

Distinctiveness Effects in Face Memory Vanish with Well-Controlled Distractors

Nicolas Davidenko (ndaviden@psych.stanford.edu)

Department of Psychology, 450 Serra Street, Building 420
Stanford, CA 94305 USA

Michael Ramscar (michael@psych.stanford.edu)

Department of Psychology, 450 Serra Street, Building 420
Stanford, CA 94305 USA

Abstract

The “distinctiveness effect” in face memory holds that distinctive faces are recognized better than typical faces, as reflected by higher hit rates and lower false alarm rates during recognition tasks. We propose that this effect is due to a systematic bias in the similarity relationships between target and distractor faces. Specifically, we claim that distinctive faces are recognized more successfully because they are on the whole more distinguishable from a set of distractor faces, but not because of any inherent advantage in the way we represent them in memory. To test this claim, we conduct two memory experiments using a low-dimensional parameterization of face space based on silhouetted face profiles (Davidenko, 2004). In Study 1, we conduct an old/new recognition task with face silhouettes to replicate the classic distinctiveness effect. In Study 2, we construct a set of face stimuli that eliminates the asymmetry of target-distractor distances between typical and distinctive faces. Under these conditions, distinctive faces lose their recognition advantage and actually produce poorer performance than typical faces. We discuss possible causes of this “typicality effect” in terms of perceptual learning.

Introduction

Recognition memory paradigms have long been used as a tool for studying the way we represent faces. In a typical experiment, a subject observes a set of target faces and later decides which faces from among a set of distractors were part of the target set. Responses are elicited as either “old/new” judgments, or as particular choices from a set of alternatives. One of the most consistent findings from this type of experiment is that distinctive faces (those that are rated as “easy to spot in a crowd”) are correctly recognized more frequently than typical faces. This “distinctiveness effect” is robust and has been widely reported in the literature (e.g., Bartlett, Hurry, & Thorley, 1984; Shepherd, Gibling, & Ellis, 1991). In such studies, distinctive faces produce more hits (correct detections) and fewer false alarms (incorrect detections) than typical faces. This difference is often interpreted to mean that distinctive faces are *better remembered* than typical faces, a conclusion that is at odds with other findings in the literature (e.g., that we are slower to identify distinctive faces as faces, as compared to typical faces; Valentine & Bruce, 1986). We suggest that a closer examination of exactly what is meant by “better

remembered” in these tasks can help resolve this apparent controversy.

It seems clear that choice of distractor faces can critically affect the difficulty of a recognition task. In particular, if a distractor is very similar to a target face, the two are likely to be confused, which can lead to an incorrect recognition judgment. It is surprising, then, that the choice of distractors in face recognition studies is rarely manipulated or even controlled. This means that the similarity relationships between targets and distractors cannot be guaranteed to be equivalent for typical and distinctive faces. In fact, because typical faces lie in a denser, more central region of face space (Valentine, 1991), they will tend to be more similar to the distractor set than distinctive faces. This raises the possibility that recognition performance on typical faces is artificially low simply because it is harder to discriminate the typical faces from the set of distractors.

We propose that this asymmetry is at the root of the distinctiveness advantage in memory, and that the relative location of distractors in face space with respect to typical and distinctive targets is sufficient to explain the effect. In other words, we suggest that there is nothing about distinctive faces *per se* that makes them more memorable, but rather that systematic differences in similarity relationships between targets and distractors for the two types of faces are responsible for performance differences. Indeed, theories of perceptual learning (e.g., Tanaka & Gauthier, 1997) should predict a memory *disadvantage* for distinctive faces, since we have less perceptual experience with peripheral regions of face space where these faces lie. This prediction would be consistent with general notions of perceptual expertise, whereby we are better able to discriminate and remember items with which we are more familiar. Confirming that the distinctiveness effect is actually an artifact of the selected distractor set would therefore help advance our understanding of face representation.

Silhouetted Face Profiles

To investigate the influence of distractor choice on the memorability of faces, we devised a low-dimensional parameterization of face space based on the shape of silhouetted face profiles (see Figure 1). A previous set of validation studies indicates that silhouettes are perceived and processed like genuine face stimuli (see Davidenko,

2004). In particular, silhouettes provide enough information for gender identification and age estimation, they elicit reliable and coherent ratings of attractiveness and distinctiveness, and they are better remembered when they are presented upright than upside down, replicating the face-inversion effect (Yin, 1969). In addition, silhouettes can be matched to their front-view gray-scale counterparts with considerable accuracy, comparable to performance with “synthetic faces” (Wilson, Loffler, & Wilkinson, 2001).




Figure 1. A silhouetted face profile obtained by thresholding and cropping a gray-scale profile face image.

