

Asymmetries in categorization, perceptual discrimination, and visual search for reference and nonreference exemplars

OLIVIER CORNEILLE

Catholic University of Louvain, Louvain-la-Neuve, Belgium

ROBERT L. GOLDSTONE and SARAH QUELLER

Indiana University, Bloomington, Indiana

and

TIMOTHY POTTER

Catholic University of Louvain, Louvain-la-Neuve, Belgium

In two experiments, we examined the representation, treatment, and attention devoted to the members of reference (i.e., club members) and nonreference (i.e., not club members) categories. Consistent with prior work on category interrelatedness (e.g., Goldstone, 1996; Goldstone, Steyvers, & Rogosky, 2003), the findings reveal the existence of asymmetric representations for reference and nonreference categories, which, however, decreased as expertise and familiarity with the categories increased (Experiments 1 and 2). Participants also more readily judged two reference exemplars as being the same than they did two nonreference exemplars (Experiment 1) and were better at detecting reference than nonreference exemplars in a set of novel, category-unspecified exemplars (Experiment 2). These findings provide evidence for the existence of a feature asymmetry in the representation and treatment of exemplars from reference and nonreference categories. Membership in a reference category acts as a salient feature, thereby increasing the perceived similarity and detection of faces that belong in the reference, in comparison with the nonreference, category.

Classic work on categorization has provided evidence that categorization enhances the perception of between-category differences and the perception of within-category resemblances. This effect has been shown to apply to the judgment of physical (e.g., Harnad, 1987; Tajfel & Wilkes, 1963) and social (e.g., Eiser & Van der Pligt, 1984; Krueger & Rothbart, 1990) stimuli and has proven larger under conditions of enhanced judgment uncertainty, such as when individuals have to map their judgment onto an unfamiliar measurement unit (Corneille, Klein, Lambert, & Judd, 2002). In past research, this effect has also been examined for judgments of multifaceted stimuli (e.g., Corneille & Judd, 1999; Goldstone, 1996; Livingston, Andrews, & Harnad, 1998) and the consequences of this bias for memory have been addressed (e.g., Corneille, Huart, Becquart, & Brédart, 2004; Huart, Corneille, & Becquart, 2005; Krueger & Clement, 1994; Taylor, Fiske, Etcoff, & Ruderman, 1978).

In this work, the categories involved have generally been given equal symbolic and attentional status. In some circumstances, however, categories may be asymmetric, in the sense that one category is the reference and category nonmembers are defined merely as lacking those features that characterize the members of the reference category. Goldstone and colleagues (Goldstone, 1996; Goldstone, Steyvers, & Rogosky, 2003) examined the consequences of these reference effects for category learning. These authors proposed that reference categories are likely to be organized around a prototype, whereas nonreference categories are likely to be distorted away from the category to which they refer. They further suggested that the latter process may result in the emergence of more caricatured representations for nonreference than for reference categories.

In Goldstone and colleagues' (Goldstone, 1996; Goldstone et al., 2003) work, the category exemplars consisted of a series of faces that were located along the left end or the right end of a continuum of morphed faces. The procedure controlled for prior familiarity with the exemplars and held constant both the differences between adjacent exemplars and within-category variability. Category reference was operationalized through club membership (the two categories were *club members* vs. *not club members*) or learning order (participants first learned about Category A members and only then learned about Category B members). The participants were asked to catego-

This research was funded in part by FNRS Grant 1.5100.03F awarded to the first author, and in part by NSF Grant 0125287 to the second author. We thank Stephanie Demoulin, Peter Freytag, Chick Judd, and Craig McGarty for their comments on an earlier version of the manuscript. Correspondence concerning this article may be addressed to O. Corneille, Catholic University of Louvain, PSP-PSOR, 10, Place du Cardinal Mercier, B-1348 Louvain-la-Neuve, Belgium (e-mail: olivier.corneille@psp.ucl.ac.be).

size exemplars into one of two categories (e.g., club vs. not club members). Feedback following each decision allowed them to progressively improve their categorization accuracy. Accuracy scores on this category-learning task revealed that the tendency to categorize extreme exemplars better than typical category exemplars was larger for the nonreference category (e.g., not club members) than for the reference category (e.g., club members). In other words, a larger caricature advantage was found in categorizing the members of the nonreference group.

Whereas this work has provided preliminary evidence for the role of category reference in the representation of categories, there are a number of questions that this line of research has left unanswered, two of which are examined in this contribution. The first issue concerns people's ability to discriminate between exemplars from reference and nonreference categories and people's willingness to report on the similarity of exemplars from these categories: Are subtle differences between members from a reference category more easily or less easily noticed than differences of the same magnitude between members of a nonreference category? Independently of perceptual discrimination, are people more or less likely to make decisions indicative of beliefs about stronger similarities among reference exemplars than to make decisions indicative of beliefs about those among nonreference exemplars? The second issue concerns people's ability to detect exemplars from reference and nonreference categories: Are people better at detecting the presence of reference exemplars among a set of distractors than at detecting the presence of nonreference exemplars?

There are conflicting accounts for whether the members of reference (club) or nonreference (not-club) categories should be expected to be more distinct from each other. On the one hand, reference category members might be expected to be more distinct, because increased attention to these members would emphasize their unique attributes. In social psychology, members of one's own group (presumably, a reference category) are generally more individuated than members of other groups (e.g., Mullen & Hu, 1989; Read & Urada, 2003). In face perception, a reliable advantage has been found for the identification of faces from one's own race over that of faces from others (Bothwell, Brigham, & Malpass, 1989). Part of these effects seems to be due to perceptual systems becoming selectively tuned to distinguish among habitually experienced faces. On the other hand, reference faces might be expected to be less distinct, because they all share a common, salient category membership. Levin (2000) has described evidence that faces belonging to salient categories are more similar to one another than are faces belonging to backgrounded categories. If reference categories are more salient than nonreference categories, faces sharing a reference category membership might be expected to be judged more similar than are faces sharing a nonreference category membership. A major goal of the present experiments was to decide between these accounts.

Given the many influences of training on featural encodings (see Palmeri, Wong, & Gauthier, 2004, for a re-

cent review), it is possible that categorization training induces asymmetric features for reference and nonreference categories. More specifically, if people focus mostly on establishing membership in the reference category, members of the reference category may develop additional psychological features, in comparison with members of the nonreference category. In turn, perceptual features that support the abstract *category member* label may be found for the members of the reference category. Several empirical predictions would follow.

