Asymmetries in Categorization, Perceptual Discrimination, and Visual Search for Reference and Non-reference Exemplars

Olivier Corneille 1

Robert L. Goldstone 2

Sarah Queller 2

&

Timothy Potter 1

1 Catholic University of Louvain, Louvain-la-Neuve

2. Indiana University, Bloomington, Indiana

Words count: 7584

Keywords: Categorization, Category learning, Face perception

Running title: Reference effects

Authors' note:

Olivier Corneille and Timothy Potter, Catholic University of Louvain, Louvain-la-Neuve, Belgium. Robert Goldstone and Sarah Queller, Indiana University at Bloomington, Indiana. We would like to thank Stephanie Demoulin, Peter Freytag, Chick Judd, and Craig McGarty, for their comments on an earlier version of this manuscript. This research was funded in part by NSF grant 0125287 to the second author. Correspondence concerning this manuscript may be addressed to Olivier 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

Abstract

Two studies examined the representation, treatment, and attention, devoted to the members of reference (i.e., Club members) and non-reference (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 non-reference categories which, however, decreased as expertise and familiarity with the categories increased (Experiment 1 and Experiment 2). Participants also more readily judged two reference than two non-reference exemplars as being the same (Experiment 1), and were better at detecting reference than non-reference 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 non-reference 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, compared to non-reference, category.

Asymmetries in Categorization, Perceptual Discrimination, and Visual Search for Reference and Non-reference Exemplars

Introduction

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). Past research also examined this effect in judgments of multifaceted stimuli (e.g., Corneille & Judd, 1999; Goldstone, 1996; Livingston, Andrews, & Harnad, 1999) and addressed the consequences of this bias for memory (e.g., Corneille, Huart, Becquart, & Bredart, 2004; Huart, Corneille, & Becquart, in press ; Krueger & Clement, 1994; Taylor, Fiske, Etcoff, & Ruderman, 1978).

In this work, the categories involved were generally 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 merely defined 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 effets for category learning. These authors proposed that reference categories are likely to be organized around a prototype, whereas non-reference 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 about non-reference than reference categories.

In Goldstone and colleagues' work, the category exemplars consisted of a series of faces that were located along the left-end or 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 versus Not-Club members) or learning order (participants first learned about Category A members, and only then learned about Category B members). Participants were asked to categorize exemplars into one of two categories (e.g., Club versus 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 better categorize extreme than typical category exemplars was larger for the non-reference (e.g., Not-Club members) than for the reference (e.g., Club members) category. In other words, a larger caricature advantage was found in categorizing the members of the non-reference group.

Whereas this work 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 left unanswered, two of which are examined in this contribution. The first issue concerns people's ability to discriminate between examplars from reference and non-reference 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 non-reference category? Independent of perceptual discrimination, are people more or less likely to make decisions indicative of beliefs about stronger similarities among reference than non-reference exemplars? The second issue concerns people's ability to detect exemplars from reference and non-reference categories: Are people better at detecting the presence of reference than non-reference exemplars among a set of distractors?

There are conflicting accounts for whether the members of reference (Club) or non-reference (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 identification of faces from one's own race than 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 may be expected to be less distinct because they all share a common, salient category membership. Levin (2000) describes evidence that faces belonging to salient categories are more similar to one another than faces belonging to backgrounded categories. If reference categories are more salient than non-reference categories, then faces sharing a reference category membership might be expected to be judged more similar than faces sharing a non-reference category membership. A major goal of the current experiments is to decide between these accounts.

Given the many influences of training on featural encodings (see Palmeri, Wong, & Gauthier, 2004 for a recent review), it is possible that categorization training induces asymmetric features for reference and non-reference categories. More specifically, if people focus mostly on establishing membership in the reference category, then members of the reference category may develop additional psychological features compared to members of the non-reference 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 non-reference categories would tend to be placed in the reference category. This is because the midway item would partially possess 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 non-distinctive item is judged to be more similar to the distinctive item (Tanaka, Giles, Kremon, & Simon, 1998). Extending this idea to the present situation, participants should more likely mis-classify a non-reference exemplar similar to the midway item into the reference category than to mis-classify a reference exemplar similar to the midway item into the non-reference 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 non-reference category. 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 their 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 (1977) 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 non-reference 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 a non-reference category. However, there should also be more incorrect "same" judgments when two different items are presented that belong to a reference as compared to non-reference category.

Third, asymmetries in a search task may be predicted. If training on a reference versus non-reference categorization task causes the reference category items to acquire additional features relative to the non-reference category, then people should be better at detecting reference than non-reference category items. This logic parallels Treisman and Gormican's (1988) argument that the letter "R" can be detected among "P"s more efficiently than a "P" can be detected among "R"s 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 than absent feature (Agostinelli, Sherman, Fazio, & Hearst, 1986; Quinlan, 2003). One might predict that feature-based asymmetries should only exist for hard-wired 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 studies

As in the original studies by Goldstone et al. (2003), we used multifaceted stimuli: 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 studies we conducted, half of the participants saw the faces lying on the left side of the continuum referred to as Club faces and saw faces lying on the right side of the continuum 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 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 participants to progressively learn to correctly assign the faces into the Club and Not-Club categories. Unlike Goldstone et al. (2003), 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) are magnified or weakened as familiarity with the categories increases. In the subsequent Perceptual Discrimination task, participants were sequentially presented with pairs of faces. On a given trial, they either saw the same face presented twice or they saw two faces that were adjacent on the morph continuum. 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 participants' ability to discriminate between adjacent faces, and to examine participants' overall readiness to judge two faces as being the same.

