

Typicality effects in face and object perception: Further evidence for the attractor field model

JAMES W. TANAKA

University of Victoria, Victoria, British Columbia, Canada

AND

OLIVIER CORNEILLE

Université Catholique de Louvain, Louvain-la-Neuve, Belgium

In a previous study, it was shown that a 50/50 morph of a typical and an atypical parent face was perceived to be more similar to the atypical parent face than to the typical parent face (Tanaka, Giles, Kremen, & Simon, 1998). Experiments 1 and 2 examine face typicality effects in a same/different discrimination task in which typical or atypical faces and their 80%, 70%, 60%, and 50% morphs were presented sequentially (Experiment 1) or simultaneously (Experiment 2). The main finding was that in both modes of presentation, atypical morphs were more poorly discriminated than their corresponding typical morphs. In Experiment 3, typicality effects were extended to the perception of nonface objects; in this instance, it was found that 50/50 morphs of birds and cars were judged to be more similar to their atypical parents than to their typical parents. These results are consistent with an attractor field model, in which it is proposed that the perception of a face or object stimulus depends not only on its fit to an underlying representation, but also on the representation's location in the similarity space.

A consistent finding in the face recognition literature is that atypical faces are more easily learned and better recognized than typical faces. The atypicality face advantage has been demonstrated for the recognition of newly familiarized faces (J. C. Bartlett, Hurry, & Thorley, 1984; Light, Kayra-Stuart, & Hollander, 1979; Vokey & Read, 1992) and famous faces (Valentine & Bruce, 1986). The "face space" model, whereby faces are defined by a large, but finite set of visual properties or dimensions (e.g., length of nose, intereye distance) (Busey, 1998; Valentine, 1991), provides the most widely accepted account of the atypicality advantage. A given face is specified as a single point in this high-dimensional face space according to its values along these dimensions. It is assumed that the face representations are normally distributed in the similarity space; a higher proportion or density of typical faces is found at the center of the space and a lower proportion of atypical faces is found in the periphery. Given that typical faces lie in a region of higher face density, there is more competition from other neighboring typical representations, making them less distinctive in memory (Valentine, 1991). Atypical faces, on the other hand, are more distinctive because they are located in a sparser region of face space, where there are fewer representations, and therefore, less competition.

The effects of typicality were directly tested in a series of experiments by Tanaka, Giles, Kremen, & Simon (1998), in which typical and atypical faces were pitted against one another using a morphing procedure. In this paradigm, equal contributions from a typical parent face and an atypical par-

ent face were combined to produce a 50/50 face morph. The main finding was that participants judged the face morph as bearing a stronger resemblance to its atypical parent than to its typical parent. Critically, no atypicality preference was found when the morph was paired with an unrelated atypical and typical face, indicating that the response did not reflect a general bias toward atypical faces (Tanaka et al., 1998, Experiment 2). Moreover, results from computer modeling simulations suggested that the atypicality bias is not an artifact of the morphing process itself (M. S. Bartlett, in press; M. S. Bartlett & Tanaka, 2007). When we tested the same stimuli used in the Tanaka et al. study, both gray-level and normalized Gabor representations of a 50/50 morph face were equally similar to their typical and atypical parents when image similarity was measured by Euclidean distance and cosines. However, after face learning in an attractor network model, the model demonstrated a systematic bias toward the atypical parent face when presented with a 50/50 morph. The behavioral and computational results indicate that the atypicality bias is not attributable to a predisposition toward atypical faces or an irregularity in the morphing procedure. Instead, the atypicality bias reveals something about the underlying relationship between typical and atypical faces in face space.