First, ambiguous items equally similar to reference and nonreference categories will tend to be placed in the reference category, because the midway item partially possesses the features of the reference category and the presence of a feature is more salient than its absence. The difference between not possessing a feature and partially possessing it is larger than the difference between partially possessing a feature and fully possessing it (Tversky, 1977). Consistent with this idea, an item that morphs midway between a distinctive and a nondistinctive item is judged to be more similar to the distinctive item (Tanaka, Giles, Kremen, & Simon, 1998). Extending this idea to the present situation, participants should more likely misclassify a nonreference exemplar similar to the midway item into the reference category than to misclassify a reference exemplar similar to the midway item into the nonreference category. Therefore, assimilation to the reference category should be obtained.

Second, items that belong to the reference category should be perceived as more similar to one another than items that belong to the nonreference category will be. Items sharing reference category membership should become more subjectively similar because of the acquired reference features that they share. This claim is based on the finding that objects become more similar to one another as the number of common features increases (Tversky, 1977). For example, a circle and a triangle become more similar if the same square pedestal is placed beneath each. Similarly, items that are placed into a common salient reference category would also become more similar to one another. Interestingly, Tversky also suggested that subjective similarity between identical items increases as the number of common features increases. Thus, *same* judgments for two identical stimuli should be more frequent for items belonging to the reference category than for items belonging to the nonreference category. Accordingly, the prediction for a perceptual discrimination task is that perceivers should be more likely to correctly respond *same* when an item that is presented twice belongs to a reference, rather than to a nonreference, category. However, there should also be more incorrect *same* judgments when two different items are presented that belong to a reference, rather than to a nonreference, category.

Third, asymmetries in a search task may be predicted. If training on a reference versus nonreference categorization task causes the reference category items to acquire additional features, relative to the nonreference category, people should be better at detecting reference than at detecting nonreference category items. This logic parallels

Treisman and Gormican's (1988) argument that the letter R can be detected among Ps more efficiently than a P can be detected among Rs because the R has a psychologically salient feature, the diagonal slash \, that P does not possess. Considerable evidence suggests that detecting an object or category is easier if it is identifiable on the basis of a present feature, rather than on the basis of an absent feature (Agostinelli, Sherman, Fazio, & Hearst, 1986; Quinlan, 2003). One might predict that feature-based asymmetries should exist only for hardwired perceptual features, such as oriented lines and colors (Treisman & Gelade, 1980). However, there are also influences of experience on what counts as a psychological feature. Highly familiar conjunctions of simple lines act as features for search tasks (Shiffrin & Lightfoot, 1997), and searching for unfamiliar objects among familiar objects is not as difficult as the converse task (Wang, Cavanagh, & Green, 1994).

Overview of the Experiments

As in the original experiments by Goldstone et al. (2003), we used multifaceted stimuli—namely, faces. A morphing program allowed us to generate faces that were previously unknown to the participants, thereby controlling for prior beliefs and expectations. The neighboring faces along a morph continuum differed by constant amounts from each other, thereby holding constant the variability of the two face categories and the physical differences between adjacent pairs of faces. In the two experiments we conducted, the faces that half of the participants saw lying on the left side of the continuum were referred to as *club* faces and those lying on the right side of the continuum were referred to as *not-club* faces (Condition 1; left-end referent). Labeling was reversed for the other half of the participants (Condition 2; right-end referent). Each category comprised an equal number of face exemplars, thereby controlling for category size, and all the faces were presented the same number of times, thereby controlling for familiarity.

In Experiment 1, participants completed a category-learning task, followed by a perceptual discrimination task. In the category-learning task, they sequentially viewed the various face exemplars and predicted the category membership of each. As in Goldstone et al. (2003), feedback was provided following each decision, and this helped the participants to progressively learn to correctly assign the faces into the club and not-club categories. Unlike in Goldstone et al., the category prototype was never presented to the participants. This modification allowed us to examine whether representational asymmetries would survive in the absence of exposure to the prototype. In addition, the statistical power of the experiment was enhanced, allowing us to detect whether asymmetries (i.e., caricature effect and assimilation toward the reference category) would be magnified or weakened as familiarity with the categories increased. In the subsequent perceptual discrimination task, the participants were sequentially presented with pairs of faces. On a given trial, they either saw the same face presented twice or saw two faces that were adjacent on the morph continuum. The participants

had to decide whether the two faces in each presented pair were the same or different. This perceptual discrimination task made it possible to examine the effect of reference on the participants' ability to discriminate between adjacent faces and to examine the participants' overall readiness to judge two faces as being the same.

In Experiment 2, the participants completed a category-learning task, followed by a visual search task. The category-learning task was similar to that in Experiment 1, except that exposure time was held constant for all the faces. In the visual search task, the participants had to decide as quickly and accurately as possible whether or not a particular target face was among a set of previously unseen faces. We manipulated whether the target face was a club or a not-club face and whether it was present or absent. This task was used to examine how adept the participants would be at correctly identifying reference and nonreference exemplars against a background of novel distractor faces.

The present experiments are a major extension to previous experiments on asymmetries in category representations due to category labeling (Goldstone, 1996; Goldstone et al., 2003). In particular, these previous experiments showed an influence of category labeling on categorization performance itself. The present experiments extended the influence of category labeling to separate tasks not directly related to categorization. Accordingly, they are consistent with the general campaign to chart the importance of categorization for tasks other than classification (Markman & Ross, 2003). Moreover, the particular tasks potentially affected have traditionally been considered to be perceptual and attentional tasks. An effect of category labeling on a simple task of deciding whether two objects are identical or of picking out an object from a set of distractors might be surprising for accounts that draw a sharp boundary between perceptual and conceptual tasks (Pylyshyn, 1999). These kinds of effects would, however, be consistent with the results in a growing literature suggesting that perceptual representations can be influenced by experience, task demands, and learned categories (Corneille et al., 2002; Goldstone, 1994b, 1998; Levin, 2000; Wang et al., 1994).

EXPERIMENT 1

Method

Participants. Two hundred eighty-two undergraduate students from Indiana University served as participants in order to fulfill a course requirement. The students were randomly assigned to the two labeling conditions.

Materials. The stimuli were faces that were generated by morphing between photographs of two bald, male, European American heads selected from Kayser (1997). Previous research has suggested that morphs generated from the two selected faces do not introduce conspicuous nonlinearities between physical and psychological scalings (Goldstone & Steyvers, 2001). The morph sequence of eight faces used is shown in Figure 1. Each of the morphs was automatically generated, using a morphing technique described by Steyvers (1999). Applying this technique, the main contours in the face images were delineated by 127 control lines. These control lines served to align the features of the two faces. In the warping phase of this

morphing algorithm, correspondences were calculated between the pixels of all the images to be morphed. Then, in the cross-dissolving phase, the grayscale values of corresponding pixels were blended to create the grayscale values of the resulting morph image. The faces on the left and right ends of Figure 1 are actual faces, and the six intermediate faces are blends of the two actual faces, with the proportion of the rightmost face beginning at 0% for the leftmost face and shifting along the series in equal 14.29% increments.