In Experiment 2, participants completed a Category Learning task, followed by a Visual Search task. The Category Learning task was similar to that of Experiment 1, except that exposure time was held constant for all faces. In the Visual Search task, 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 participants would be at correctly identifying reference and non-reference exemplars against a background of novel distractor faces.

The current experiments are a major extension to previous results 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 current experiments extend 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 are traditionally considered to be perceptual and attentional tasks. An effect of category labeling on the simple task of deciding whether two objects are identical, or picking out an object from a set of distractors, might be surprising on accouts that draw a sharp boundary between perceptual and conceptual tasks (Pylyshyn, 1999). These kind of effects would, however, be consistent with a growing literature suggesting that perceptual representations can be influenced by experience, task demands, and learned categories (Corneille et al., 2002; Goldstone, 1998; Levin, 2000; Wang, Cavanagh, & Green, 1994)

Experiment 1

Method

Participants. Two hundred and 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 did not introduce conspicuous nonlinearities between physical and psychological scalings (Goldstone & Steyvers, 2001). The morph sequence of 8 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 gray scale values of corresponding pixels were blended to create the gray scale values of the resulting morph image. The faces on the left and right ends of Figure 1 are actual faces, and the 6 intermediate faces are blends of the two actual faces, with the proportion of the right-most face beginning at 0% for the left-most face, and shifting along the series in equal 14.29% increments.

----------------------------Insert Figure 1 about here ------------------------------

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 through 4 with Faces 2 and 3 straddling the prototype. Similarly, the right-end category consists of Faces 5 through 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 (one pixel = .034 cm), and measured 14.48 cm tall by 11.68 cm wide. Each face was photographed against a dark background and displayed on a white Apple Imac computer screen. The average viewing distance was 46 cm.

Procedure. The stimuli were divided into Club 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 later factor (i.e., Condition) is basically a counterbalancement and it will be ignored in the remainder of this manuscript (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), Club 3 (i.e., Face 3 in in condition 1 and Face 6 in condition 2), Club 4 (i.e., Face 4 in condition 1 and Face 5 in condition 2), Not-Club 4 (i.e., Face 5 in in condition 1 and Face 4 in condition 2), Not-Club 3 (i.e., Face 6 in condition 1 and Face 3 in condition 2), Not-Club 2 (i.e., Face 7 in condition 1 and Face 2 in condition 2), and Not-Club 1 (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: "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 seconds. The blank interval between trials was 1 second. 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 X 6 cm square in the center of the screen. Participants were given breaks every 80 trials. During these breaks, participants were informed of their accuracy and speed during the preceding block.

In the second phase of the experiment, the Perceptual Discrimination task, participants were instructed that they would see displays with two faces on the screen. Their task was to decide if the faces were exactly identical or differed in any way at all. Participants were warned that all of the faces would be highly similar to one another. Participants pressed the “s” key to indicate a “same” response and “d” to indicate a “different” response. The computer gave 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 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 participants responded. Immediately after pressing "s" or "d", feedback was provided, and after 1.5 seconds, the screen was erased. The blank interval between trials was 1 second.

Results

We divided 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 analyses, we averaged across the multiple observations collected for the same configuration (e.g., when considering how participants categorized Club 1 in Block 1, we averaged across data obtained on the ten presentations of Club 1 in Block 1).

Category Learning task

We first ran a MANOVA on the categorization accuracy scores for the eight levels of the Face factor with 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 by Block interaction was also found, $F(14,3766)=9.27$, $p<.001$. Additional analyses conducted on a dichotomous Club factor (i.e., Club versus Not-Club faces) clarified the meaning of the latter effects : categorization was more accurate for Club faces (M=0.85, SD=0.069) than for Not-Club faces (M=0.82, SD=0.083), $F(1,269)=66.08$, $p<.001$, 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 Block 1 and 2). Because lower accuracy scores reflect assignments to the alternative category, this finding may be re-interpreted 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.

----------------------------Insert Figure 2 about here ------------------------------

We also examined whether Goldstone et al. (2003)'s caricature advantage for Not-Club faces could be replicated in the context of the present study that involved no prototype, and whether this effect would prove sensitive to the Block factor. Because no actual prototype was presented in the present study, 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 accuracy for Not-Club Face 1 minus the accuracy for the typical Not-Club face and subtracting from that accuracy the result of Club Face 1 accuracy minus the typical Club face accuracy. This stronger caricature advantage for the Not-Club faces as compared to 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*block1=0.066 ; *SD*block1=0.238 ; *M*block2=0.017 ; *SD*block2=0.127 ; *M*block3=0.017 ; *SD*block3=0.098; all *p*s<.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 participants' responses, we averaged the mean percentage of correct decisions across the two trial types (same and different) within each of these 6 levels of the FaceLevels factor. So, for instance, accuracy at Club1_1-2 was an average across the mean percentage of correct 'same' 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 conventional level of significance, $F(2,538)=2.87$, $p=.059$, with accuracy scores increasing across Blocks. A main effect of FaceLevel 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, $F(1,269)=0.01$, *ns*.

