retina (Figure 1, top row). It would not be useful to form a categorical relation that separates oblique angles into more than one category because there is a risk that the orientation between an object's parts, when seen rotated in depth, will cross a categorical boundary (thus activating a different memory representation and impairing the recognition process). However, parallel and perpendicular do not change as an object rotates in depth (Figure 1, bottom two rows).

Coding parallel, perpendicular, and oblique for the purpose of recognition would be advantageous, however, because these categories are less easily disrupted by rotations in depth than are other categories of orientation. For example, as a viewer's perspective on all three objects depicted in Figure 1 changes (the 0° version of each object compared with the 30° version of each object), the categorical relations between the two parts of the object will not change in that parallel will always appear parallel, perpendicular will always appear perpendicular, and oblique will always appear oblique (although the degree of obliqueness will change).

All of the arguments hitherto presented for including parallel, perpendicular, and oblique as categorical relations in an object recognition system are entirely computational. Although great success has been enjoyed by computer models of recognition that explicitly code parallel, perpendicular, and oblique (e.g., Hummel & Bie-

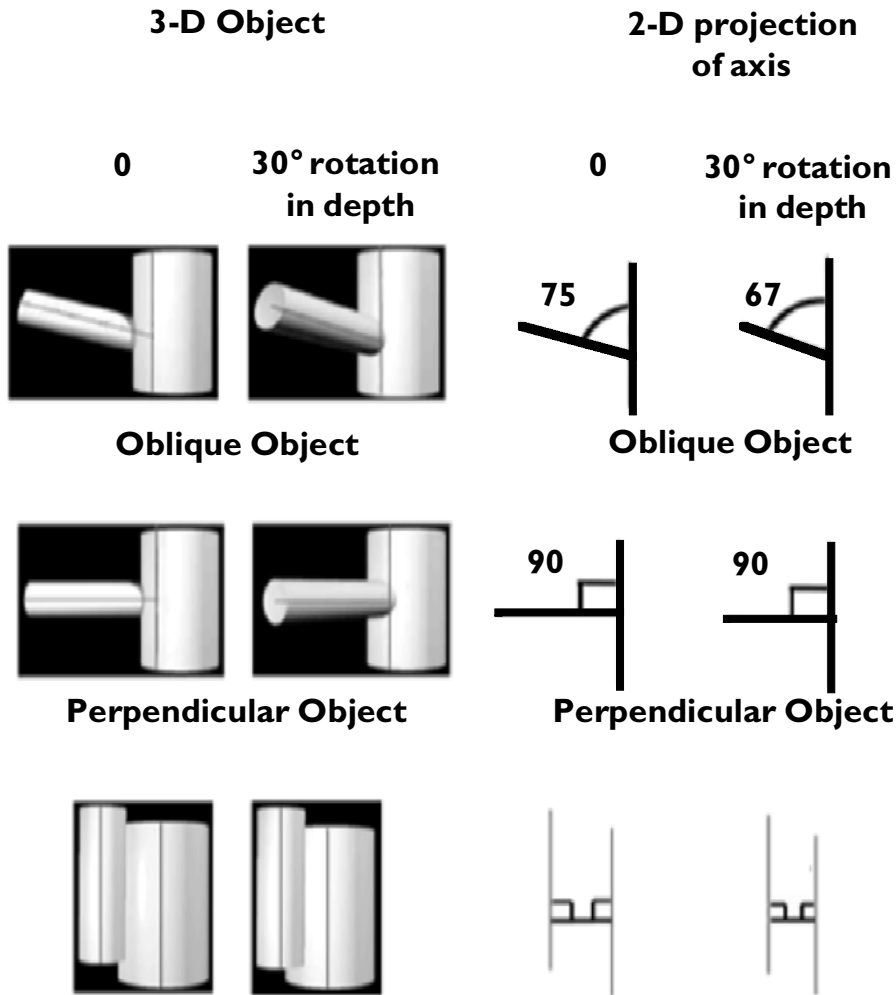


Figure 1. Illustration of how the degree of obliqueness between the parts of an object changes as an object rotates in depth. Three objects with different degrees of orientation between their parts (designated the *oblique object*, the *perpendicular object*, and the *parallel object*) have been rotated 30° in depth toward the viewer. The major axes of the parts of the objects have been superimposed onto each object and have been presented in isolation to the right. Note how as the objects are rotated in depth, the categorical relations between the object parts do not change, although the degree of obliqueness does change (as indicated by the different oblique angles).

derman, 1992), the question of whether the human object recognition system computes these categories between the parts of an object has never been tested empirically.

Alternative Ways of Coding Orientation

Despite the advantages of coding the orientation between an object's parts categorically (via the specific categories parallel, perpendicular, and oblique), and because of the lack of empirical data that speak to the question of how orientation might be coded, several alternative, noncategorical ways of coding orientation are possible. The two principal alternative ways of coding orientation are discussed briefly in the following sections.

No coding of orientation. One possibility is that the memory representation of an object is completely invariant to the orientation of the object's parts. That is, the same representation in memory will be activated no matter how an object's parts are oriented. Although this may at first glance seem unlikely, a system that does not code orientation would be advantageous because it would allow for a greater degree of generalization among objects that differ in the orientation of their parts (as opposed to delegating objects that differ in the orientation of their parts into three separate categories). It would be computationally inefficient to represent the orientation of the parts of certain nonrigid objects, such as the human body,

or objects attached to some kind of hinge (e.g., a tank turret; Kosslyn, 1994). Thus, it is possible that the object recognition system does not code relative orientation at all.

Metric coding of orientation. A second possibility is that orientation is only coded metrically. That is, the orientation of an object's parts could be coded along a continuum and not partitioned into separate categories. While coding orientation metrically loses many of the advantages mentioned above of parsing continuous information into discrete categories, the precise nature of the metrically defined memory representation could allow for easier discriminations between two objects that differ only in a small degree of orientation change in one of their parts. Whether such a precise representation is generally useful for the purposes of everyday vision is an open question and a subject of some debate in the literature (see Biederman & Gerhardstein, 1995; Tarr & Bülthoff, 1995). As noted previously, categorical coding and metric coding are not necessarily mutually exclusive, and it is certainly possible that both play a role in object perception (for a discussion, see Hummel & Stankiewicz, 1998).

Three experiments were conducted to investigate the nature of the coding of orientation between the parts of an object. The logic behind the experiments derives from the methods used to look for categorical perception. Stimuli were novel objects that varied along a continuum: the orientation between the objects' parts. In Experiment 1, participants were given a physical discrimination task in which they had to determine whether two sequentially presented objects were physically identical. In Experiments 2 and 3, participants were given a classification task in which they had to determine whether two objects were the same in terms of the identity of their parts, while ignoring any changes in orientation among the parts that might occur. The same set of stimuli were used in all three experiments.