The prototype for each category can be defined as the central face within the category’s set of four faces. The actual prototypes are not part of the stimulus set. The left-end category consists of Faces 1–4, with Faces 2 and 3 straddling the prototype. Similarly, the right-end category consists of Faces 5–8, with Faces 6 and 7 straddling the prototype. The caricature of a category is defined as the face that is least like the faces from the other category. Each face was displayed in grayscale with 256 possible brightness values per pixel (1 pixel = 0.034 cm) and measured 14.48 cm tall × 11.68 cm wide. Each face was photographed against a dark background and was displayed on a white Apple iMac computer screen. The average viewing distance was 46 cm.

Procedure. The stimuli were divided into club members and not club members. The dividing line between club members and not club members is shown by the vertical line in Figure 1. For half of the participants, those in Condition 1, the first four faces were club members, and the last four faces were not club members. For the other half of the participants, those in Condition 2, the first four faces were not club members, and the last four faces were club members. The latter factor (i.e., condition) is basically a counterbalancing, and it will be ignored in the remainder of this article (mirror effects were actually obtained within both condition groups). Thus, for the sake of clarity, we will consider eight type of faces here : Club1 (i.e., Face 1 in Condition 1 and Face 8 in Condition 2), Club2 (i.e., Face 2 in Condition 1 and Face 7 in Condition 2), Club3 (i.e., Face 3 in Condition 1 and Face 6 in Condition 2), Club4 (i.e., Face 4 in Condition 1 and Face 5 in Condition 2), Not-Club4 (i.e., Face 5 in Condition 1 and Face 4 in Condition 2), Not-Club3 (i.e., Face 6 in Condition 1 and Face 3 in Condition 2), Not-Club2 (i.e., Face 7 in Condition 1 and Face 2 in Condition 2), and Not-Club1 (i.e., Face 8 in Condition 1 and Face 1 in Condition 2). The experiment was divided into two phases: the category-learning task and the perceptual discrimination task.

For the category-learning task, the participants were instructed as follows: “You will see faces appear on the screen. Half of them belong to a certain club, while the remaining half do not. If you think that a face belongs to the club, press the ‘Y’ key for ‘Yes.’ If you think that it does not belong to the club, press the ‘N’ key for ‘No.’” Next, each trial began with a face appearing on the screen. The face

remained on the screen until the participant pressed the Y or N key. Immediately after pressing one of the keys, feedback was given to the participant. A ✓ or an X indicated whether or not the participant was correct or incorrect, respectively. In addition, written feedback was provided in the form of “Yes, this face is a club member;” “No, this face is not a club member;” “Yes, this face is not a club member;” or “No, this face is a club member.” The feedback was erased from the screen after 1.5 sec. The blank interval between trials was 1 sec. The category-learning task included 30 repetitions of the eight faces shown in Figure 1, for a total of 240 trials. The order of the 240 trials was randomized. The placement of a face’s center was also randomized within a 6 × 6 cm square in the center of the screen. The participants were given breaks every 80 trials. During these breaks, the participants were informed of their accuracy and speed during the preceding block.

In the second phase of the experiment, the perceptual discrimination task, the participants were instructed that they would see displays with two faces on the screen. Their task was to decide whether the faces were exactly identical or differed in any way at all. The participants were warned that all of the faces would be highly similar to one another. The participants pressed the s key to indicate a *same* response and the d key to indicate a *different* response. The computer gave the participants trial-by-trial feedback by presenting either a ✓ or an X for correct and incorrect responses, respectively. On each trial, the two faces to be compared were selected from the set of faces used during category learning. A pair of faces was presented simultaneously on the screen, separated both horizontally and vertically by 5 cm. The vertical displacement prevented the participants from directly comparing face features at a particular height on the screen. Each participant made 270 *same/different* judgments, equally divided into *same* and *different* trials. On *same* trials, one of the eight faces in Figure 1 was randomly selected and presented twice. On *different* trials, one of the seven pairs of adjacent faces in Figure 1 was randomly selected and displayed. The pair of faces remained on the screen until the participants responded. Immediately after the s or d key had been pressed, feedback was provided, and after 1.5 sec, the screen was erased. The blank interval between trials was 1 sec.

Results

We divided the participants’ categorization responses into three blocks of 80 trials each. We then removed from the analyses those participants (*n* = 12) who had not achieved 70% correct categorizations at the end of the third and last categorization block. In all the analyses, we averaged across the multiple observations collected for the same configura-

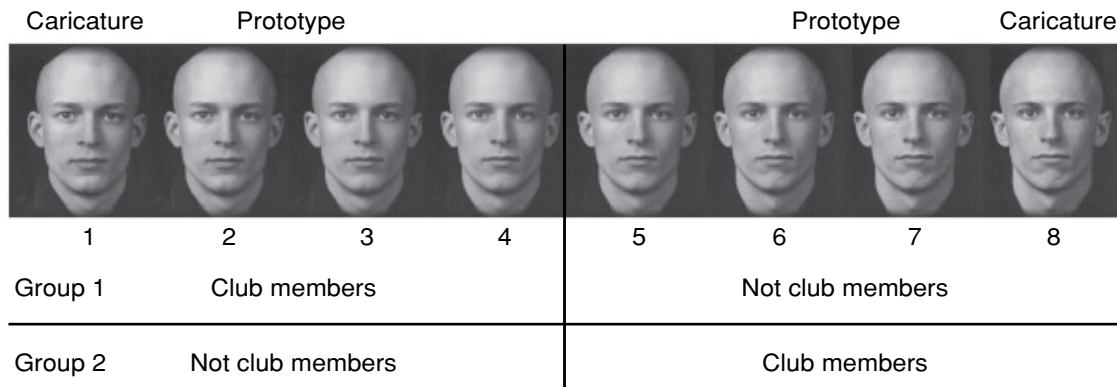


Figure 1. In Experiment 1, a morph sequence of eight faces was divided into two categories. For half of the participants, the left five faces belonged to a category of *club members*, and the remaining faces were labeled as *not club members*. For the other half of the participants, these labels were reversed.

tion (e.g., when considering how participants categorized Club 1 in Block 1, we averaged across data obtained on the 10 presentations of Club 1 in Block 1).