Independent of perceptual discrimination, there may be a decisional component that contributes to same/different judgments. To evaluate the decisional component of participants' responses, we averaged across the mean percentage of "same" responses across the two trial types (same and different) within each of the 6 levels of the FaceLevel factor. So, 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 FaceLevel emerged, $F(5,1345)=55.82$, $p<.001$, with percentage same scores decreasing as faces approched category boundaries (the quadratic trend is $F(1,269)=198.54$, $p<.001$). Additional analyses conducted on a dichotomous Club factor also revealed that participants were more likely to call a pair of faces the same when these faces pertained to the Club category ($M=0.622$, $SD=0.10$) than to the Not-Club category ($M=0.603$, $SD=0.096$), $F(1,269)=11.60$, $p<.001$. This effect emerged independent of the actual similarity of these faces (otherwise, a main effect of Club would have been obtained on the Accuracy scores examined above), and independent of the Block factor (otherwise, a Club by 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 FaceLevel factor, when collapsing across blocks, are reported in Figure 3. As can be seen on 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 the literature on categorical perception (e.g., Harnad, 1987). Categorical perception is classically 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 ‘same’ responses supports our prediction that subjective similarity is enhanced for the Club faces.

----------------------------Insert Figure 3 about here ------------------------------

Discussion

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 study. Like Goldstone and colleagues (Goldstone et al., 2003), we found a relative caricature advantage when categorizing the Not-Club members as compared to when categorizing the Club members. That is, the caricature was more accurately categorized than the prototype to a larger extent for Not-Club members than for Club members. This caricature advantage for the Not-Club category, however, diminished over time (although is 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 mis-classify non-members as belonging to the club than they were to mis-classify club members as not belonging to the club. This effect is consistent with the idea that reference and non-reference exemplars display a feature asymmetry. As explained in the introduction, one possibility is that the club members become associated with a shared "Club" feature but that the non-members 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 occured. Interestingly, this effect was found to decrease as 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 participants' judgments of within category similarity were asymmetric for reference and non-reference categories. Participants' ability to discriminate between two category members did not differ for reference as compared to non-reference categories. Nonetheless, an asymmetry was revealed in that participants were more willing overall to claim that two reference category members were the same than they were to claim that two non-reference category members were the same. Together, these results suggest that the depth of encoding for the reference and non-reference exemplars was equivalent but that 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 non-reference members. According to classic work by Tversky (1977) on feature asymmetry, similarity increases with the addition of a common feature.

Experiment 2 more directly tested this hypothesis of a feature asymmetry in the representation of reference and non-reference categories. In Experiment 2, 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 already 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 & Gormican, 1988). Thus, for example, searching for a R in a field of P's is faster than searching for an P in a field of R's. 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 of Experiment 2, we examined participants' ability to correctly determine the presence or absence of reference versus non-reference 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 non-reference exemplars, then participants should be better (and possibly faster) at detecting reference than non-reference faces in a background of novel distractor faces.

To equate exposure to reference and non-reference exemplars prior to the Visual Search task, all 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 non-reference categories would survive a tight control of exposure time.

Experiment 2

Method

Participants. One hundred and 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.

Procedures. 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 1300 msec, and then the display was erased, and participants were asked to categorize the face as either belonging to the club or not. The feedback timing was identical to that of Experiment 1. During the feature search phase, 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 participants for 2 seconds. Then, a display with 7 faces was presented and the target was included as one of the faces. Absent trials followed the same procedure except the target face was not included among the 7 faces. The non-target distractors were not selected from the remaining faces of 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 (1994b). Participants pressed the “Y” key to indicate presence of the target face, and the “N” key to indicate absence. The computer gave participants trial-by-trial feedback by presenting either a "✓" check or an "X" for correct and incorrect responses, respectively. After 1.5 seconds, the screen was erased. The blank interval between trials was 1 second. For each search display, each of the faces was photographed against a black background. The 7 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 cm X 3.5 cm. This radius of the entire circle of faces was 15.5 cm. The average viewing distance was 46 cm.

----------------------------Insert Figure 4 about here ------------------------------

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 analyses, we averaged across the multiple observations collected for a same factorial event (e.g.,when considering how 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 for 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 by 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, $F(2,222)=5.93$, $p<.005$.

Figures 5 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. Thus, consistent with Experiment 1, assimilation toward the reference category was obtained and this effect decreased with category learning.

----------------------------Insert Figure 5 ------------------------------

We also examined whether caricature effects could be replicated in the context of the present study that involved no prototype, and which kept constant exposure times to the faces. The scoring for the caricature advantage was the same as in Experiment 1. As it 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*block1=0.11 ; *SD*block1=0.3 ; *M*block2=0.07 ; *SD*block2=0.23 ; *M*block3=0.02 ; *SD*block3=0.08 ; all *p*s<.008).

Visual Search task

We considered separately the percentage 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 ($M=0.861$, $SD=.096$) than the Not-Club faces ($M=0.834$, $SD=.104$), $F(1,111)=10.44$, $p<.002$. 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, $F(7,777)=1.7$, ns.

In summary, participants more accurately reported on the presence of reference than non-reference faces but they did not differ in the accuracy with which they reported the absence of reference versus non-reference faces. The mean percentage 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.

----------------------------Insert Figure 6 about here ------------------------------

We also analyzed 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 participants' exposure to the faces was held constant across the various face presentations in this study. 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 non-reference faces. One may also note here that representational asymmetries were unlikely to result from a deeper encoding of the reference faces, as 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 Search findings brings further support to our hypothesis for a feature asymmetry in the representation of reference and non-reference categories. We posited that reference exemplars share a feature that is lacking for the non-reference exemplars. The finding of better detection of the reference relative to nonreference exemplars 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 than extreme values of the continuum. 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 criteria 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.