Distance–Density Hypothesis and the Attractor Field Model

The atypicality results are consistent with the predictions of the distance–density hypothesis (Krumhansl,

1978), in which it has been claimed that the similarity between exemplars is a function not only of their metric distances in the similarity space, but also of the spatial density of representations in the surrounding configuration. The distance–density relationship can be represented in an attractor field model. According to this approach, representations attract perceptual inputs by carving out valleys or basins of activations in the similarity space (Tank & Hopfield, 1987). In recognition, inputs decrease their computational energy by following the path that leads to the nearest attractor basin in representational space. The stimulus need not be a perfect fit with its underlying representation, only a close enough approximation to fall within the boundaries of

the representation's attractor field. The potential activation of any given representation is proportional to the size or span of its attractor field, so that representations with large attractor fields will capture a broader range of stimulus inputs than representations with smaller attractor fields. As shown in Figure 1, by virtue of their location in a sparser subregion of similarity space, the attractor fields of atypical representations are relatively large in comparison with the smaller attractor fields of typical faces, which are located in a denser subregion. In this model, face morphs are represented as points along the vector that connect a typical face and an atypical face. As shown in Figure 1, the atypicality bias is a consequence of the 50/50 morph lying closer to the larger attractor

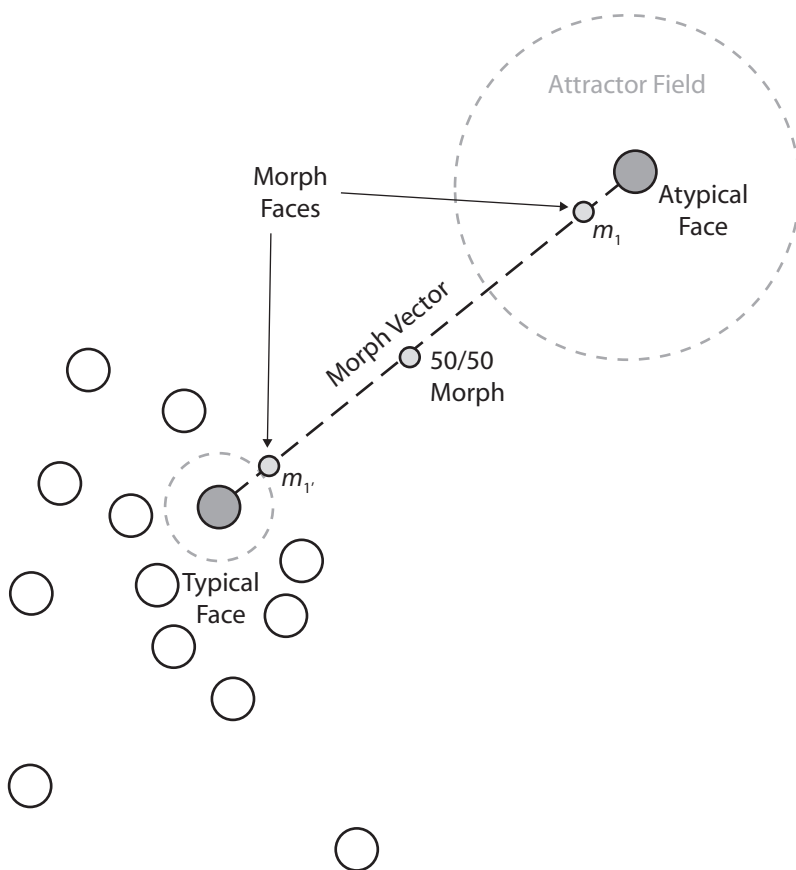


Figure 1. Diagram of the attractor field model. The atypical and typical representations of interest are depicted as filled gray circles. Located in a sparse subregion of face space, the atypical face representation has a larger attractor field relative to the smaller attractor field of the typical representation, situated in a denser subregion. The morph vector is indicated by the dashed line connecting the atypical and typical parent face representations. The morph representations are located as points along the vector. The 50/50 morph is located at the midpoint of the vector and is equal from its atypical and typical parent representations. The atypicality bias is a result of the 50/50 morph lying closer to the boundary of the atypical representation's expanded attractor field than to the boundary of the typical representation. In the diagram, distance between morph m_1 and its atypical parent is equal to the distance between morph m_1 and its typical parent representation. According to the attractor field hypothesis, discriminations between m_1 and its atypical parent should be more difficult than discriminations between m_1 and its typical parent, given that m_1 lies inside the expanded attractor field of the atypical representation, whereas m_1 lies outside the boundary of the more restricted attractor field of the typical representation.

field of the atypical representation than to the smaller attractor field of the typical representation.