Category-learning task. We first ran a MANOVA on the categorization accuracy scores, with the eight levels of the face factor within the three blocks as within-participants factors. Main effects for face [$F(7,1883) = 492.3, p < .001$] and for block [$F(2,538) = 411.05, p < .001$] were obtained. A face \times block interaction was also found [$F(14,3766) = 9.27, p < .001$]. Additional analyses conducted on a dichotomous club factor (i.e., club vs. not-club faces) clarified the meaning of the latter effects: categorization was more accurate for club faces ($M = .85, SD = .069$) than for not-club faces ($M = .82, SD = .083$) [$F(1,269) = 66.08, p < .001$], and this effect decreased across blocks [$F(2,538) = 15.89, p < .001$].

Figure 2 reports the mean categorization accuracy scores across the various levels of the face factor for the first, second, and third training blocks. As can be seen, accuracy scores were on average higher for the club faces. This difference decreased over blocks (mainly between Blocks 1 and 2). Because lower accuracy scores reflect assignments to the alternative category, this finding may be reinterpreted as follows: Consistent with the predictions, not-club faces were assimilated to the club category more than club faces were assimilated to the not-club category, and this relative assimilation toward the club category decreased across blocks.

We also examined whether Goldstone et al.'s (2003) caricature advantage for not-club faces could be replicated in the context of the present experiment, which involved no prototype, and whether this effect would prove sensitive to the block factor. Because no actual prototype was presented in the present experiment, we approximated accuracy scores for the typical club face by averaging across accuracy scores for Club Face 2 and Club Face 3, and we approximated accuracy scores for the typical not-club face

by averaging across accuracy scores for Not-Club Face 2 and Not-Club Face 3. Then, we established a caricature advantage score by computing the result of the accuracy for Not-Club Face 1 minus the accuracy for the typical not-club face and subtracting from that the result of accuracy for Club Face 1 minus the typical club face accuracy. This stronger caricature advantage for the not-club faces, as compared with the club faces, decreased across the levels of the block factor [$F(2,538) = 7.99, p < .001$], although the advantage remained significant within each of the three blocks (M for Block 1 = .066, $SD = .238$; M for Block 2 = .017, $SD = .127$; M for Block 3 = .017, $SD = .098$; all $ps < .03$).

Perceptual discrimination task. We had eight possible scores for the trials involving the same face presented twice and seven possible scores for the trials involving two different faces. We were interested in the impact of category reference on perceived within-category variability, so the analysis was conducted after dropping out the Club4/Not-Club4 pair that crossed the category boundary, as well as the corresponding Club4/Club4 and Not-Club4/Not-Club4 pairs. The within-category variability was analyzed using *same/different* judgments from the following data: Club1_1-2 (Club1/Club1 score and Club1/Club2 score), Club2_2-3 (Club2/Club2 score and Club2/Club3 score), Club3_3-4 (Club3/Club3 score and Club3/Club4 score), Not-Club3_3-4 (Not-Club3/Not-Club3 score and Not-Club3/Not-Club4 score), Not-Club2_2-3 (Not-Club2/Not-Club2 score and Not-Club2/Not-Club3 score), Not-Club1_1-2 (Not-Club1/Not-Club1 score and Not-Club1/Not-Club2 score).

To evaluate the perceptual component of the participants' responses, we averaged the mean percentage of correct decisions across the two trial types (*same* and *different*) within each of these six levels of the face levels factor. So, for instance, accuracy at Club1_1-2 was an average across the mean percentage of correct *same*




Figure 2. Mean categorization scores (percentages of correct answers) and standard errors as a function of face and block.

decisions for Club1/Club1 and correct *different* decisions for Club1/Club2 trials. The six accuracy scores obtained within each of the three discrimination blocks were examined as within-participants factors in a MANOVA. A main effect of block approached the conventional level of significance [$F(2,538) = 2.87, p = .059$], with accuracy scores increasing across blocks. A main effect of face level was also obtained [$F(5,1345) = 6.96, p < .001$], with accuracy scores increasing as faces approached category boundaries [the quadratic trend is $F(1,269) = 21.15, p < .001$]. Importantly, additional analyses conducted on a dichotomous club factor revealed that accuracy scores did not differ as a function of club [$F(1,269) = 0.01, n.s.$].

Independent of perceptual discrimination, there may be a decisional component that contributes to *same/different* judgments. To evaluate the decisional component of the participants' responses, we averaged across the mean percentage of *same* responses across the two trial types (*same* and *different*) within each of the six levels of the face level factor—for instance, the percentage *same* score at Club1_1-2 averaged across the mean percentage of correct *same* decisions on Club1/Club1 trials and the mean percentage of incorrect *same* decisions on Club1/Club2 trials. The six percentage *same* scores obtained within each of the three discrimination blocks were examined as within-participants factors in a MANOVA. A main effect of face level emerged [$F(5,1345) = 55.82, p < .001$], with percentage *same* scores decreasing as faces approached category boundaries [the quadratic trend is $F(1,269) = 198.54, p < .001$]. Additional analyses conducted on a dichotomous club factor also revealed that the participants were more likely to call a pair of faces the same when these faces pertained to the club category ($M = .622, SD = .10$) than when they pertained to the not-club category ($M =$

$.603, SD = .096$) [$F(1,269) = 11.60, p < .001$]. This effect emerged independently of the actual similarity of these faces (otherwise, a main effect of club would have been obtained on the accuracy scores examined above) and independently of the block factor (otherwise, a club \times block interaction would have been obtained).

The mean percentage of *same* responses and the mean perceptual discrimination scores obtained across the various levels of the face level factor, when collapsing across blocks, are reported in Figure 3. As can be seen from this figure, the percentage of *same* responses varied positively as a function of both club membership and face extremity. In contrast, perceptual discrimination scores varied as a function of face extremity only. The impact of face extremity is consistent with results in the literature on categorical perception (e.g., Harnad, 1987). Categorical perception has classically been defined as a better perceptual discrimination for stimuli lying closer to the categorical boundaries, and this effect was recently reported for face stimuli (e.g., Levin & Beale, 2000). More important to our present research interests, the impact of the club factor on the percentage of *same* responses supports our prediction that subjective similarity would be enhanced for the club faces.