General discussion

The research presented here suggests that reference-based category asymmetry is much broader than previously envisioned. Not only does category reference produce asymmetry in categorization decisions (Experiment 1 and 2), but it also produces asymmetry in judgments of within-category similarity (Experiment 1) and asymmetry in attention to the presence of reference versus non-reference category members (Experiment 2). As in prior work (Goldstone et al., 2002), Experiment 1 and 2 replicated the finding that the advantage for categorizing the caricature versus the prototype was stronger in the nonreference than reference category. Categorization asymmetry was also indicated by stronger assimilation of the non-reference 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 non-reference status, there is, nonetheless, a higher probability of judging two reference, compared to non-reference, 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 relative to non-reference 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 non-reference categories (Levin, 2000). As 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 non-reference exemplars toward the reference category, enhanced judgments of similarity for reference than non-reference exemplars, and better detection of reference than non-reference exemplars.

More generally, the present findings also confirm that non-reference 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 Experiment 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 non-reference exemplars away or toward the reference category. Recall that Goldstone and colleagues proposed that reference categories are organized around a referent category prototype, whereas non-reference categories are organized around a non-referent caricature. The hypothesized result of this difference in representational organization was a relative advantage for categorizing the caricature for the non-reference category as compared to 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 by comparison of each exemplar encountered to 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 non-reference category exemplars, categorization accuracy increases for exemplars that are farthest 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 non-reference category.

This same idea of comparison of all exemplars to the reference category prototype can account for participants' same/different judgments as well. Assuming that all exemplars were compared to the reference prototype during the early stages of the categorization process, it makes sense that participants formed less differentiated representations for the reference category than the non-reference 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 the distances between the non-reference 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 intra-categorical variations for the reference than for the non-reference category. This may in turn have enhanced the probability for "same" decisions for the reference exemplars.

Finally, the finding that participants detected the presence of reference category members more readily than non-reference category members might also be accounted for by comparing all exemplars to 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 non-reference category caricature.

Both these 'referent prototype only' and the 'referent prototype + non-referent caricature' accounts of the data thus lead to the same set of predictions as to the asymmetries emerging in the representation, treatment, and attention devoted to reference and non-reference categories. Clearly, our studies do not provide, and were certainly not aimed at providing, a test of this idea of organization all exemplars around the reference category prototype as compared to the combination of more established claims of 1) organization of the reference category around its prototype and of the non-reference 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 ‘prototype referent 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 on the Perceptual Discrimination and Visual Search tasks.

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 the high-level cognitive processes associated with classification. The minimal nature of our category labeling manipulations seems impressive. The asymmetry effets obtained here were not caused by minority status, exemplar frequency, one's own perspective regarding in-groups and out-groups, or even familiarity or exposure time. Rather, the labels alone, and the reference status 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.

Bruner, J. S. (1957). On perceptual readiness. Psychological Review, 64, 123-152.

Corneille, O., & Judd, C. M. (1999). Accentuation and sensitization effects in the categorization of multi-faceted stimuli, *Journal of Personality and Social Psychology, 77*, 927-941.

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 and Social Psychology, 86*, 236-250.

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 verlag.

Furl, N. Phillips, P. J., & O'Toole, A. (2002). Face recognition algorithms and the other-race efffect : Computational mechanisms for a developmental contact hypothesis. <u>Cognitive Science</u>, <u>26</u>, 797-815.

Goldstone, R. (1994a). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General, 123*, 178-200.

Goldstone, R. L. (1994b). An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, & Computers, 26*, 381-386.

Goldstone, R. L. (1995). Effects of categorization on color perception. *Psychological Science, 6*, 298-304.

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. (1987). Introduction : psychophysical and cognitive aspects of categorical perception : A critical overview. In S. Harnad (Ed.), *Categorical perception: The groundwork of cognition*. New York : Cambridge university Press.

Huart, J., Corneille, O., & Becquart, E. (in press). Face-based Categorization, Context-based Categorization, and Distortions in the Recollection of Gender Ambiguous Faces. *Journal of Experimental Social Psychology*.

Kayser, A. (1997). *Heads*. New York: Abbeville Press.

Krueger, J., & Clement, R. W. (1994). Memory-based judgments about multiple categories: A revision and extension of Tajfel's accentuation theory. *Journal of Personality and Social Psychology; 67*, 35-47.

Krueger, J., & Rothbart, M. (1990). Contrast and accentuation effects in category learning. *Journal of Personality and 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 and 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, and 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 and 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.

Read, S. J., & Urada, D. I. (2003). A neural network simulation of the outgroup homogeneity effect. *Perconality and Social Psychology Review*, *7*, 146-169.

Quinlan, P. T. (2003). Visual feature integration theory: Past, present, and future. *Psychological Bulletin*, *129*, 643-673.

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, Volume* 36. San Diego: Academic Press. (pp. 45-82).

Steyvers, M. (1999). Morphing techniques for generating and 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 and 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.

Figures Caption :

Figure 1. In Experiment 1, a morph sequence of 8 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.

Figure 2 : Mean categorization scores (% correct answers) and Standard Errors, as a Function of Face and Block.

Figure 3 : Mean % of « Same » answers and Mean % Accuracy, and associated Standard Errors, in Perceptual discrimination as a Function of FaceLevel.

Figure 4. A sample display of the feature search task from Experiment 2 (Note : this is a Target-present trial : Face 1 is presented at the left-end of the display)

Figure 5 : Mean categorization scores (% correct answers) and Standard Errors as a Function of Face and Block.

Figure 6 : Mean % correct detections and Standard Errors as a Function of Face and Trial Type.