Although the attractor field model predicts an atypicality *bias* in perceptual similarity, the model predicts an inverse, atypicality *disadvantage* in perceptual discrimination. That is, presented with inputs whose differences are equivalent in their physical magnitudes, the larger attractor field of the atypical representation should be less sensitive to differences than the more restricted attractor field of the typical representation. The hypothesized discrimination disadvantage is illustrated in Figure 1, in which morph m_1 falls within the attractor field of the atypical representation and therefore is more likely to be confused with the atypical representation. In contrast, morph m_1 falls outside the attractor field of the typical representation and therefore is more likely to be differentiated from the typical exemplar. Thus, despite the equivalence of their physical differences, the dissimilarities are better differentiated by the more restricted attractor field of the typical representation than by the larger field of the atypical representation. This prediction was tested in the following experiments, in which participants made a same/different discrimination to typical or atypical faces and their incremental morphs presented either sequentially (Experiment 1) or simultaneously (Experiment 2).

EXPERIMENT 1

Face Typicality Effects in a Sequential Perceptual Matching Task

In Experiment 1, we tested the typicality effects in a sequential match-to-sample task in which a typical or an atypical study face was presented for 2 sec, followed by a test face. In the *same* condition, the typical (atypical) study and test faces were identical. In the *different* condition, the test face was a morph of the atypical and typical parent faces, as shown in Figure 2. For the atypical *different* trials, the morph contained an 80%, 70%, 60%, or 50% contribution from its atypical parent face. For typical *different* trials, the morph contained an 80%, 70%, 60%, or 50% contribution from its typical parent. If the morphs more strongly resemble their atypical parent than their typical parent, participants should have more difficulty detecting an atypical morph as *different* than they would have with a typical morph.

Method

Participants. Thirty-six undergraduates from an introductory psychology class at Oberlin College participated in this experiment for course credit. All of the participants had normal or corrected-to-normal vision.

Stimuli. Four typical male faces and four atypical male face stimuli from a previous study (Tanaka et al., 1998) were used as the parent face images. Using the Morph 2.5 program, we morphed each of the four typical faces with each of the four atypical faces, yielding 16 typical–atypical face combinations. To construct a morph of the atypical and typical parent faces, we identified corresponding control points on the two parent images (e.g., the corner of the left eye on Parent Face Image 1 and Parent Face Image 2). The number of control points for facial features was kept constant, with 12 points on the mouth, 7 points on each eye, 6 points on the nose, 5 points on each eyebrow, and 22 points for the outline of the face. Depending

on the desired level of morphing, we defined new control points for the morph face by moving the specified distance along the vector connecting the control points in parent images. The locations of intervening pixels were linearly interpolated across the surface on the basis of the position of the nearest control point (Wolberg, 1990). We then employed a fade process in which the brightness values for each corresponding pixel were weighted according to the contribution of each parent image.

Face morphs were generated by varying the relative contribution of the atypical and typical parent face images to the face morph. For each atypical–typical face pairing, seven morph faces were generated, containing an 80%, 70%, 60%, or 50% contribution from the atypical (or typical) parent face. A family of seven morph faces was created for each of the 16 typical–atypical face pairings, yielding a total of 112 face morph stimuli. To minimize any potential artifacts introduced by the morphing process, we replaced the original atypical and typical parent images by face morphs containing a 99% contribution from either the atypical or the typical parent. Faces subtended visual angles of 2.5° and 4° in the vertical and horizontal dimensions, respectively. Face stimuli were presented on a computer monitor with a resolution of 72 dots per inch (dpi).