Stimuli

The stimuli used in all three experiments were line drawings of simple three-part objects (Figure 2). As can be seen

in Figure 2, each object is composed of two large parts, each corresponding to one of Biederman's (1987) geons, connected by a smaller elongated "arm" geon to allow for easy rotations. The arm geon was attached to one of the larger geons (designated the "base" geon) via a small circle that allowed the arm to rotate without producing emergent object features (i.e., features, such as extra vertices, that might be produced by changing the physical characteristics of the stimuli). Different objects were produced by changing the identity of the base geon, the other large geon (hereafter referred to as the "top" geon), or both. The arm geon always remained the same.

A total of 12 different objects were constructed by combining four different base geons and three different top geons. The four different base geons were a brick, a cone, a cylinder, and a wedge. The three different top geons were a cone, a cylinder, and a wedge. For both experiments, the arm (and the top geon) were rotated to produce four different versions of each object: 0°, 30°, 60°, and 90°, all relative to the major axis of the base geon. An example of the four versions of one object is shown in Figure 2. Note that the physical difference between successive versions of the object are identical (i.e., steps of 30°).

EXPERIMENT 1 Physical Discrimination

Harnad's (1987) first criterion for categorical perception—within-category physical discriminations are more difficult than between-category physical discriminations—was tested in Experiment 1. In Experiment 1, the participant had to decide whether two sequentially presented objects were physically identical. Thus, for any change, either in the identity of the parts or in the orientation of one of the parts, the correct response would be *different*. See Figure 3 for examples of same and different objects used in Experiment 1.

Experiment 1 had four conditions. The trials of interest in Experiment 1 are the *different* object trials. In the following description of the conditions, the two objects





Figure 2. Examples of each of the four versions of an object (0° version, 30° version, 60° version, and 90° version). Note that the 30° and 60° versions of the object occupy the same category (oblique).

These objects are all the same



These objects are all different




Figure 3. Examples of objects that are considered *same* or *different* for Experiment 1. Note that all of the *same* objects are physically identical. All of the *different* objects differ either in at least one geon or in the orientation of the object's parts.

seen by a participant during the trial sequence differed only in relative orientation, not in the geons that compose the objects. In the parallel-versus-oblique condition, the participant compared the 0° and 30° versions of an object. In the oblique-versus-oblique condition, the participant compared the 30° and 60° versions of an object. In the perpendicular-versus-oblique condition, the participant compared the 90° and 60° versions of an object. Finally, in the parallel-versus-perpendicular condition, the participant compared the 0° and 90° versions of an object.

If relative orientation is coded categorically for physical discrimination via the categories parallel, perpendicular, and oblique, the 30° and 60° objects should activate the same representation in memory because both objects are composed of the same set of parts and have the same relative orientation relationship (i.e., oblique) between their parts. However, the 0° and 30° objects and the 90° and 60° objects should activate different representations in memory, because although both sets of objects are composed of the same set of parts, each has a different relative orientation relationship between their parts (parallel and oblique, and perpendicular and oblique, respectively).

Therefore, if relative orientation is coded categorically, it should be more difficult, in terms of response times and error rates, for participants to appreciate that the 30° and 60° objects (the oblique-vs.-oblique condition) are

physically different relative to the 0° and 30° objects (the parallel-vs.-oblique condition) and the 90° and 60° objects (the perpendicular-vs.-oblique condition). This pattern of results should occur even though the absolute change in orientation is the same in all cases (i.e., 30°). Because theories that code relative orientation only predict a nonlinear cost function, these theories do not make a strong prediction regarding the parallel-versus-perpendicular condition.

If, however, relative orientation is only coded metrically, then performance should depend only on the absolute difference in orientation between the parts of the objects. This prediction assumes that the difficulty in making physical discrimination judgments scales with the similarity between the two stimuli. Thus, the three conditions in which the angular disparity between the two objects being compared is 30° (the parallel-vs.-oblique, oblique-vs.-oblique, perpendicular-vs.-oblique conditions) should produce roughly equal amounts of cost (relative to the identical condition) in terms of response times and error rates because the magnitude of the orientation change is the same in all three conditions. The parallel-versus-perpendicular condition should thus produce faster response times and more errors, because the angular disparity between the objects being compared is 90° .

If only the parts' identities are coded and no orientation information is included in the memory representation for physical discrimination, then there should be no

difference between all conditions regardless of the angular disparity between the objects' parts. Although no coding of orientation for physical discrimination seems unlikely, it is certainly possible that no coding may play a role at some level of stimulus representation (e.g., perhaps for object recognition).

Method

Participants. Forty-eight college students who reported normal or corrected-to-normal vision participated in the experiment. Participants received extra credit in an introductory psychology course or developmental psychology course at Iowa State University for taking part in the experiment.

Apparatus. Stimuli were presented on a 17-in. color Macintosh computer monitor with a resolution of 832×624 pixels and a vertical refresh rate of 75 Hz. Participants responded using a two-button response box attached to a National Instruments NB-DIO-24 interface board that gave ± 0.5 -msec response time accuracy. The experiment was controlled by Picture Perception Lab Software (Kholmeyer, 1992) running on a Quadra 800 Macintosh computer.

For the experiment, 12 objects were created by combining four different base geons and three different top geons. The base geons were a brick, a cone, a cylinder, and a wedge, and the top geons were a brick, a cone, and a cylinder. Next, four versions of each of the 12 objects were created by rotating the arm and top geon to each of the rotations used in the experiment (i.e., 0° , 30° , 60° , and 90°) for a total of 44 stimuli. Each stimulus was drawn at a stroke width of 3 pixels. The largest objects, at the distance viewed by the participant in the experiment, extended horizontally 6.2° and vertically 6.0° of visual angle.

Procedure. The participant sat 1.3 m from the computer screen during the experiment.

During each trial, participants were shown a fixation cue for 504 msec, followed by the first object for 168 msec, followed by a mask for 756 msec, followed by the second object for 168 msec, followed by a mask for 126 msec.

The participant's task was to press the left button on a response box if the two objects were physically identical in every way and to press the right button on the response box if the two objects were not physically identical. Participants were allowed to use only the fingers from one hand to press both buttons on the response box. Before the experiment began, each participant was read a standard set of directions and was shown examples of objects that were the same or different.

Participants initiated each trial sequence by pressing a mouse button in response to a prompt on the computer screen. Participants were told before the experiment began to strive for 90% accuracy. Each participant completed 20 practice trials before the experiment proper began. No data were collected on the practice trials.