Figure 1 :

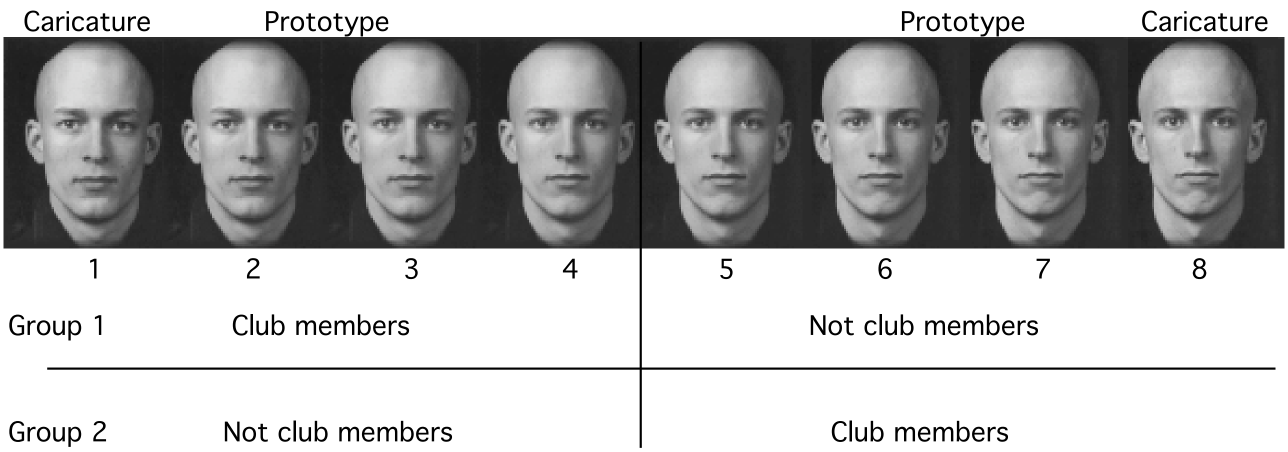
Caricature
Prototype
Prototype
Caricature
1
2
3
4
5
6
7
8
Group 1
Club members
Not club members
Group 2
Not club members
Club members

Figure 2 :

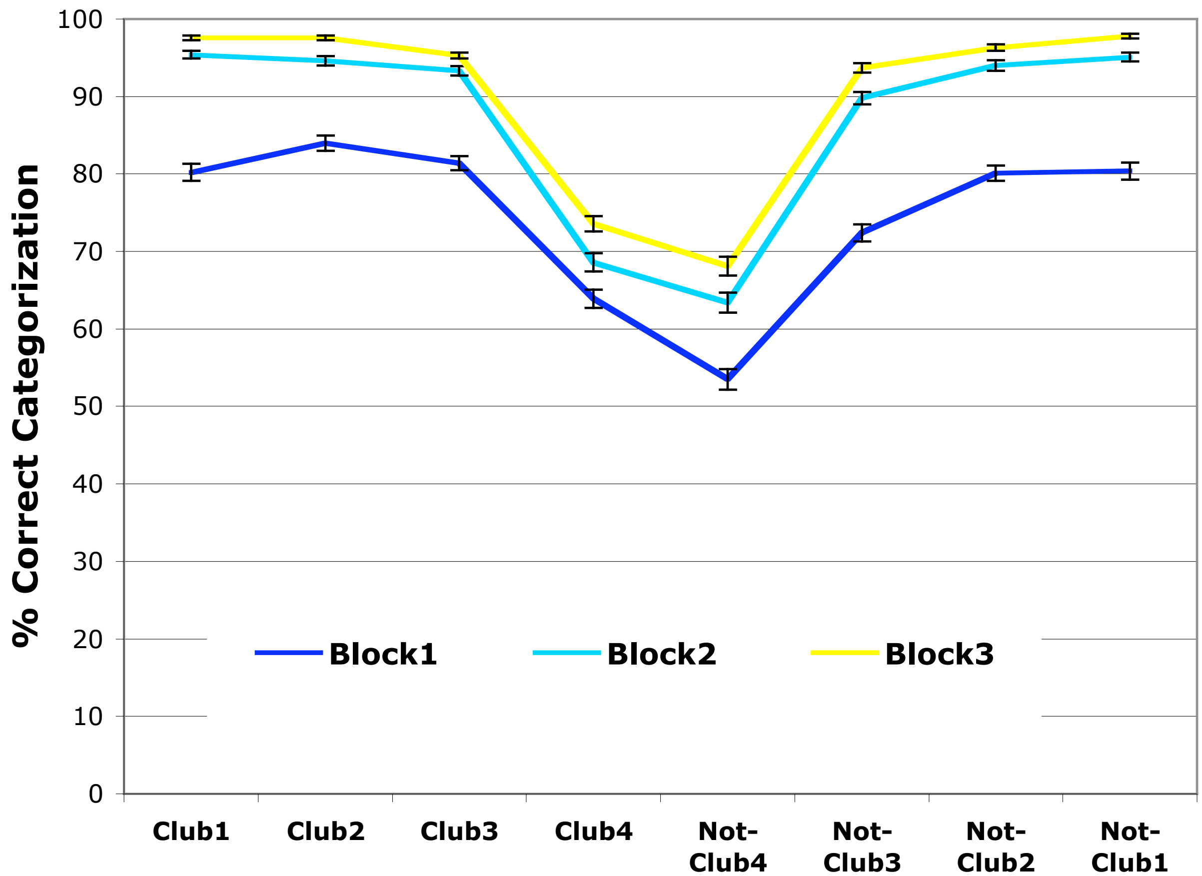
% Correct Categorization
100
90
80
70
60
50
40
30
20
10
0
Block1
Block2
Block3
Club1
Club2
Club3
Club4
Not-Club4
Not-Club3
Not-Club2
Not-Club1

Figure 3 :