Procedure. The participants viewed two sequentially presented faces; their task was to decide whether the faces were the *same* or *different* and to press the appropriately marked key on the keyboard. In the task instructions, it was emphasized to participants that they

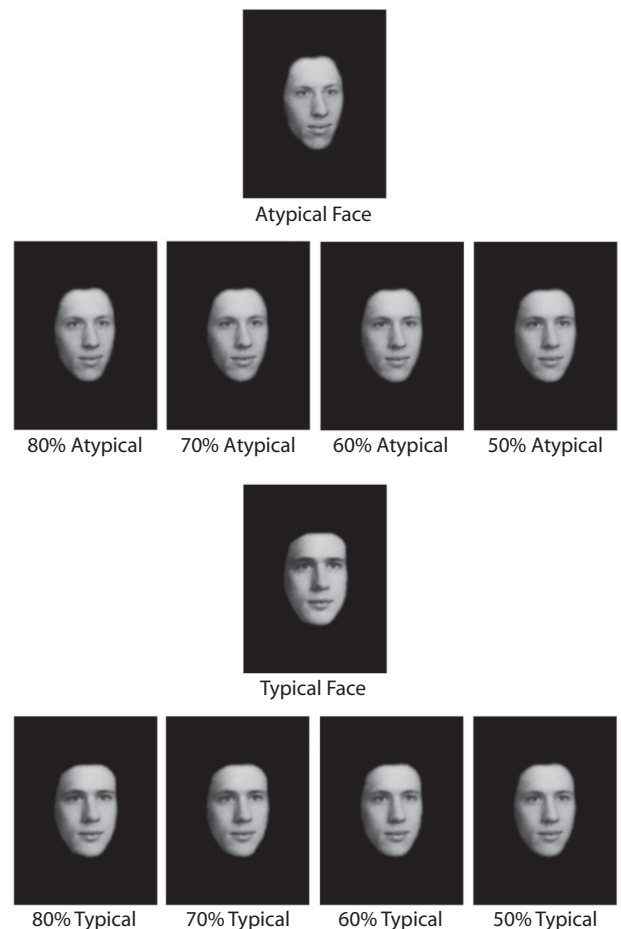


Figure 2. (A) An example of an atypical parent face and a morph image containing 80%, 70%, 60%, and 50% of its atypical parent face. (B) An example of a typical parent face and a morph image containing 80%, 70%, 60%, and 50% of its typical parent face.

should respond “same” only if the probe and test faces were physically identical. For each trial, a typical (or an atypical) study face was centrally shown for 2 sec, followed by a blank screen for a 1-sec interstimulus interval, and a central test face for 1 sec. The test face was either identical to the study face or an 80%, 70%, 60%, or 50% morph of the study face. There was an intertrial interval ITI of 2 sec.

For the *different* typical trials, the typical study face was matched with a typical–atypical face morph that was an 80%, 70%, 60%, or 50% morph of the typical parent face. For the *different* atypical trials, the atypical probe face was matched with an atypical–typical face morph that was an 80%, 70%, 60%, or 50% morph of the atypical parent face. The four typical and four atypical faces generated 16 morphing combinations, and at four levels of morphing, produced a total of 64 *different* atypical trials and 64 *different* typical trials. Each typical and atypical study face was matched for an equivalent number of *same* trials, yielding a total of 256 trials (128 *same* trials and 128 *different* trials). The trials were presented randomly on Macintosh computers using the SuperLab 1.5 experimental package.

Results and Discussion

To examine the effects of face typicality on discrimination performance, we compared the percentage of correct rejections (i.e., responding “different” when the study face was paired with a morph test face) to atypical study faces with the percentage of correct rejections to typical study faces. As shown in Figure 3, the percentage of correct rejections for atypical morphs was less than the percentage of correct rejections for typical morphs at every level of parental contribution, with the exception of the 50% morph level. A repeated measures ANOVA by participants, with face type (atypical, typical) and morph level (80%, 70%, 60%, 50%) as the main factors, showed a reliable effect of face type [$F(1,35) = 14.231, p < .01$] and morph level [$F(3,105) = 270.850, p < .01$] and a reliable face type \times morph level interaction [$F(3,105) = 6.402, p < .01$]. Planned comparisons between the atypical and typical responses showed reliable differences at every level of morphing ($p < .001$), with the exception of the 50/50 level ($p > .10$).