Each participant viewed 288 trials, half of which were *different* trials and half of which were *same* trials. Of the 144 *different* trials, 96 trials involved an orientation change and no geon change, and 48 trials involved a geon change, and no orientation change. The 48 geon-change and no-orientation-change trials were included to encourage the participants to encode the parts of the objects; otherwise participants could accomplish the task by ignoring the parts of the objects and only looking for a change in orientation.

On the 96 orientation-change trials, the first and second objects in the trial sequence were the 0° and 30° objects one fourth of the time (24 trials), the 30° and 60° objects one fourth of the time (24 trials), the 90° and 60° objects one fourth of the time (24 trials), and the 0° and 90° objects one fourth of the time (24 trials). Additionally, each of the 12 "object classes" (i.e., a unique set of geons, such as the cone-cylinder object depicted in Figure 2) appeared twice (once each as the first and second object in the trial sequence) on eight trials. Each of the four rotations (0° , 30° , 60° , 90°) appeared equally often over the course of the orientation-change trials (half of the time as the first object in the trial sequence and half of the time as the second object in the trial sequence), and each "object class" appeared equally often over the course of the orientation-change trials.

On the 48 *different* trials in which the first and second stimuli in the trial sequence had different geons, each of the 12 object classes appeared twice (once each as the first and second object in the trial

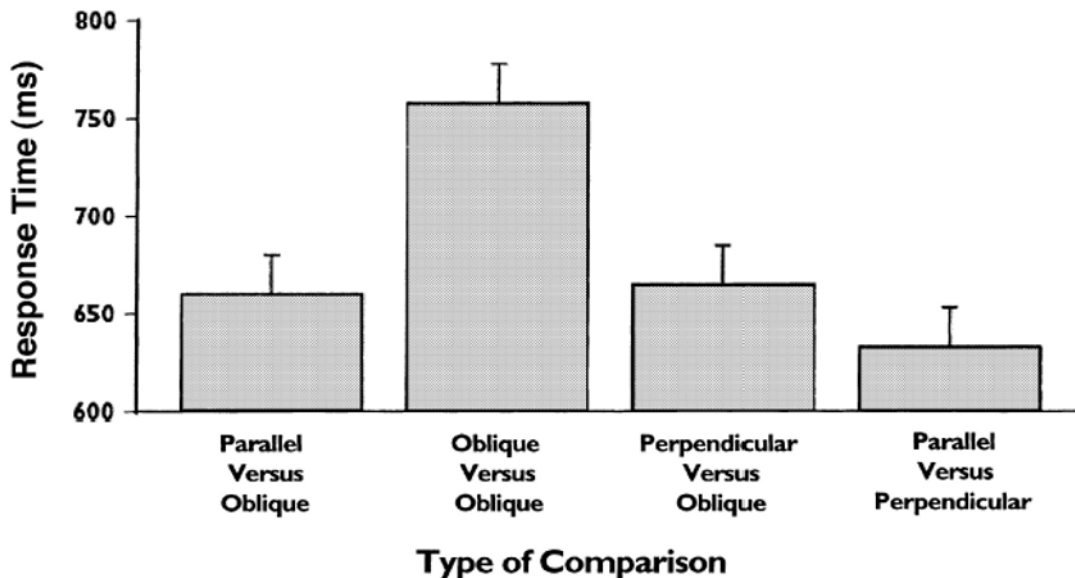


Figure 4. Mean response time (in milliseconds) for the correct responses for each condition in Experiment 1, in which there was an orientation change and no geon change. Error bars represent the standard error of the mean (Loftus & Loftus, 1988).




Figure 5. Mean error percentage for each condition in Experiment 1, in which there was an orientation change and no geon change. Error bars represent the standard error of the mean.

sequence) on four trials, and each rotation appeared twice (once each as the first and second object in the trial sequence) on one trial. On the 144 *same* trials, each of the 12 object classes appeared twice on 12 trials, and each rotation appeared twice on three trials.

Results

The results of Experiment 1 can be seen in Figures 4 and 5. An overall performance criterion of 70% correct was established for Experiment 1, and the data from participants who failed to meet this criterion were replaced by data from new participants. Because of the relatively difficult nature of the task employed in Experiment 1 (Figure 5), the data from 6 participants were replaced by data from new participants. The *same* object trials had a mean response time of 631 msec and a mean error rate of 20.7%.

All statistical hypotheses in this paper were tested with a two-tailed alpha level of .05. One-way within-subjects analyses of variance (ANOVAs) were conducted both on the response time and error rate data with type of comparison (parallel vs. oblique, oblique vs. oblique, perpendicular vs. oblique, parallel vs. perpendicular) as the single independent variable. Both analyses included only the *different* trials in which there was an orientation change (and no geon change) and included only trials in which the participant made a correct response.

For the response time data, a reliable effect of type of comparison was found [$F(3,47) = 10.0$, $MS_e = 14,496$, $p < .0001$]. In order to test the theories of how the orientation between the parts of an object might be discriminated (i.e., metric coding, categorical coding, or no coding), post hoc comparisons using the least significant difference (LSD) test were made.

For the response time data, the value of LSD at $\alpha = .05$ was 49 msec. The oblique-versus-oblique condition produced reliably slower response times than any other condition. No other pairwise comparisons were statistically reliable.

For the error rate data, a reliable effect of type of comparison was found [$F(3,47) = 14.04$, $MS_e = 4.7$, $p < .0001$], and the value of LSD at $\alpha = .05$ was 8.8%. Thus, the oblique-versus-oblique condition produced reliably more errors than did all other conditions. In addition, the parallel-versus-oblique condition and the perpendicular-versus-oblique condition produced reliably more errors than did the parallel-versus-perpendicular condition. No other pairwise comparisons were statistically reliable. Although the response time data showed no difference between the parallel-versus-perpendicular condition and the parallel-versus-oblique condition (or the perpendicular-vs.-oblique condition), note that the trend in the response time data is similar to the statistically reliable trend in the error rate data.

Discussion

The results of Experiment 1 suggest that categorical information is included in the representation used to physically discriminate objects, thereby satisfying Harnad's (1987) first criterion for categorical perception. In Experiment 1, participants found it more difficult to physically discriminate between two objects when the orientation change between their parts fell within the same category (i.e., oblique) than when the orientation change crossed a categorical boundary. The results of Experiment 1 therefore indicate that orientation is coded categorically for physical discrimination.

It is interesting to note that participants found Experiment 1 very difficult, especially the oblique-versus-oblique condition. This unexpected level of difficulty in Experiment 1 may possibly be attributable to the representational similarity of the objects being compared. In the oblique-versus-oblique condition, participants were required to discriminate between two objects that would activate virtually the same representation in memory according to theories that posit categorical coding of orientation (e.g., Biederman, 1987). In the other three conditions, participants were required to discriminate between two objects that would activate virtually the same representation in memory except for the category of relative orientation being coded. Given that support was found for categorical coding, the difficulty in comparing two representationally similar objects may have led to the relative difficulty of the task.