Discussion

The results from Experiment 1 are informative in several respects. First, Experiment 1 replicated prior research on category asymmetry, even though the prototypical category members were never presented in this experiment. Like Goldstone and colleagues (Goldstone et al., 2003), we found a relative caricature advantage when the not club members were categorized, as compared with when the club members were categorized. That is, the caricature was more accurately categorized than the prototype to a larger

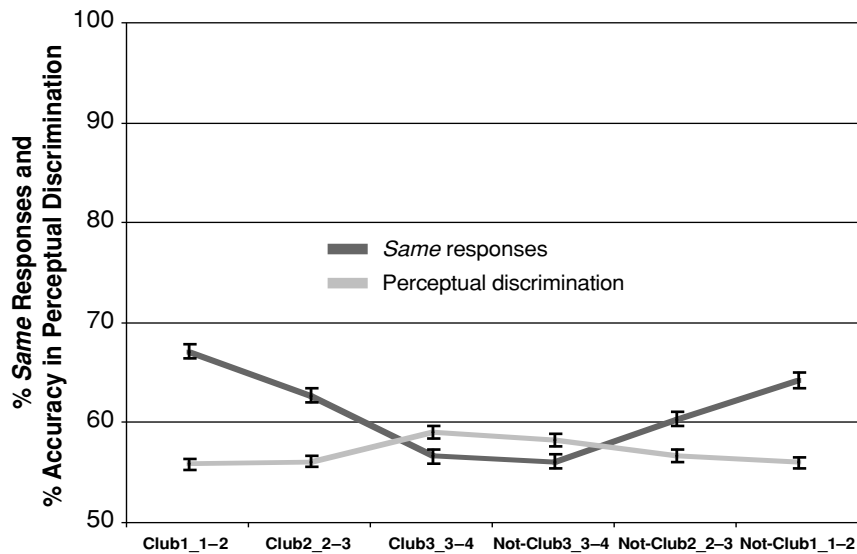


Figure 3. Mean percentages of *same* answers and mean percentages of accuracy, and associated standard errors, in perceptual discrimination as a function of face level.

extent for not club members than for club members. This caricature advantage for the not-club category, however, diminished over time (although it was reliably present in all three blocks). In Goldstone's terms, this finding would suggest that as training progresses, perceivers move from an isolated (and prototypical) representation of club members, paired with an interrelated (and caricatural) representation for not club members, toward a representation that is more isolated (and prototypical) for both categories.

The categorization results also supported the prediction that the not club members would be assimilated into the club category. Specifically, perceivers were more likely to misclassify nonmembers as belonging to the club than they were to misclassify club members as not belonging to the club. This effect is consistent with the idea that reference and nonreference exemplars display a feature asymmetry. As was explained in the introduction, one possibility is that the club members become associated with a shared *club* feature but that the nonmembers do not become associated with as salient a common feature. Because the difference between not possessing a feature and partially possessing it is larger than the difference between partially possessing a feature and fully possessing it, assimilation toward the club category may have occurred. Interestingly, this effect was found to decrease as the participants learned to more accurately categorize the exemplars. Category learning thus progressively decreased the impact of feature asymmetry on categorization.

Experiment 1 also provided a novel test of whether the participants' judgments of within-category similarity were asymmetric for reference and nonreference categories. The participants' ability to discriminate between two category members did not differ for reference, as compared with nonreference, categories. Nonetheless, an asymmetry was revealed in that the participants were more willing overall to claim that two reference category members were the same than they were to claim that two nonreference category members were the same. Together, these results suggest that the depth of encoding for the reference and the nonreference exemplars was equivalent but that the participants were more willing to report similarity for the reference exemplars. Again, this effect seems consistent with the hypothesis of a feature asymmetry for reference and nonreference members. According to the classic work by Tversky (1977) on feature asymmetry, similarity increases with the addition of a common feature.

In Experiment 2, we more directly tested this hypothesis of a feature asymmetry in the representation of reference and nonreference categories. In Experiment 2, the participants again learned about the categories via club/not-club categorization training, just as they did in Experiment 1. In Experiment 2, however, the category-learning task was followed by a visual search task. As has already been mentioned, performance in visual search tasks is known to be better when the searched-for item has an added feature that is not present in the distractor items than when the searched-for item exhibits the absence of a feature that is present in the distractor items (e.g., Treisman & Gormi-

can, 1988). Thus, for example, searching for a R in a field of Ps is faster than searching for an P in a field of Rs. This is such a replicable effect that the feature search task has also been employed to provide evidence that one set of items shares a feature that the other set of items does not (for an example with face stimuli, see Levin, 2000).

In the visual search task in Experiment 2, we examined the participants' ability to correctly determine the presence or absence of reference versus nonreference faces presented in a field of previously unseen faces. The use of a visual search task thus allowed for a straightforward test of the feature asymmetry hypothesis: If reference exemplars are defined by a feature that is lacking in the nonreference exemplars, the participants should be better (and possibly faster) at detecting reference than at detecting nonreference faces in a background of novel distractor faces.

To equate exposure to reference and nonreference exemplars prior to the visual search task, all the faces were shown for a fixed time during the category-learning task. Incidentally, this modification allowed us to examine whether representational asymmetries for the reference and nonreference categories would survive a tight control of exposure time.

EXPERIMENT 2

Method

Participants. One hundred twenty-five undergraduate students from Indiana University served as participants in order to fulfill a course requirement. The students were randomly assigned into the two labeling conditions.

Procedure. This experiment was divided into a category-learning task and a visual search task. For the category-learning task, the stimuli were identical to those used in Experiment 1, and the procedure was the same, except for the differences noted here. There were 216 trials, consisting of 27 repetitions of each of the eight faces. Each face to be categorized was shown for 1,300 msec; then the display was erased, and the participants were asked to categorize the face as either belonging to the club or not. The feedback timing was identical to that in Experiment 1. During the feature search phase, the participants saw 128 trials, consisting of 16 repetitions of each of the faces in Figure 1. The 16 repetitions were evenly divided into randomized *absent* and *present* trials. On present trials, one of the faces from Figure 1 was selected as a target. It was shown to the participants for 2 sec. Then a display with seven faces was presented, and the target was included as one of the faces. Absent trials followed the same procedure, except that the target face was not included among the seven faces. The nontarget distractors were not selected from the remaining faces shown in Figure 1 but, rather, were chosen from a set of 16 additional bald heads. The distractors were borrowed from Kayser (1997) and were selected to be approximately equally similar to the endpoint faces in Figure 1. Similarity was quantitatively based on our similarity assessments, using a technique described by Goldstone (1994a). The participants pressed the Y key to indicate presence of the target face and the N key to indicate absence. The computer gave the participants trial-by-trial feedback by presenting either a ✓ or an X for correct and incorrect responses, respectively. After 1.5 sec, the screen was erased. The blank interval between trials was 1 sec. For each search display, each of the faces was photographed against a black background. The seven faces were displayed in equal intervals around a circle. An example of a search display is shown in Figure 4. The target face, when present, was equally likely to appear in any one of the locations. Each of the faces

in a search display was 4×3.5 cm. This radius of the entire circle of faces was 15.5 cm. The average viewing distance was 46 cm.