Silhouettes can be accurately reproduced by interpolating smooth splines across a set of 16 informative points along the contour of a face profile. A principal components analysis on the location of these points from a large collection of silhouettes (taken from the FERET Face Database; Phillips, Moon, Rizvi, & Rauss, 1998) results in a 32-dimensional parameterization of silhouette face space, which effectively requires only 16 dimensions for veridical reconstruction (Davidenko, in preparation). Using this parameterization, we can easily create artificial but realistic silhouettes in specified regions of face space. In particular, we can manipulate target and distractor silhouettes for use in a series of recognition experiments, precisely controlling their relative distances.

In the first of two studies, we test whether the classic distinctiveness effect replicates with face silhouettes. In the second study, we manipulate the placement of distractors in silhouette face space so that their distance from typical and distinctive face silhouettes is matched. Under these conditions, we examine whether any recognition advantage for distinctive face silhouettes persists once the asymmetry in target-distractor similarities is eliminated, or whether perceptual learning of typical regions of face space results in a typicality advantage.

Study 1. Replicating the Distinctiveness Effect

Method

The design of this study was a standard old/new task using parameterized face silhouettes. Participants were 12 Stanford undergraduates with normal or corrected vision, who participated in the experiment for course credit.

Stimuli. Using Matlab, 210 silhouettes were constructed online for each participant, by sampling randomly from the

multi-normal distribution of face silhouettes derived from the parameterization of silhouette face space (Davidenko, 2004). Typical and distinctive silhouettes were defined by a median split on the distance in silhouette face space between a silhouette and the overall norm (with the larger distances corresponding to the distinctive silhouettes). In previous studies, we have found a high correlation between this geometric measure and human ratings of distinctiveness (Davidenko, in prep.). Of the 105 typical silhouettes, 70 were randomly designated as targets and 35 as distractors, and likewise for the 105 distinctive silhouettes. Examples of typical and distinctive silhouettes are shown in Figure 2.




Figure 2. Typical (A) and distinctive (B) silhouettes.

Procedure. Participants completed a total of 35 trials. In each trial, they observed a sequence of 4 “training” silhouettes on a computer screen, followed by an 8-second retention interval. They then observed a sequence of 4 “test” silhouettes (2 novel, and 2 from the training set) in random order. For each test silhouette, they entered ‘y’ or ‘n,’ indicating whether or not they believed the silhouette had been shown in the preceding training set.

Results

Each response was coded as a hit, miss, false alarm, or correct rejection. A 2x2 repeated measures ANOVA revealed a significant interaction between Typicality and Response Type ($F_{1,11} = 27.24, p < .001$). Replicating previous studies using front-view face images (e.g., Shepherd et al., 1991), distinctive face silhouettes received significantly higher hits (.63 vs. .54; 2-tailed paired t-test = 2.78, $p < .05$) and significantly lower false alarms (.21 vs. .33; 2-tailed paired t-test = 2.57, $p < .05$) than typical faces (see Figure 3). Separate d' scores for performance on typical and distinctive silhouettes were computed for each participant. A paired t-test revealed a significant recognition advantage for distinctive silhouettes over typical silhouettes across participants, with the average d' for distinctive silhouettes = 1.17 and d' for typical silhouettes = .59 (paired t-test = 4.89, $p < .001$). This replication of the distinctiveness effect provides further validation of silhouettes as genuine face stimuli.




Figure 3. Hits and False alarms for typical and distinctive silhouettes. All differences are significant to $p < .05$.

Face-space explanation

As mentioned earlier, the distribution of typical and distinctive faces in face space should have consequences on the set-wise distances between distractors and the two types of targets (see Valentine, 1991). In particular, we should expect randomly chosen distractors to be on the whole closer to a set of typical faces than to a set of distinctive faces. To show this, we simulated random sets of faces researchers might use in a standard face memory experiment and compared the relevant distances in face space. In each simulation, we sampled 30 typical and 30 distinctive faces, and 30 random distractor faces from one of two versions of face space: a silhouette-like face space with 16 dimensions, and front-view-like face space with 400 dimensions (as estimated by Penev and Sirovich, 2000). We measured pairwise distances between each distractor and each target face. Figure 4 shows the mean distances between distractors and the two types of targets. Standard errors were calculated by running 1000 simulations. In both face spaces, we obtain the same reliable difference: distractors are statistically closer to typical faces than to distinctive faces (all differences significant to $p < .01$). The results of these simulations support the notion that distinctive faces may elicit more hits and fewer false alarms because they are more distinguishable from the distractors. The conclusion leaves open the question of whether distinctive faces would still be more memorable than typical faces if these differences were somehow eliminated. In the next study, we address this question by constructing equally-spaced distractors for each target silhouette.