However, Experiment 1 did yield partial support for metric coding of orientation, indicating that at least some metric information is indeed coded for physical discrimination. Not surprisingly, a reliable difference was found in the error rate data between the parallel-versus-perpendicular condition (in which the angular disparity between the two objects was 90°) and all other conditions. Additionally, the above-chance performance in the oblique-versus-oblique condition points to partial metric coding of some sort. Thus, although the results indi-

cate that categorical coding is occurring in the physical discrimination of relative orientation, metric coding is occurring as well.

The no-coding-of-orientation hypothesis was not supported by Experiment 1. If orientation were not being coded at all for physical discrimination, then there should have been no reliable difference between any of the conditions.

Experiment 1 yielded evidence for categorical coding of relative orientation for physical discrimination, thus confirming that the orientation between an object's parts satisfies Harnad's (1987) first criterion for categorical perception. The next step was to investigate whether relative orientation is coded categorically for object recognition, Harnad's second criterion for categorical perception.

EXPERIMENT 2 Object Classification

Experiment 2 was designed to test Harnad's (1987) second criterion for categorical perception (i.e., a stimulus continuum should be classified into several discrete categories) by examining how relative orientation is coded during object recognition. In Experiment 2 a classification task was employed in which the participant had to decide whether two objects (like those depicted in Figure 6) sequentially presented on a computer screen were




Figure 6. Examples of objects that are considered *same* or *different* for Experiment 2. Note that all of the *same* objects are identical; the two geons and differ only in the orientation of the object's parts. All of the *different* objects differ in at least one geon.

composed of the same set of parts. The participant was instructed that any differences in relative orientation were to be ignored. Thus, all of the objects in the top row of Figure 6 are the same by this criterion because the identity of the objects' parts does not change (i.e., all objects are composed of a cone on top of a cylinder). All of the objects in the bottom row of Figure 6 are different because they differ in at least one geon.

Experiment 2 had five conditions. The trials of interest in Experiment 2 are the positive (*same* object) trials. Thus in the following description of the conditions, the two objects seen by a participant during the trial sequence differed only in relative orientation, not in the geons that composed the objects. In the parallel-versus-oblique condition, the participant compared the 0° and 30° versions of an object. In the oblique-versus-oblique condition, the participant compared the 30° and 60° versions of an object. In the perpendicular-versus-oblique condition, the participant compared the 90° and 60° versions of an object. In the parallel-versus-perpendicular condition, the participant compared the 0° and 90° versions of an object. Finally, in the identical condition, the participant compared two objects that were the same in every way. The parallel-versus-perpendicular condition was included only for reasons of counterbalancing, since categorical theories make no strong prediction about this condition (see the introduction). Without the parallel-versus-perpendicular condition, the 30° and 60° versions of each object would be presented twice as often as the 0° and 90° versions of each object.

If relative orientation is coded categorically via the categories of parallel, perpendicular, and oblique for object recognition, then the 30° and 60° objects should activate the same representation in memory because both objects are composed of the same set of parts and have the same relative orientation relationship (i.e., oblique) between their parts. However, the 0° and 30° objects and the 90° and 60° objects should activate different representations in memory, because although both sets of objects are composed of the same set of parts, each has a different relative orientation relationship between their parts (parallel and oblique, and perpendicular and oblique, respectively).

Therefore, if relative orientation is coded categorically, it should be easier, in terms of response times and error rates, for participants to appreciate that the 30° and 60° objects (the oblique-vs.-oblique condition) are the *same* (i.e., are composed of the same set of parts) than it should be for the 0° and 30° objects (the parallel-vs.-oblique condition) and the 60° and 90° objects (the perpendicular-vs.-oblique condition). This pattern of results should occur even though the absolute change in orientation is the same in all cases (i.e., 30°).

If, however, relative orientation is coded only metrically, then performance should depend only on the absolute difference in orientation between the parts of the objects. Thus, the three conditions in which the angular

disparity between the two objects being compared is 30° (the parallel-vs.-oblique, oblique-vs.-oblique, and perpendicular-vs.-oblique conditions) should produce roughly equal amounts of cost (relative to the identical condition) in terms of response times and errors rates because the magnitude of the orientation change is the same in all three conditions. The parallel-versus-perpendicular condition should thus produce slower response times and more errors, because the angular disparity is 90°.

If only the parts' identities are coded, and no relation information is included in the memory representation, there should be no difference between all conditions regardless of the angular disparity.

Method

Participants. Thirty-two college students who reported normal or corrected-to-normal vision participated in the experiment. Participants received extra credit in an introductory psychology course or developmental psychology course at Iowa State University for taking part in the experiment.

Apparatus. The apparatus for collecting data and the stimuli used in Experiment 2 were the same as in Experiment 1.

Procedure. Unless otherwise noted, the procedure was identical to that of Experiment 1. The participant's task in Experiment 2 was to press the left button on the response box if the two objects were composed of the same set of parts (ignoring any orientation changes that might occur) and to press the right button on the response box if the identity of one or more of the parts changed. On each trial, the participant first saw a fixation cue for 504 msec, followed by the first object for 168 msec, followed by a mask for 2,016 msec, followed by the second object for 168 msec, followed by a mask for 126 msec.

Each participant viewed 288 trials, half of which were positive (the two objects were composed of the same set of parts) and half of which were negative (the two objects were composed of a different set of parts). The order in which all trials appeared was generated at random. On the 144 positive trials, each of the 0°, 30°, 60°, and 90° objects were presented 36 times as the first object in the trial sequence. The second object in the trial sequence was selected so that each possible object was equally likely to occur. So, if the first object in the trial sequence was the 0° object, the 0°, the 30°, and the 90° object each followed 12 times as the second object in the sequence.

On the negative trials (*different* trials), each rotation was seen approximately equally often (the actual objects that were presented were chosen randomly). On one third of the negative trials, the top geon changed identity (48 trials); on one third of the negative trials, the base geon changed identity (48 trials); and on the remaining one third of the negative trials, both the top and the base geons changed identity (48 trials).