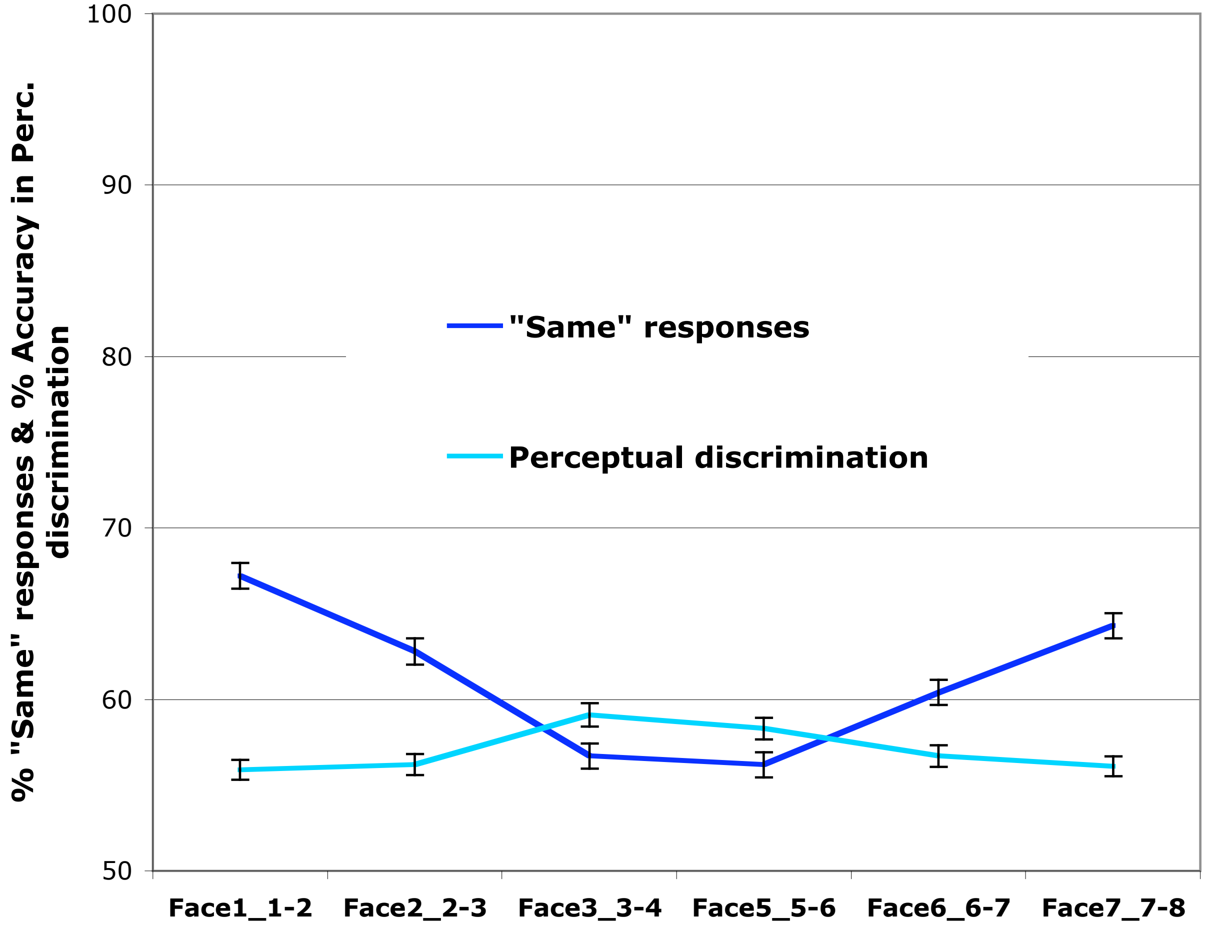
100
90
80
70
60
50
% "Same" responses & % Accuracy in Perc. discrimination
"Same" responses
Perceptual discrimination
Face1_1-2
Face2_2-3
Face3_3-4
Face5_5-6
Face6_6-7
Face7_7-8

Figure 4 :