Study 2a. Deconfounding Distinctiveness and Isolation from Distractors: A Typicality Effect

Method

We conducted a short-delay recognition task in a 3-alternative-forced-choice (3AFC) paradigm. Participants were 16 Stanford undergraduates who participated for course credit.

Stimuli. Using Matlab, 100 typical and distinctive target silhouettes were constructed in the same way as in Study 1, by sampling from the multi-normal distribution of silhouettes from our parameterized silhouette face space. Two distractors were constructed specifically for each target silhouette by varying the values on 3 of its 16 dimensions. The distractors were, in essence, small translations of the target silhouettes in silhouette face space. The magnitude of these translations was kept constant for all targets, while the direction of the translations varied randomly across targets. This ensured that, statistically, distractors were neither more nor less distinctive than their respective targets (which could, over time, have become a cue for participants as to the oldness or newness of the presented item)

Mean distance from distractors




Figure 4. Mean distance from distractors to target faces. Distances are statistically larger for distinctive targets, in both 16-dimensional and 400-dimensional face spaces.

Procedure. Participants completed 100 trials in which they observed a target silhouette for 2.5 seconds, a random line mask for 2 seconds, and a set of 3 test silhouettes. They were asked to decide which of the 3 test silhouettes was the one they had just seen. A sample trial is shown in Figure 5.




Figure 5. A sample trial from Study 2a.

Results

Performance was coded as percent identification of the target silhouette. Mean performance across participants was

61% for typical silhouettes and 56% for distinctive silhouettes, revealing a significant *disadvantage* for distinctive silhouettes (2-tailed paired t-test = 2.20, $p < .05$). In other words, we find a *typicality effect*, whereby typical face silhouettes are recognized with more accuracy than distinctive face silhouettes.

One way to interpret this result is that due to a lifetime of experience with faces, we become more sensitive to and better at discriminating between faces in dense, over-learned regions of face space. This interpretation would be consistent with reports of an own-race advantage in face recognition (e.g., Tanaka, Kiefer, & Bukach, 2004). However, there is an alternative explanation for our finding. It is possible that participants developed specific strategies for recognizing silhouettes during the course of the experiment, and that this learning was biased toward typical silhouettes. As shown schematically in Figure 6, the region of silhouette face space occupied by the set of typical target silhouettes observed during the experiment is smaller and more dense than the region occupied by the set of distinctive silhouettes. Because the silhouettes are embedded in a 16-dimensional space, these differences are actually more pronounced than is apparent in the diagram.




Figure 6. Typical silhouettes observed in Study 2a occupy a smaller, denser area of silhouette face space than distinctive silhouettes.

To test whether the typicality effect shown in Study 2a is due to a lifetime of perceptual learning or, more trivially, to online learning during the experiment, we conducted a variation of the study where the size and density of the regions corresponding to typical and distinctive silhouettes were matched.

Study 2b. Short-Term or Long-Term Perceptual Learning?

Method

This study was conducted similarly to Study 2a, except for the way in which distinctive silhouettes were constructed. To avoid differences in the size and density of the two regions of silhouette face space, we defined the set of distinctive silhouettes as a translation of the set of typical silhouettes, in a particular direction in face space. That is, the set of distinctive silhouettes was identical to the set of

typical silhouettes, except that it was centered in a different location of the space, still within the realm of normal-looking silhouettes. To avoid specific item effects, we ran two between-subject conditions where we reversed the direction of translation (see Figure 7 for an illustration). Participants were Stanford undergraduates who participated for course credit, of which 16 were randomly assigned to Condition 1, and 14 to Condition 2.




Figure 7. Regions of face space occupied by typical and distinctive silhouettes used in Study 2a, conditions 1 and 2.

Results

Just as in Study 2a, performance was measured as percent correct in a 3AFC task. In both conditions, participants were more successful at recognizing typical silhouettes than distinctive silhouettes. In condition 1, recognition of typical vs. distinctive faces was 64% vs. 59% (2-tailed paired t-test = 2.32, $p < .05$), and in condition 2 it was 62% vs. 57% (2-tailed paired t-test = 2.42, $p < .05$). Figure 8 summarizes the results of Studies 2a and 2b.




Figure 8. Recognition performance in 3AFC task in Studies 2a and 2b (conditions 1 and 2).

By matching the sizes of the regions in silhouette face space corresponding to the typical and distinctive target sets, we can reject the alternative hypothesis suggested earlier. While we cannot rule out the possibility that participants are learning over the course of the experiment, this alone cannot account for the superior performance on

typical silhouettes in Study 2b. The regions of silhouette face space occupied by typical and distinctive silhouettes are the same size and density, and thus perceptual learning should not favor one type of silhouette over the other. The result provides strong evidence that prior experience with central regions of face space (corresponding to typical faces) improve our ability to discriminate and represent these faces in memory.