Results

As in Experiment 1, we removed from the analyses those participants ($n = 13$) who had not achieved 70% correct categorizations at the end of the third and last categorization block. In all the analyses, we averaged across the multiple observations collected for a same factorial event (e.g., when considering how the participants categorized Club2 in Block 3, we averaged the data obtained for the nine presentations of Club2 in Block 3).

Category-learning task. We first ran a MANOVA on the categorization accuracy scores, with the eight levels of the face factor within the three blocks as within-participants factors. Main effects for face [$F(7,777) = 130.77, p < .001$] and for block [$F(2,222) = 238.98, p < .001$] were obtained. A face \times block interaction was also found [$F(14,1554) = 6.09, p < .001$]. Additional analyses conducted on a dichotomous club factor clarified the meaning of the latter effects: Categorization was more accurate for club faces ($M = 0.829, SD = 0.075$) than for not-club faces ($M = 0.798, SD = 0.094$) [$F(1,111) = 15.08, p < .001$], and this effect decreased across blocks [$F(2,222) = 5.93, p < .005$].

Figure 5 shows the mean categorization accuracy scores across the various levels of the face factor for the first, second, and third training blocks. As can be seen, accuracy scores were, on average, higher for the club faces. This difference decreased over blocks. Thus, consistent with the results of Experiment 1, assimilation toward the reference category was obtained, and this effect decreased with category learning.




Figure 4. A sample display of the feature search task in Experiment 2. Note: This is a target-present trial; Face 1 is presented at the left end of the display.

We also examined whether caricature effects could be replicated in the context of the present experiment, which involved no prototype and kept exposure times to the faces constant. The scoring for the caricature advantage was the same as that in Experiment 1. As was the case in Experiment 1, this stronger caricature advantage for the not-club faces, relative to the club faces, decreased across the levels of the block factor [$F(2,222) = 6.21, p < .003$] but remained significant within each of the three blocks (M for Block 1 = .11, $SD = .3$; M for Block 2 = .07, $SD = .23$; M for Block 3 = .02, $SD = .08$; all $ps < .008$).

Visual search task. We considered separately the percentages of correct decisions for the face-present and face-absent trials as a function of face. We first considered the face-present trials. A main effect of face was found [$F(7,777) = 2.90, p < .006$]. Further analyses conducted on a dichotomous club factor confirmed our hypothesis for better detection of the club faces ($M = .861, SD = .096$) than of the not-club faces ($M = .834, SD = .104$) [$F(1,111) = 10.44, p < .002$]. Interestingly, although not related to the present research interests, a quadratic trend was also obtained on these trials [$F(1,111) = 4.01, p < .05$], with better detection scores, on average, for stimuli lying closer to the category boundaries. For the face-absent trials, no effect of the face factor was obtained [$F(7,777) = 1.7, n.s.$].

In summary, the participants reported more accurately on the presence of reference faces than on the presence of nonreference faces, but they did not differ in the accuracy with which they reported the absence of reference versus nonreference faces. The mean percentages of correct decisions for the face-present and face-absent trials across the eight levels of the face factor are reported in Figure 6, which offers a finer-grained illustration for the aforementioned effects.

We also analyzed the participants' response latencies as a function of face. For each of the eight faces separately, and for the target-absent and target-present trials separately, we removed response latencies that were associated with incorrect answers and response latencies that were three SDs above or below the mean response time for that face. These response time analyses failed to produce any significant effect for the club factor.

Discussion

The categorization results obtained in Experiment 1 were replicated in Experiment 2, even though the participants' exposure to the faces was held constant across the various face presentations in this experiment. Thus, it cannot be argued that representational asymmetries (i.e., caricature effects, assimilation to the reference category) emerged because of a difference in exposure times for reference and nonreference faces. One may also note here that representational asymmetries were unlikely to result from a deeper encoding of the reference faces, since no reference effect was obtained on the perceptual component of the perceptual discrimination task in Experiment 1.

Beyond this successful replication for representational asymmetries in these challenging conditions, the visual

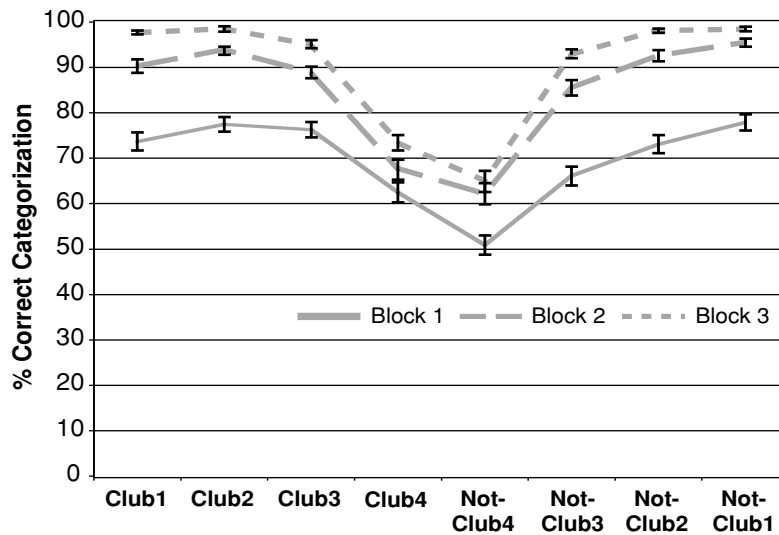


Figure 5. Mean categorization scores (percentages of correct answers) and standard errors as a function of face and block.

search findings brings further support to our hypothesis that there is a feature asymmetry in the representation of reference and nonreference categories. We posited that reference exemplars share a feature that is lacking for the nonreference exemplars. The finding of better detection of the reference exemplars, relative to nonreference ones, is consistent with this hypothesis. Importantly, this effect could not be attributed to a general tendency for reporting the presence of reference faces. If this had been the case, lower accuracy scores would have been obtained for the reference faces on the face-absent trials. This clearly was not the case (see Figure 6).

Although not at the focus of the present contribution, a quadratic trend was found on the face factor, with relatively

better performances, on average, for faces lying at moderate values of the continuum than for those lying at extreme values. This effect, which was obtained in an experimental setting that offered a tight control for face exposure, seems to provide original support for the categorical perception hypothesis (e.g., Harnad, 1987). Specifically, stimuli lying closer to the category boundaries may benefit from a perceptual discrimination advantage (Experiment 1), a higher decision criterion for responding *same* (Experiment 1), and a detection advantage (Experiment 2). One possibility for the latter advantage, however, is that boundary stimuli, because of enhanced classification uncertainties, were more deeply encoded in the category-learning task, resulting in better detection subsequently.