The ANOVA by items showed reliable main effects of face type [$F(1,30) = 7.518, p < .01$] and morph level [$F(3,30) = 160.578, p < .001$], but no interaction between face type and morph level [$F(3,90) = 1.963, p > .10$]. Planned comparisons between the atypical and typical items showed reliable differences at the 80% and 60% morph levels ($p < .05$), but not at the 70% and 50% levels ($p > .10$).

The nonparametric measure of A' was used as an index of discrimination by computing the number of hits minus the number of false alarms for each participant.

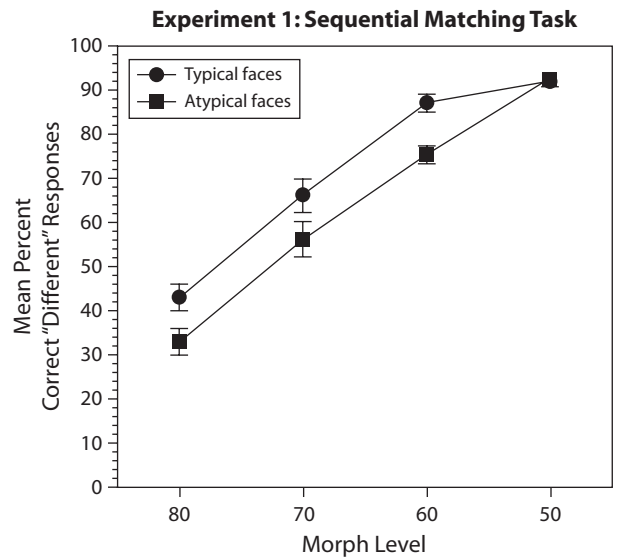


Figure 3. The mean percentage of correct “different” responses for pairs of atypical and typical parent faces and morphs at 80%, 70%, 60%, and 50% levels of morphing in a sequential same/different task in Experiment 2.

“Hits” were the trials in which the participant correctly responded “same” to trials containing two identical faces, and “false alarms” were trials in which the participant incorrectly responded “same” to trials containing a typical (or atypical) face and a morph. A' scores were calculated for the atypical and typical faces at the four morph levels. The ANOVA showed that the main effect of morph level was reliable [$F(3,105) = 182.161, p < .001$], but the main effect of face type was not [$F(1,35) = 2.320, p > .10$]. However, face type reliably interacted with morph level [$F(3,105) = 3.059, p < .05$]. Direct comparisons of the A' values showed that the differences between typical faces and atypical faces were reliable at the 60% and 70% levels of morphing ($p < .05$), but not at the 50% and 80% levels. (See Table 1.)

The main finding of Experiment 1 was that atypical morphs were less accurately discriminated than typical morphs. Hence, the previously reported atypicality bias in a perceptual similarity task (Tanaka et al., 1998) produces an atypicality *disadvantage* in a perceptual discrimination task. According to an attractor field model, the atypicality disadvantage is attributed to the expanded attractor field of the atypical representation, in which atypical morphs fall within the boundaries of the representation and, as a

Table 1
Mean A' Values for Atypical and Typical Face Discriminations in a Sequential Same/Different Matching Task in Experiment 1 at the Four Levels of Morphing

	Morph Level								Mean
	80%		70%		60%		50%		
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	
Atypical face	.715	.014	.806	.023	.889	.010	.949	.006	.840
Typical face	.744	.015	.837	.015	.917	.008	.934	.007	.858
Mean	.729		.822		.903		.941		

consequence, are more poorly differentiated. On the other hand, a typical morph of corresponding physical distance falls outside the smaller attractor field of the typical representation and is therefore better discriminated.