Discussion

The results reported in this paper may appear at first glance contradictory. In Study 1, we successfully replicated the distinctiveness effect with face silhouettes and confirmed that distinctive faces elicit higher discriminability, as measured by hits and false alarms in a recognition task. In Studies 2a and 2b, we provided evidence that *typical* faces are actually recognized with better accuracy, as reflected by a higher percent correct on a 3AFC task. How can we reconcile these two results?

The simplest answer is that the choice of distractors in recognition tasks, which is often overlooked and rarely manipulated, can dramatically influence performance in recognition tasks. In forced choice as well as old/new paradigms, the choice of distractors can determine which of two sets of items will be discriminated better than the other. For example, if the set of distractors is statistically more similar to set A than to set B, we should expect items in set B to be more distinguishable from distractors, and thus successfully recognized more often. Conversely, if the set of distractors is more similar to set B, we should expect better performance for items in set A. It is misleading to report that one set of items is more memorable than another without describing the structure of the distractor set, in particular, the relative proximity of the distractor set to the two sets being compared.

Perhaps a more informative way of reporting the memorability of items is to measure the *accuracy* of our representations. In Studies 2a and 2b, we attempted to do just that. By constructing distractors that were equally spaced from typical and distinctive targets, we were able to obtain a fair comparison of performance between the two sets. We found in three separate conditions that typical face silhouettes are remembered more *accurately* than distinctive faces silhouettes; that is, typical faces are successfully recognized more often than distinctive faces, when subjected to equivalent modifications.

The ‘typicality effect’ we report with faces should not be too surprising. A variety of results from human perception and memory report similar advantages for well-learned regions of representational space. For example, prototypical colors are discriminated better than less prototypical ones (Lucy & Shweder, 1979); realistic chess patterns are better remembered by experts than random assortments of chess pieces (Chase & Simon, 1973). One reason this effect has not been reported in faces is the lack of a well-accepted parameterization of face space.

To construct the stimuli for Studies 2a and 2b, we relied on a full parameterization of silhouette face space that allowed us to both sample faces from a realistic population

and also modify the images in precisely controlled ways. Without this precise control over the construction of face stimuli, we would not be able to isolate the effect of distractors in recognition memory. The studies presented here hopefully demonstrate the value of constructing parameterized stimulus spaces for the study of human memory and representation. We predict, in particular, that the ‘typicality effect’ reported with face silhouettes will generalize, not only to front-view face representations, but also to other object categories whose items are centrally distributed.

Acknowledgments

This research was funded in part by an NSF Graduate Research Fellowship. The authors would like to thank Dan Yarlett for his helpful suggestions.

References

- Bartlett, J. C., Hurry, S., Thorley, W. (1984). Typicality and familiarity of faces. *Memory & Cognition*, 12, 219-228.
- Chase, W. G., & Simon, H. A. (1973). Perception in Chess. *Cognitive Psychology*, 4, 55-81.
- Davidenko, N. (Manuscript in preparation). Silhouetted face profiles: a low-dimensional parameterization of face space.
- Davidenko, N. (2004). Modeling face-shape representation using silhouetted face profiles [Abstract]. *Journal of Vision*, 4(8), 436a.
- Lucy, J., and Shweder, R. (1979). Whorf and his critics: Linguistic and nonlinguistic influences on color memory. *American Anthropologist*, 81, 581-615.
- Penev, P. S. and Sirovich, L. (2000). The Global Dimensionality of Face Space. *Proceedings of IEEE International Conference on Automatic Face and Gesture Recognition*, pp. 264-270.
- Phillips, P., Moon, H., Rizvi, S., & Rauss, P. (1998). The FERET evaluation. In H.W. et al., editor, *Face Recognition: From Theory to Applications*, 244-261.
- Shepherd, J. W., Gibling, F., & Ellis, H. D. (1991). The effects of distinctiveness, presentation time and delay on face recognition. *European Journal of Cognitive Psychology*, 3(1), 137-145.
- Tanaka, J. W., & Gauthier, I., (1997). Expertise in object and face recognition. *Mechanisms of Perceptual Learning*, Vol.36 (eds. Medin, Goldstone & Schyns), 83-125.
- Tanaka, J.W., Kiefer, M., & Bukach, C. M. (2004). A holistic account of the own-race effect in face recognition: evidence from a cross-cultural study, *Cognition*, 93, B1-B9.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion and race in face recognition. *Quarterly Journal of Experimental Psychology*, 43A, 161-204.
- Wilson, H.R., Loffler, G., & Wilkinson, F. (2002). Synthetic faces, face cubes, and the geometry of face space. *Vision Research*, 42(27), 2909-2923.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81, 141-145.