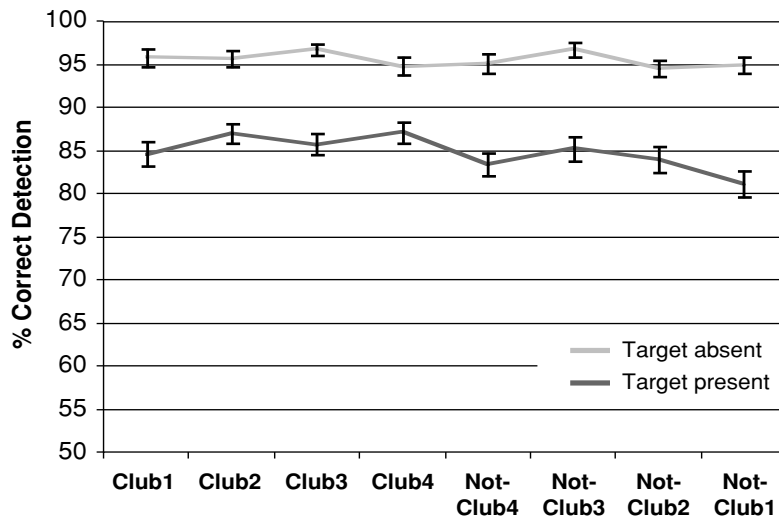


Figure 6. Mean percentages of correct detections and standard errors as a function of face and trial type.

GENERAL DISCUSSION

The research presented here suggests that reference-based category asymmetry is much broader than has previously been envisioned. Not only does category reference produce asymmetry in categorization decisions (Experiments 1 and 2), but also it produces asymmetry in judgments of within-category similarity (Experiment 1) and asymmetry in attention to the presence of reference versus nonreference category members (Experiment 2). As in prior work (Goldstone et al., 2003), Experiments 1 and 2 replicated the finding that the advantage for categorizing the caricature versus the prototype was stronger in the nonreference than in the reference category. Categorization asymmetry was also indicated by stronger assimilation of the nonreference exemplars into the reference category than vice versa. These categorization asymmetries occur even when no reference prototype is actually presented to the participants, and they survive under controlled exposure time. The present work also demonstrated that reference-based categorization asymmetries decrease as category learning progresses. A further notable contribution of the present work is that category reference does not result in an asymmetry in perceptual discrimination but does produce a decisional asymmetry: Although within-category perceptual discrimination accuracy did not vary as a function of reference versus nonreference status, there is, nonetheless, a higher probability of judging two reference, as compared with nonreference, category exemplars to be the same. In another novel extension of reference-based category asymmetry, we found that category reference facilitates the detection of reference exemplars, relative to nonreference exemplars. Finally, we demonstrated that detection is facilitated for stimuli lying at the category boundaries, highlighting a detection advantage component of categorical perception. Overall, this set of findings seems consistent with the hypothesis of a feature asymmetry in the representation of reference and nonreference categories (Levin, 2000). As was discussed in the introduction, the existence of a feature advantage for reference exemplars would lead to the three major results examined and obtained in this contribution: assimilation of nonreference exemplars toward the reference category, enhanced judgments of similarity for reference than for nonreference exemplars, and better detection of reference than of nonreference exemplars.

More generally, the present findings also confirm that nonreference categories are organized *in relation to* reference categories, whereas reference categories are more isolated from their conceptual neighbors. As a matter of fact, the assimilation effect obtained in Experiments 1 and 2 appears quite consistent with this notion. It seems unreasonable, however, to argue for a definitive answer as to whether category relatedness results in the organization of nonreference exemplars away or toward the reference category. Recall that Goldstone and colleagues proposed that reference categories are organized around a referent

category prototype, whereas nonreference categories are organized around a nonreferent caricature. The hypothesized result of this difference in representational organization was a relative advantage for the categorizing of the caricature for the nonreference category, as compared with the reference category.

Looking at the pattern of results, however, it is also possible that all exemplars are organized around the reference prototype (and only that referent). In this conceptualization, category learning progresses through a comparison of each exemplar encountered with the reference category prototype (i.e., the club prototype). For reference category exemplars, categorization accuracy decreases with distance from the reference prototype, producing maximal accuracy at the category prototype. For nonreference category exemplars, categorization accuracy increases for exemplars that are furthest from the reference category prototype. That is, it is easier to exclude an exemplar that is very dissimilar from the club prototype from the club than it is to exclude an exemplar that is more similar to the club prototype from the club. This process would produce both assimilation toward the reference category and maximal categorization accuracy at the caricature face for the nonreference category.

This same idea of comparison of all exemplars with the reference category prototype can account for participants' *same/different* judgments as well. Assuming that all the exemplars were compared with the reference prototype during the early stages of the categorization process, it makes sense that the participants formed less differentiated representations for the reference category than for the nonreference category. This is because the distances between the reference exemplars and the reference prototype have a much smaller range (i.e., from 0.5 to 1.5 in the present study) than do the distances between the nonreference exemplars and the reference prototype (from 2.5 to 5.5 in the present study). In other words, the constant reference to the reference prototype may have resulted in the perception of smaller intracategorical variations for the reference than for the nonreference category. This, in turn, may have enhanced the probability for *same* decisions for the reference exemplars.

Finally, the finding that the participants more readily detected the presence of reference category members than that of nonreference category members might also be accounted for by the comparison of all the exemplars with the reference category prototype during category learning. In this conceptualization, the prototypical reference category member is accessed and referred to on every categorization trial. Thus, additional *experience* with the reference category prototype might make exemplars similar to this well-learned reference prototype more detectable than exemplars that are less similar to the reference prototype. Although the pattern of results across the eight individual faces is not entirely consistent with this idea, the enhanced detection of faces at the category boundary (a categorical perception effect) may be occluding a

more detailed pattern of best detection at the reference category prototype and worst detection at the nonreference category caricature.

Both the *referent prototype only* and the *referent prototype + nonreferent caricature* accounts of the data thus lead to the same set of predictions with regard to the asymmetries emerging in the representation, treatment, and attention devoted to reference and nonreference categories. Clearly, our experiments do not provide, and were certainly not aimed at providing, a test of this idea of the organization of all exemplars around the reference category prototype, as compared with the combination of more established claims of (1) organization of the reference category around its prototype and of the nonreference category around its caricature, (2) enhanced similarity of reference category members due to a common club feature, and (3) better detection of reference category members because of the added club feature. Still, the referent prototype only account seems, in some ways, more parsimonious and provides one additional advantage. With this account, we do not have to explain why a feature adds similarity at the decisional stage but not at the perceptual discrimination stage. This account remains speculative, however, and in the absence of complementary evidence, the feature asymmetry hypothesis seems better suited to account for the data obtained for the perceptual discrimination and visual search tasks.