Results

The results of Experiment 2 can be seen in Figures 7 and 8. The *different* object trials in Experiment 2 had a mean response time of 825 msec and a mean error rate of 17.6%. Only data from trials in which the participant made a correct response were included in the response time ANOVA. For the response time data, a reliable effect of type of comparison was found [$F(4,31) = 6.9$, $MS_e = 2,985$, $p < .0001$]. In order to test the theories of how the orientation between the parts of an object might

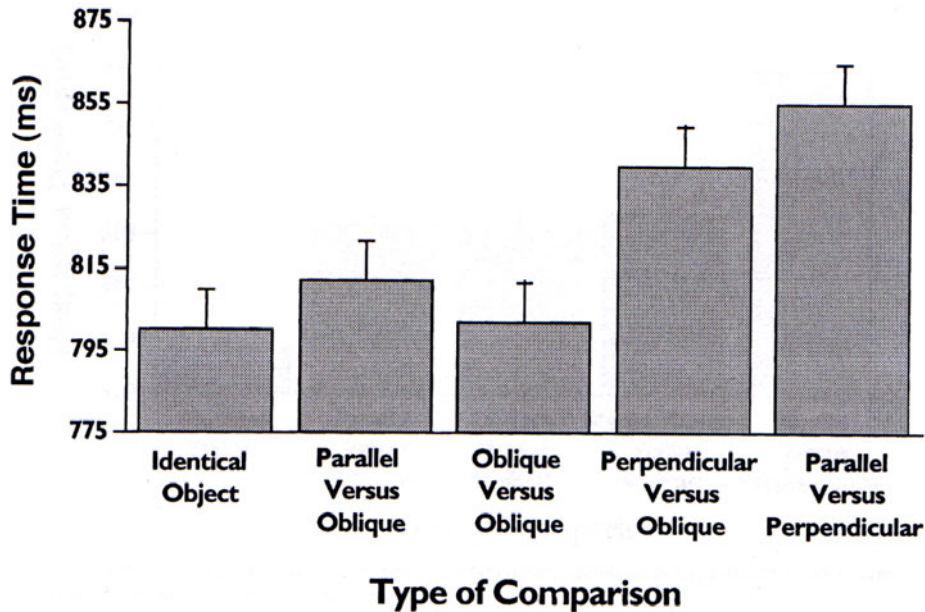


Figure 7. The mean response time (in milliseconds) for each condition in Experiment 2. Error bars represent the standard error of the mean.

be coded (i.e., metric coding, categorical coding, and no coding), post hoc comparisons with the LSD test were conducted.

For the response time data, the value of LSD at $\alpha = .05$ was 28 msec. The LSD comparisons indicated that the perpendicular-versus-oblique condition produced reliably longer response times than did the oblique-versus-oblique condition and the identical condition. Addition-

ally, the parallel-versus-perpendicular condition produced reliably longer response times than did the parallel-versus-oblique, the oblique-versus-oblique, and the identical conditions. No other pairwise comparisons were statistically reliable.

For the error rate data, a reliable effect of type of comparison was found [$F(4,31) = 5.6, MS_e = .004, p < .001$]. Again, post hoc comparisons were made with LSD. The

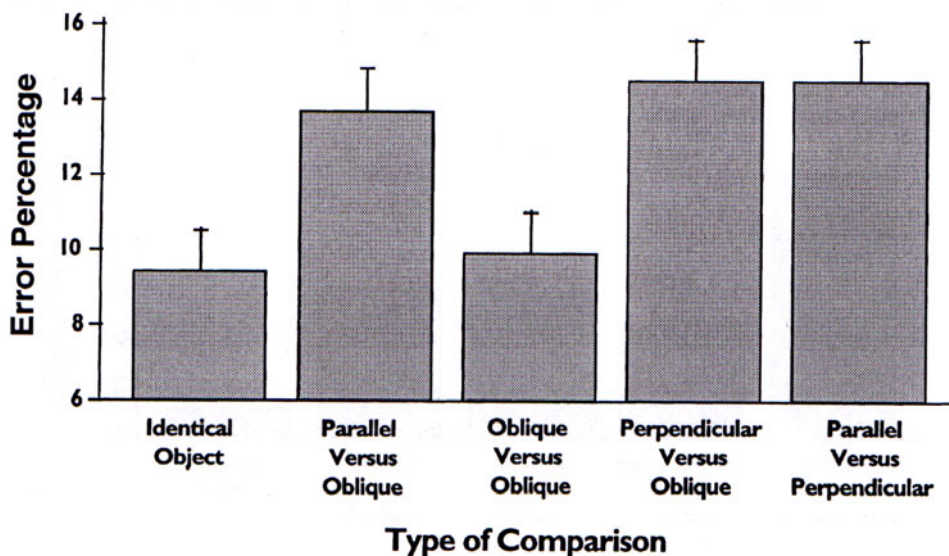


Figure 8. The mean error percentage for each condition in Experiment 2. Error bars represent the standard error of the mean.

value of LSD at $\alpha = .05$ was 3.2%. The LSD comparisons indicated that the oblique-versus-oblique condition produced reliably fewer errors than did the parallel-versus-oblique condition, the perpendicular-versus-oblique condition, and the parallel-versus-perpendicular condition. Likewise, the identical condition produced reliably fewer errors than did the parallel-versus-oblique condition, the perpendicular-versus-oblique condition, and the parallel-versus-perpendicular condition. No other pairwise comparisons were statistically reliable. Although no reliable difference between the oblique-versus-oblique condition and the parallel-versus-oblique condition was found in the response time data, note that the trend in the response time data is identical to the statistically reliable trend in the error rate data.

Discussion

Among the three possibilities for the coding of relative orientation, the pattern of data from Experiment 2 most closely matches that predicted by categorical coding. It was more difficult for participants to judge that two objects were the same in terms of the identity of their parts when the orientation between those parts went from perpendicular to oblique (in terms of both the errors and response times) or from parallel to oblique (in terms of the errors) relative to when the orientation between the parts went from oblique to oblique. This nonlinearity in the costs associated with a change in orientation occurred even though the physical distance between the two objects being compared was 30° in all three conditions. Thus, the results show that the representation used to recognize objects does include categorical information.

However, the results of Experiment 2 also suggest that the representation used to recognize objects does include some metric information as well. The parallel-versus-perpendicular condition produced reliably longer response times than the parallel-versus-oblique condition and the identical condition. Although this trend was not evident in the error rate data, it does indicate that the absolute degree of angular disparity in the parts between the two objects being compared does affect the speed of response, in addition to category membership. Thus, it appears that both metric and categorical information is included in the representation used to recognize objects.

Like Experiment 1, Experiment 2 did not support the no-coding-of-orientation hypothesis. If the object recognition system does not code orientation at all, then no differences should have been found between any of the conditions in Experiment 2. The results of Experiment 2 confirm that people do code the orientation of an object's parts for the purposes of object recognition, as indicated by the reliable performance differences found between the various conditions.