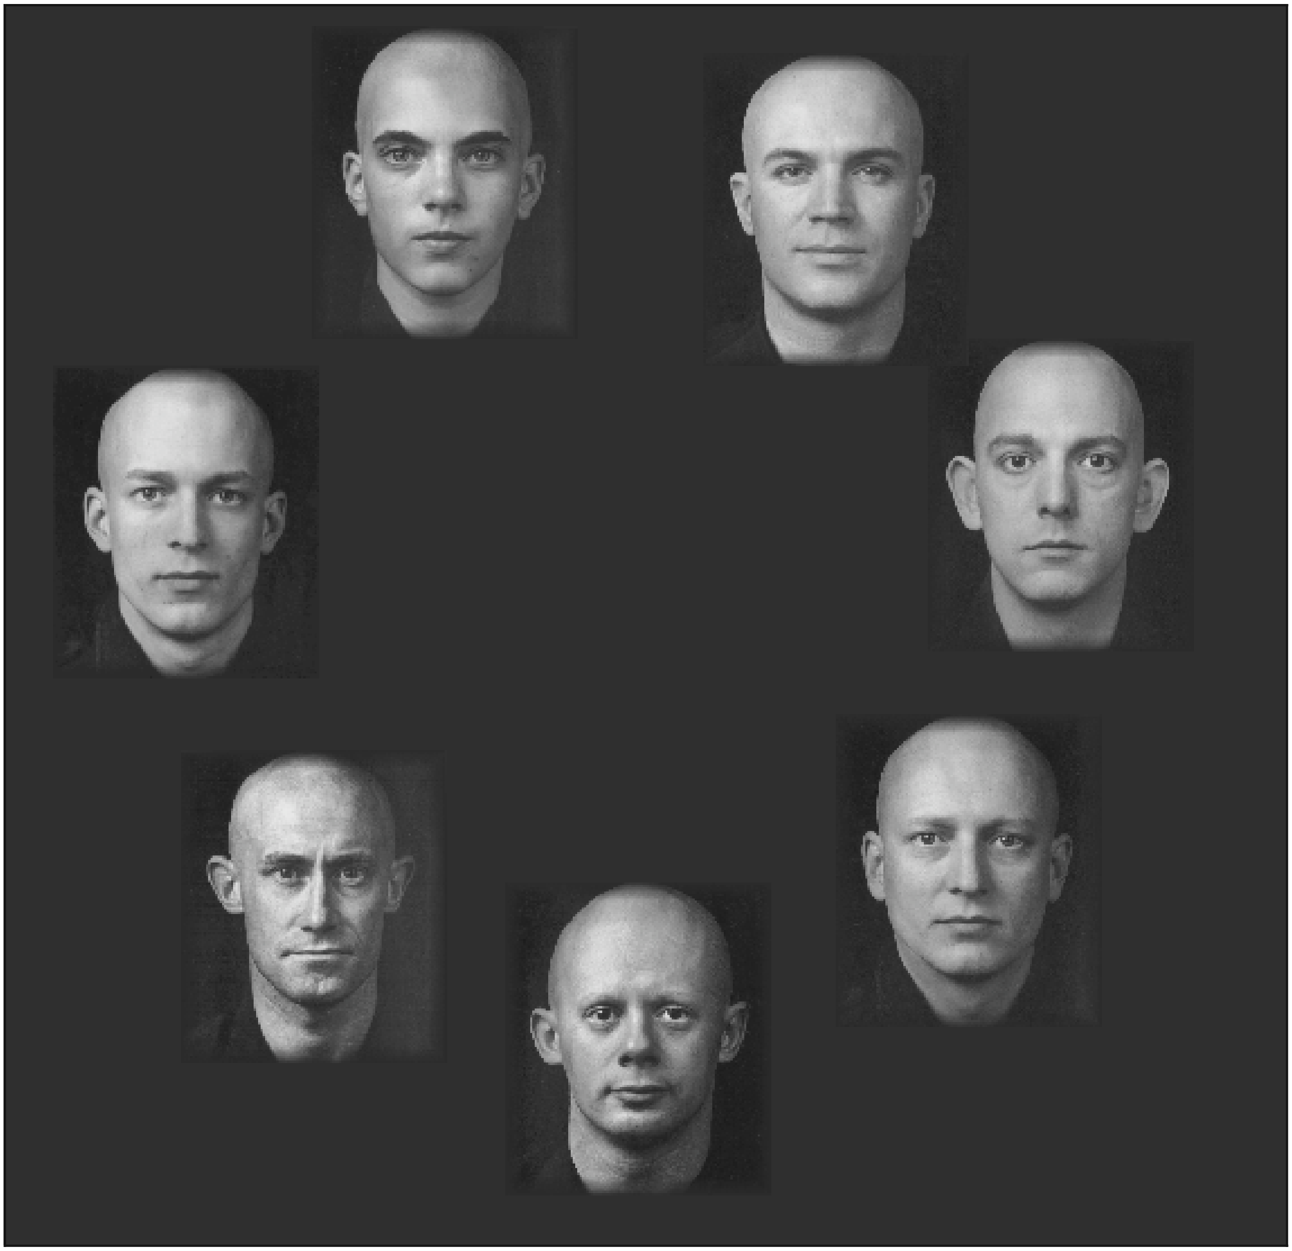

Figure 5

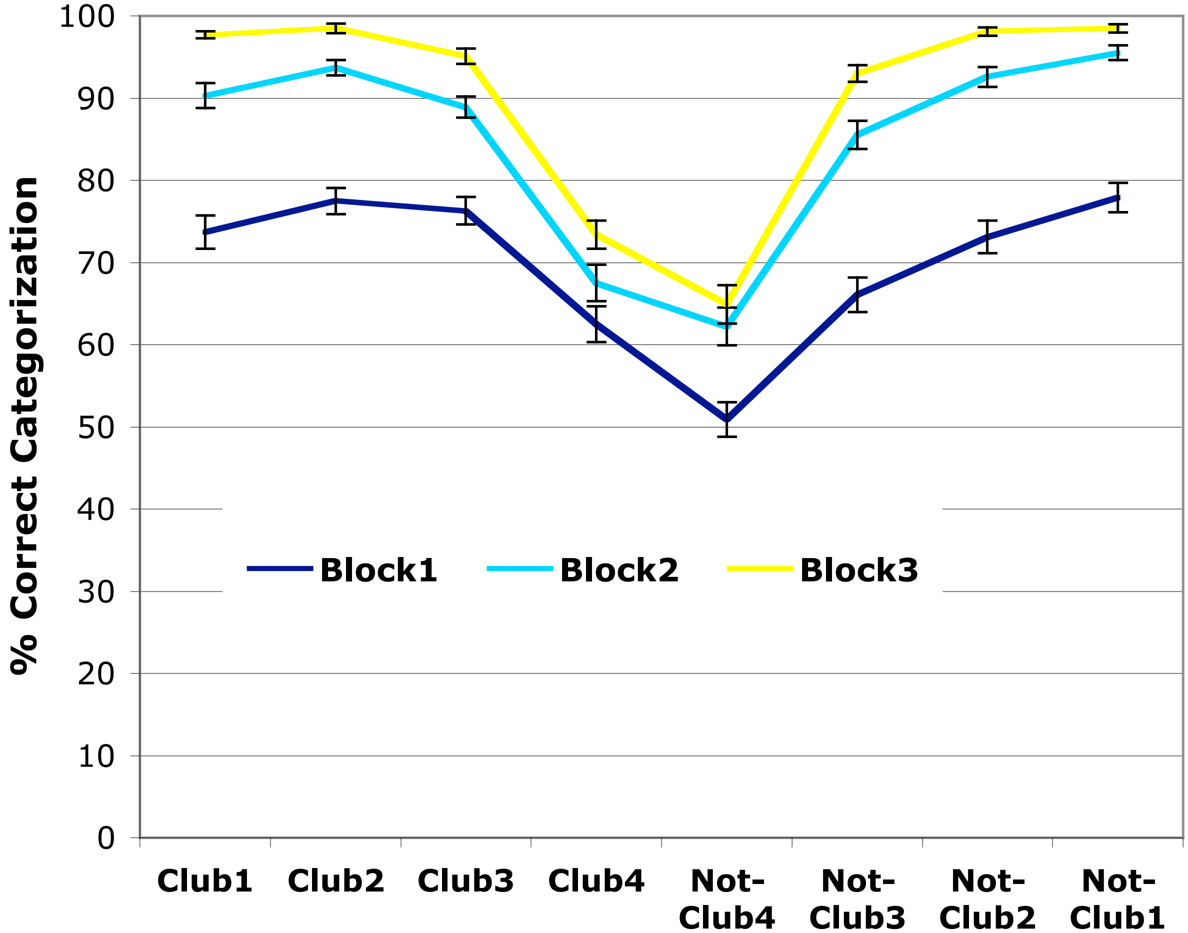
100
90
80
70
60
50
40
30
20
10
0
% Correct Categorization
Block1
Block2
Block3
Club1
Club2
Club3
Club4
Not-Club4
Not-Club3
Not-Club2
Not-Club1