Before concluding, one should note a possible limitation of the present research: The categories being tested here may have been perceived as variants of two different faces, not the sets of different exemplars typically associated with classification experiments. This might matter if reference effects are limited to situations in which participants are led to focus on subtle differences in shape that represent variants of one thing. Although this limitation is common in the context of categorization work that relies on morphed stimuli, further research should be done to examine whether the effects obtained here can be generalized to categories involving stimuli morphed across multiple source images (see Corneille & Judd, 1999, for a morphing procedure allowing this kind of test).

CONCLUSION

We found that a simple category-labeling manipulation affects not only categorization performance, but also performance on tasks normally thought to be based on perceptual and attentional processes, rather than on the high-level cognitive processes associated with classification. The minimal nature of our category-labeling manipulations seems impressive. The asymmetry effects obtained here were not caused by minority status, exemplar frequency, a participant's own perspective regarding in-groups and out-groups, or even familiarity or exposure time. Rather, the labels alone, and the reference status that they conveyed, sufficed to induce asymmetries in both classification and perceptual performance.

REFERENCES

- AGOSTINELLI, G., SHERMAN, S. J., FAZIO, R. H., & HEARST, E. S. (1986). Detecting and identifying change: Additions versus deletions. *Journal of Experimental Psychology: Human Perception & Performance*, *12*, 445-454.
- BOTHWELL, R. K., BRIGHAM, J. C., & MALPASS, R. S. (1989). Cross-racial identification. *Personality & Social Psychology Bulletin*, *15*, 19-25.
- CORNEILLE, O., HUART, J., BECQUART, E., & BRÉDART, S. (2004). When memory shifts towards more typical category exemplars: Accentuation effects in the recollection of ethnically ambiguous faces. *Journal of Personality & Social Psychology*, *86*, 236-250.
- CORNEILLE, O., & JUDD, C. M. (1999). Accentuation and sensitization effects in the categorization of multi-faceted stimuli. *Journal of Personality & Social Psychology*, *77*, 927-941.
- CORNEILLE, O., KLEIN, O., LAMBERT, S., & JUDD, C. M. (2002). On the role of familiarity with units of measurement in producing categorical accentuation: Tajfel and Wilkes (1963) revisited and replicated. *Psychological Science*, *4*, 380-383.
- EISER, J. R., & VAN DER PLIGT, J. (1984). Accentuation theory, polarization and the judgement of attitude statement. In J. R. Eiser (Ed.), *Attitudinal judgement* (pp. 43-63). New York: Springer.
- GOLDSTONE, R. [L.] (1994a). An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, & Computers*, *26*, 381-386.
- GOLDSTONE, R. [L.] (1994b). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178-200.
- GOLDSTONE, R. L. (1996). Isolated and interrelated concepts. *Memory & Cognition*, *24*, 608-628.
- GOLDSTONE, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585-612.
- GOLDSTONE, R. L., & STEVYERS, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, *130*, 116-139.
- GOLDSTONE, R. L., STEYVERS, M., & ROGOSKY, B. J. (2003). Conceptual interrelatedness and caricatures. *Memory & Cognition*, *31*, 169-180.
- HARNAD, S. (ED.) (1987). *Categorical perception: The groundwork of cognition*. New York: Cambridge University Press.
- HUART, J., CORNEILLE, O., & BECQUART, E. (2005). Face-based categorization, context-based categorization, and distortions in the recollection of gender ambiguous faces. *Journal of Experimental Social Psychology*, *41*, 598-608.
- KAYSER, A. (1997). *Heads*. New York: Abbeville.
- KRUEGER, J., & CLEMENT, R. W. (1994). Memory-based judgments about multiple categories: A revision and extension of Tajfel's accentuation theory. *Journal of Personality & Social Psychology*, *67*, 35-47.
- KRUEGER, J., & ROTHBART, M. (1990). Contrast and accentuation effects in category learning. *Journal of Personality & Social Psychology*, *59*, 651-663.
- LEVIN, D. T. (2000). Race as visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross race recognition deficit. *Journal of Experimental Psychology: General*, *129*, 559-574.
- LEVIN, D. T., & BEALE, J. M. (2000). Categorical perception occurs in newly learned faces, other-race faces, and inverted faces. *Perception & Psychophysics*, *62*, 386-401.
- LIVINGSTON, K. R., ANDREWS, J. K., & HARNAD, S. (1998). Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *24*, 732-753.
- MARKMAN, A. B., & ROSS, B. H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592-613.
- MULLEN, B., & HU, L. (1989). Perceptions of ingroup and outgroup variability: A meta-analytic integration. *Basic & Applied Social Psychology*, *10*, 233-252.
- PALMERI, T. J., WONG, A. C.-N., & GAUTHIER, I. (2004). Computational approaches to the development of perceptual expertise. *Trends in Cognitive Sciences*, *8*, 378-386.

- PYLYSHYN, Z. (1999). Is vision continuous with cognition? The case of impenetrability of visual perception. *Behavior & Brain Sciences*, **22**, 341-423.
- QUINLAN, P. T. (2003). Visual feature integration theory: Past, present, and future. *Psychological Bulletin*, **129**, 643-673.
- READ, S. J., & URADA, D. I. (2003). A neural network simulation of the outgroup homogeneity effect. *Personality & Social Psychology Review*, **7**, 146-169.
- SHIFFRIN, R. M., & LIGHTFOOT, N. (1997). Perceptual learning of alphanumeric-like characters. In R. L. Goldstone, P. G. Schyns, & D. L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 36, pp. 45-82). San Diego: Academic Press.
- STEYVERS, M. (1999). Morphing techniques for manipulating face images. *Behavior Research Methods, Instruments, & Computers*, **31**, 359-369.
- TAJFEL, H., & WILKES, A. L. (1963). Classification and quantitative judgment. *British Journal of Psychology*, **54**, 101-114.
- TANAKA, J. W., GILES, M., KREMEN, S., & SIMON, V. (1998). Mapping attractor fields in face space: The atypicality bias in face recognition. *Cognition*, **68**, 199-220.
- TAYLOR, S. E., FISKE, S. T., ETCOFF, N. L., & RUDERMAN, A. (1978). Solo status as a psychological variable: The power of being distinctive. *Journal of Personality & Social Psychology*, **36**, 778-793.
- TREISMAN, A., & GELADE, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, **12**, 97-136.
- TREISMAN, A., & GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, **95**, 15-48.
- TVERSKY, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.
- WANG, Q., CAVANAGH, P., & GREEN, M. (1994). Familiarity and pop-out in visual search. *Perception & Psychophysics*, **56**, 495-500.

(Manuscript received January 21, 2005 ;
revision accepted for publication April 19, 2005.)