Although Experiments 1 and 2 featured different presentation conditions (i.e., the interstimulus interval [ISI] was shorter in Experiment 1 than in Experiment 2), the purpose of the present research was to test each of two components of categorical perception individually, not

to directly compare the classification task with the physical discrimination task. Much of the research that has reported categorical perception effects has featured different stimulus presentation conditions (and often different paradigms) for the physical discrimination and identification tasks. For example, Calder et al. (1996) and Etcoff and Magee (1992) used an ABX physical discrimination task in which stimuli were presented for 750 and 1,000 msec, respectively (with an intervening 1,000-msec ISI), whereas in their classification task they presented a single stimulus in free vision (see also Campbell, Woll, Benson, & Wallace, 1999; de Gelder, Teunisse, & Benson, 1997; Siegel & Siegel, 1977; Strange & Dittmann, 1984; Zatorre & Halpern, 1979, for other examples of research reporting categorical perception that have featured different presentation conditions for the classification and discrimination tasks).

In Experiment 2, we were interested in examining the characteristics of the long-term memory representation used for object identification. However, some evidence suggests that the representation governing short-term memory for objects may have characteristics that are different from those of the representation governing long-term memory for objects. For example, Ellis and Allport (1986) noted that when the ISI between two comparison objects is reasonably brief, participants base their response on a metrically specific memory representation. At longer ISIs, however, this temporary representation is no longer available and a more permanent representation that includes categorical information is used to drive the participant's response. Ellis and Allport supported their conclusion with some empirical data; at an ISI of 500 msec, participants' comparison of two objects was impaired by varying the metric characteristics of the objects, while at an ISI of 2,000 msec, no such impairment was detected. Similarly, Biederman and Gerhardstein (1993) noted that at short ISIs, participants in a *same-different* classification task may be able to match the second object against an iconic or short-term memory representation of the first object that is presumably metrically specific. However, at longer ISIs, this temporary representation is no longer available and participants must rely on a more permanent long-term memory representation (see Murray, 1999, for a discussion).

If the short-term memory representation for objects does not contain categorical information, as suggested by Biederman and Gerhardstein (1993) and Ellis and Allport (1986), then it may be case that the categorical effects found in Experiment 2 of the present research will disappear simply by reducing the time that participants must hold a representation of the first object in memory. The purpose of Experiment 3 was to test this possibility. Experiment 3 was thus identical to Experiment 2, except that the ISI was reduced to 756 msec. If participants code the objects in Experiment 3 using a metrically specific short-term memory representation, then a set of results predicted by the metric coding of orientation hypothesis should be found (i.e., recognition time should depend




Figure 9. The mean response time (in milliseconds) for each condition in Experiment 3. Error bars represent the standard error of the mean.

only on the difference in absolute orientation). If, however, relative orientation information is also contained in the short-term memory representation used to recognize objects, then the results of Experiment 3 should mirror those of Experiment 2.

EXPERIMENT 3

Method

Participants. Forty-eight college students who reported normal or corrected-to-normal vision participated in the experiment. Participants received extra credit in an introductory psychology course

or developmental psychology course at Iowa State University for taking part in the experiment. Prior to the experiment, a performance criterion of 70% correct was established and data from participants who did not meet this criterion were replaced by data from new participants. Under this criterion, data from 5 participants were replaced by new data.

Apparatus. The apparatus for collecting data and the stimuli used in Experiment 3 were the same as in Experiment 2.

Procedure. Unless otherwise noted, the procedure was identical to that of Experiment 2. On each trial in Experiment 3, participants saw a fixation cue for 504 msec, followed by the first object for 168 msec, followed by a mask for 756 msec, followed by the second object for 168 msec, followed by a mask for 126 msec. As in Experiment 2, the participant's task was to indicate via a button-




Figure 10. The mean error percentage for each condition in Experiment 3. Error bars represent the standard error of the mean.

press whether two objects were composed of the same set of parts while ignoring any orientation change between those parts.

Results

The results of Experiment 3 are shown in Figures 9 and 10. The *different* object trials had a mean response time of 728 msec and a mean error rate of 20.0%.

Only data from trials in which the participant made a correct response were included in the response time ANOVA. For the response time data, a reliable effect of type of comparison was found [$F(4,47) = 6.9$, $MS_e = 2,985$, $p < .0001$]. In order to test the theories of how the orientation between the parts of an object might be coded (i.e., metric coding, categorical coding, and no coding), post hoc comparisons LSD test were conducted.

For the response time data, the value of LSD at $\alpha = .05$ was 23 msec. The LSD comparisons indicated that the identical condition produced reliably faster response times than did all other conditions, and the parallel-versus-perpendicular condition produced reliably slower response times than did all other conditions. No other pairwise comparisons were statistically reliable.

For the error rate data, a reliable effect of type of comparison was found [$F(4,47) = 12.7$, $MS_e = .004$, $p < .0001$]. Again, post hoc comparisons were made with LSD. The value of LSD at $\alpha = .05$ was 2.6%. The LSD comparisons indicated that the identical condition produced fewer errors than all other conditions, and the parallel-versus-perpendicular condition produced more errors than all other conditions. No other pairwise comparisons were statistically reliable.

Discussion

Among the three possibilities for the coding of orientation between the parts of an object, the pattern of data from Experiment 3 closely matches that predicted by metric coding. It was no more difficult for participants to judge that two objects were the same in terms of the identity of their parts when the orientation between those parts went from perpendicular to oblique or from parallel to oblique than when the orientation between the parts went from oblique to oblique. Likewise, participants found it relatively easy to judge that two identical objects were composed of the same set of parts, but relatively difficult to classify two objects as being the same when the angular disparity was 90°. Thus, these results support the idea that the orientation between the parts of an object is coded metrically at brief ISIs.

GENERAL DISCUSSION

The purpose of the present experiments was to investigate whether the orientation between an object's parts constitutes an instance of categorical perception as defined by Harnad (1987). Three experiments were conducted examining how relative orientation is coded for physical discrimination (Harnad's first criterion for categorical perception) and for object classification (Har-

nad's second criterion for categorical perception). In Experiment 1, participants found it easier to judge that two objects were physically different when they had a different relative orientation relationship (i.e., parallel vs. oblique or perpendicular vs. oblique) than when two objects had the same relative orientation relationship (i.e., oblique vs. oblique). In Experiment 2, participants found it more difficult to classify two objects as being composed of the same set of parts when the relative orientation of the parts crossed a categorical boundary than when the relative orientation of the parts did not cross a categorical boundary. In contrast, Experiment 3 showed that when the length of time between the presentation of the first and second object was reduced, categorical coding of relative orientation for object recognition did not occur. Thus, categorical information does appear to be included in the representation used to recognize objects, but only at extended ISIs.