Figure 6 :

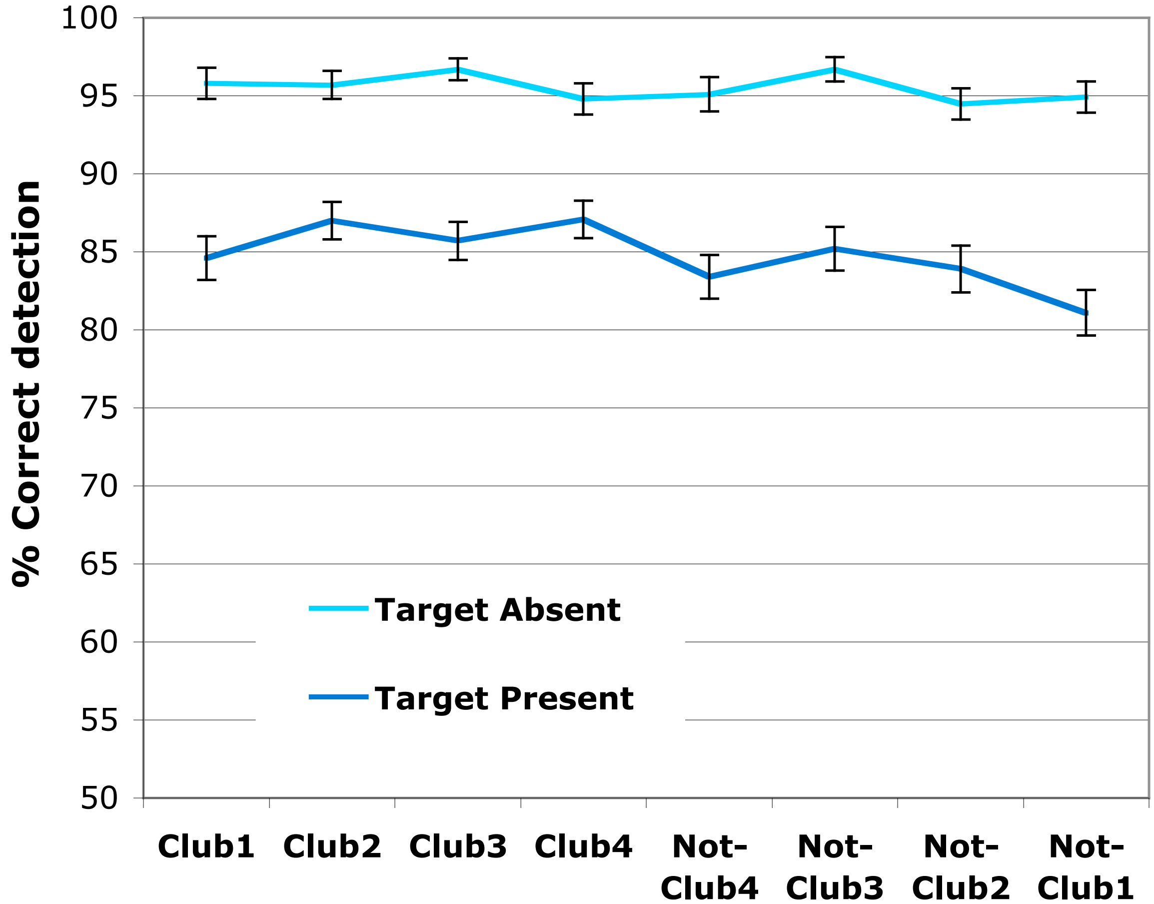
100
95
90
85
80
75
70
65
60
55
50
% Correct detection
Target Absent
Target Present
Club1
Club2
Club3
Club4
Not-Club4
Not-Club3
Not-Club2
Not-Club1