The results of the classification tasks (Experiments 2 and 3) thus suggest that the way in which the orientation between an object's parts is coded for object recognition differs as a function of temporal duration. The results suggest that in early processing, only metric information is included in the representation used to code relative orientation, whereas categorical information is not coded until later. These results are consistent with previous work in which metric invariance was found at long ISIs in both object-naming tasks (e.g., Biederman & Cooper, 1992) and object-matching tasks (Biederman & Gerhardstein, 1993; Ellis & Allport, 1986; but see Murray, 1999). Further, some models of shape recognition that code relative orientation as part of the representation used to recognize objects (i.e., Hummel & Biederman, 1992; Hummel & Stankiewicz, 1998) code orientation differently at different levels of stimulus representation. In the Hummel and Biederman model, for example, the orientation between an object's parts is coded metrically in early stages of processing (specifically, Layer 2 in their seven-layer neural network model), and this metric information is then used to build a categorical representation of relative orientation in later stages of processing (i.e., Layer 5). The results of Experiment 3 may therefore reflect judgments based on a representation analogous to Layer 2 in the Hummel and Biederman model, whereas the results of Experiment 2 may reflect judgments based on a representation derived from Layer 5.

Taken together, the results of Experiments 1 and 2 meet Harnad's (1987) two criteria for determining that a perceptual event is an instance of categorical perception. These results thus suggest that categorical perception does extend beyond simple, lower level stimuli (such as colors and phonemes) and can apply to higher level aspects of visual perception. Other recent studies have shown that categorical perception also occurs in the perception of facial expression (Calder et al., 1996; Etcoff & Magee, 1992) and facial identity (Beale & Keil, 1995). On the basis of Experiments 1 and 2 in the present research, categorical perception also applies to certain aspects of

object perception. Whether other types of within-object categorical relations besides relative orientation (i.e., relative position of an object's parts or the relative size of an object's parts) also are coded categorically for physical discrimination is still an open question.

The categorical perception of relative orientation appears to become apparent only when the classification task places demands on long-term memory. However, this finding does not necessarily weaken the claim that relative orientation is perceived categorically. In categorical perception research reported in the literature (see, e.g., Calder et al., 1996; Etcoff & Magee, 1992), classification is usually measured by having participants compare a stimulus to a long-term memory representation (often by presenting a single stimulus in free vision and having participants name the stimulus). Therefore, the finding in the present research that relative orientation meets Harnad's (1987) two criteria only for categorical perception when long ISIs are employed is consistent with previous findings.

Some researchers (e.g., Studdert-Kennedy, Liberman, Harris, & Cooper, 1970) have argued that categorical perception implies a participant's complete inability to physically discriminate stimuli from within the same category (thus producing at-chance performance on a physical discrimination task). This criterion was not met in Experiment 1. Although some debate still exists on this issue, many authors (e.g., Eimas, Miller, & Jusczyk, 1987; Etcoff & Magee, 1992; Harnad, 1987) have argued that at-chance performance for within-category physical discrimination is too stringent a criterion for categorical perception and that few, if any, perceptual events actually meet this criterion. In fact, most instances of categorical perception reported in the literature showed reasonably good within-category discrimination performance, but reliably better between-category discrimination performance (see, e.g., Macmillan, Kaplan, & Creelman, 1977). Most authors agree that this does not weaken the concept of categorical perception (see Beale & Keil, 1995; Harnad, 1987; but also see Massaro, 1987). Physical discrimination performed above chance, these authors have argued, implies only that some metric information is included in the representation used for the physical discrimination of certain sets of stimuli.

In the present research, only orientations between 0° and 90° were tested. However, the category "oblique," broadly defined, would also include orientations greater than 90° (e.g., 120°). Thus, the results of these experiments do not extend beyond angles greater than 90°, and it may be the case that oblique angles less than 90° are coded differently in memory than oblique angles greater than 90°. Hummel and Biederman (1992), for example, posited that "oblique and pointing left" is coded in memory as a category separate from "oblique and pointing right," but this assumption has never been tested empirically. Additional research is therefore needed to further define what constitutes the "oblique" category.

The finding that metric information is included in the representation used to physically discriminate and classify objects is consistent with both categorical theories (e.g., Hummel & Biederman, 1992; Hummel & Stankiewicz, 1998) and metric theories (e.g., Bulthoff & Edelman, 1992; Tarr & Pinker, 1990, 1991; Ullman, 1989, 1996). For example, Hummel and Stankiewicz proposed a model that codes the spatial relations between an object's parts using both categorical and metric information. Their model assumes an output function that is non-linear across categorical boundaries (e.g., the boundary between parallel and oblique). Thus, categorical information does not have to be coded at the complete exclusion of metric information, an interpretation consistent with the present findings. However, the finding that categorical information is also included in the representation used to physically discriminate objects is compatible only with categorical theories, since metric theories assume that no categorical coding occurs at any level of stimulus representation.

Although a large volume of literature has accumulated that speaks to the perception of spatial relations between two or more objects (see, e.g., Hellige & Michimata, 1989; Kosslyn et al., 1989), very little empirical work has been done before to directly investigate the coding of relations within the same object. Several authors (e.g., Humphreys & Riddoch, 1994, 1995; Saiki & Hummel, 1999) have suggested that the coding of between-object relations is different from the coding of within-object relations, and so it is unclear what, if any, the relationship between the two processes might be. In the research literature, the term *spatial relations* is usually used synonymously with *relative position* (see, e.g., Cave & Kosslyn, 1993; Cooper & Wojan, 2000; Hummel & Stankiewicz, 1996); however, the results of the present experiments suggest that other types of spatial relations (e.g., orientation) are coded as well.

REFERENCES

- APPELLE, S. (1972). Perception and discrimination as a function of stimulus orientation: The oblique effect in man and animals. *Psychological Bulletin*, **78**, 266-278.
- BEALE, J. M., & KEIL, F. C. (1995). Categorical effects in the perception of faces. *Cognition*, **57**, 217-239.
- BIEDERMAN, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.
- BIEDERMAN, I., & COOPER, E. E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 121-133.
- BIEDERMAN, I., & GERHARDSTEIN, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 1162-1182.
- BIEDERMAN, I., & GERHARDSTEIN, P. C. (1995). Viewpoint-dependent mechanisms in visual object recognition: Reply to Tarr and Bulthoff (1995). *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 1506-1514.
- BORNSTEIN, M. H. (1987). Perceptual categories in vision and audition. In S. Harnad (Ed.), *Categorical perception* (pp. 287-300). Cambridge: Cambridge University Press.

- BORNSTEIN, M. H., & KORDA, N. O. (1984). Discrimination and matching within and between hues measured by reactions times: Some implications for categorical perception and levels of information processing. *Psychological Research*, **46**, 207-222.
- BULTHOFF, H., & EDELMAN, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences*, **89**, 60-64.
- CALDER, A. J., YOUNG, A. W., PERRETT, D. I., ETCOFF, N. L., & ROWLAND, D. (1996). Categorical perception of morphed facial expressions. *Visual Cognition*, **3**, 81-117.
- CAMPBELL, R., WOLL, B., BENSON, P., & WALLACE, S. (1999). Categorical perception of face actions: Their role of sign language and in communicative facial displays. *Quarterly Journal of Experimental Psychology*, **52A**, 67-95.
- CAVE, C. B., & KOSSLYN, S. M. (1993). The role of parts and spatial relations in object identification. *Perception*, **22**, 229-248.
- COOPER, E. E., & BIEDERMAN, I. (1993). *Metric versus viewpoint-invariant shape differences in visual object recognition*. Poster presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Sarasota, FL.
- COOPER, E. E., & WOJAN, T. J. (2000). Differences in the coding of spatial relations in face and object identification. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 470-488.
- DE GELDER, B., TEUNISSE, J., & BENSON, P. (1997). Categorical perception of facial expressions: Categories and their internal structure. *Cognition & Emotion*, **11**, 1-23.
- EIMAS, P. D., MILLER, J. L., & JUSCZYK, P. W. (1987). On infant speech perception and the acquisition of language. In S. Harnad (Ed.), *Categorical perception* (pp. 161-195). Cambridge: Cambridge University Press.
- ELLIS, R., & ALLPORT, D. A. (1986). Multiple levels of representation for visual objects: A behavioral study. In A. G. Cohn & J. R. Thomas (Eds.), *Artificial intelligence and its applications* (pp. 245-257). New York: Wiley.
- ESSOCK, E. A. (1980). The oblique effect of stimulus identification considered with respect to two classes of oblique effects. *Perception*, **9**, 37-46.
- ETCOFF, N. L., & MAGEE, J. J. (1992). Categorical perception of facial expressions. *Cognition*, **44**, 227-240.
- HARNAD, S. (1987). Introduction: Psychophysical and cognitive aspects of categorical perception: A critical overview. In S. Harnad (Ed.), *Categorical perception* (pp. 1-25). Cambridge: Cambridge University Press.
- HELLIGE, J. B., & MICHIMATA, C. (1989). Categorization versus distance: Hemispheric differences for processing spatial information. *Memory & Cognition*, **17**, 770-776.
- HUMMEL, J. E., & BIEDERMAN, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, **99**, 480-517.
- HUMMEL, J. E., & STANKIEWICZ, B. J. (1996). Categorical relations in shape perception. *Spatial Vision*, **3**, 201-236.
- HUMMEL, J. E., & STANKIEWICZ, B. J. (1998). Two roles for attention in shape perception: A structural description model of visual scrutiny. *Visual Cognition*, **5**, 49-79.
- HUMPHREYS, G. W., & RIDDOCH, M. J. (1994). Attention to within-object and between-object spatial representations: Multiple sites for visual selection. *Cognitive Neuropsychology*, **11**, 207-241.
- HUMPHREYS, G. W., & RIDDOCH, M. J. (1995). Separate coding of space within and between perceptual objects: Evidence from unilateral neglect. *Cognitive Neuropsychology*, **12**, 283-311.
- KHOLMEYER, S. W. (1992). Picture Perception Lab: A program for picture experiments on the Macintosh II. *Behavior Research Methods, Instruments, & Computers*, **24**, 67-71.
- KOSSLYN, S. M. (1994). *Image and brain*. Cambridge, MA: MIT Press.
- KOSSLYN, S. M., KOENIG, O., BARRETT, A., CAVE, C. B., TANG, J., & GABRIELI, J. D. E. (1989). Evidence for two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 723-735.
- LIBERMAN, A. M., HARRIS, K. S., HOFFMAN, H. S., & GRIFFITH, B. C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, **54**, 358-368.
- LOFTUS, G. R., & LOFTUS, E. F. (1988). *Essence of statistics* (2nd ed.). New York: Random House.
- MACMILLAN, N. A., KAPLAN, H. L., & CREELMAN, C. D. (1977). The psychophysics of categorical perception. *Psychological Review*, **84**, 452-471.
- MASSARO, D. W. (1987). Categorical partition: A fuzzy-logical model of categorization behavior. In S. Harnad (Ed.), *Categorical perception* (pp. 254-283). Cambridge: Cambridge University Press.
- MURRAY, J. E. (1999). Orientation-specific effects in picture matching and naming. *Memory & Cognition*, **27**, 878-889.
- ROSCH, E. (1975). The natural mental codes for color categories. *Journal of Experimental Psychology: Human Perception & Performance*, **1**, 303-322.
- SAIKI, J., & HUMMEL, J. E. (1999). *Interactions between connectedness and axis-alignment in orientation sensitivity: Evidence for axis-based specification of within-object spatial relations*. Manuscript submitted for publication.
- SIEGEL, J. A., & SIEGEL, W. (1977). Categorical perception of tonal intervals: Musicians can't tell sharp from flat. *Perception & Psychophysics*, **21**, 399-407.
- SMITH, E. E., & MEDIN, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- SNOWDON, C. T. (1987). A naturalistic view of categorical perception. In S. Harnad (Ed.), *Categorical perception* (pp. 332-354). Cambridge: Cambridge University Press.
- STRANGE, W., & DITTMANN, S. (1984). Effects of discrimination training on the perception of /r-/l/ by Japanese adults learning English. *Perception & Psychophysics*, **36**, 131-145.
- STUDDERT-KENNEDY, M., LIBERMAN, A. M., HARRIS, K. S., & COOPER, F. S. (1970). Motor theory of speech perception: A reply to Lane's critical review. *Psychological Review*, **77**, 234-249.
- TARR, M. J., & BULTHOFF, H. (1995). Is human object recognition better described by geon structural descriptions of by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 1494-1505.
- TARR, M. J., & PINKER, S. (1990). When does human object recognition use a viewer-centered reference frame? *Psychological Science*, **1**, 253-256.
- TARR, M. J., & PINKER, S. (1991). Orientation-dependent mechanisms in shape recognition: Further issues. *Psychological Science*, **2**, 207-209.
- ULLMAN, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, **32**, 193-254.
- ULLMAN, S. (1996). *High-level vision: Object recognition and visual cognition*. Cambridge, MA: MIT Press.
- ZATORRE, R. J., & HALPERN, A. R. (1979). Identification, discrimination, and selective adaptation of simultaneous musical intervals. *Perception & Psychophysics*, **26**, 384-395.

(Manuscript received March 3, 1999;
revision accepted for publication February 28, 2000.)