However, there were limitations to the sequential matching task used in Experiment 1. First, the task required immediate perceptual memory, to the extent that the participant was asked to compare the test face with a previously presented study face. Because it has been demonstrated that atypical faces are more memorable than typical faces (J. C. Bartlett et al., 1984; Light et al., 1979; Vokey & Read, 1992), it is conceivable that factors related to memory encoding and retrieval might account for differences in discrimination. Moreover, because there was no masking stimulus between the target and probe faces, differences in aftereffects created by atypical and typical target faces might have affected discrimination of the test faces. The issues of memory load and aftereffects were addressed in Experiment 2.

EXPERIMENT 2

Face Typicality Effects in a Simultaneous Perceptual Matching Task

To minimize memory demands and eliminate aftereffects, we employed a simultaneous matching task in Experiment 2. In this experiment, the participants made same/different judgments to pairs of typical or atypical faces and their morphs that were shown together, followed by a pattern mask. In the *same* condition, the two faces were either identical typical or atypical parent faces. In the *different* condition, the atypical or typical parent face was matched with an 80%, 70%, 60%, or 50% face morph of the atypical (or typical) parent face. According to the attractor field hypothesis, it was predicted that it should be more difficult for participants to detect differences between atypical parents and their morphs than to detect differences between typical parents and their morphs.

Method

Participants. Thirty-three undergraduates from an introductory psychology class at Oberlin College participated in this experiment for course credit. All of the participants had normal or corrected-to-normal vision.

Stimuli. The stimuli used in this experiment were the same atypical and typical parent face and faces morphs used in Experiment 1.

Procedure. In this task, two faces were presented side by side, and the participant's task was to decide whether the faces were the *same* or *different*. The participants were told that in order for them to respond "same," the faces should be physically identical. The faces were presented to the left and right of the center of the screen. On each trial, a fixation cross appeared for 500 msec, followed by a 500-msec central presentation of a face stimulus, followed by a 250-msec line mask that was replaced by a response prompt screen with the words "Same or Different?" At the appearance of the response screen, the participants indicated their response by pressing the appropriately labeled key. There was an ITI of 1.5 sec.

Following the design of Experiment 1, the *different* typical trials were composed of a typical face paired with an 80%, 70%, 60%, or 50% morph of the typical parent face. For the *different* atypical trials, the atypical probe face was paired with an 80%, 70%, 60%, or 50% morph of the atypical parent face. The left and right locations

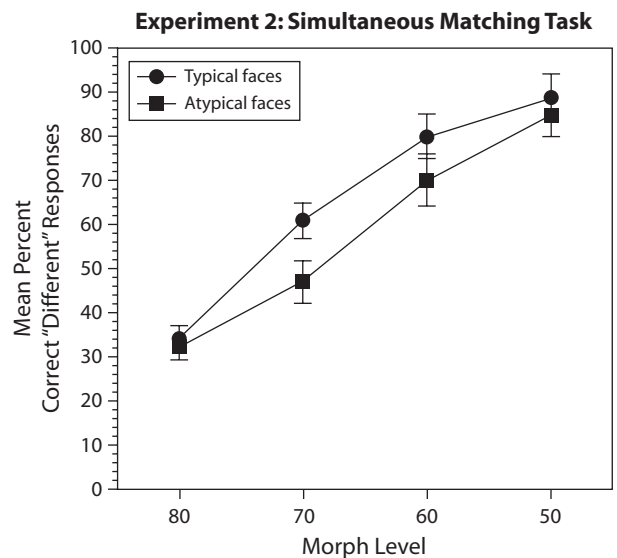


Figure 4. The mean percentage of correct "different" responses for pairs of atypical and typical parent faces and morphs at 80%, 70%, 60%, and 50% levels of morphing in a simultaneous same/different task in Experiment 3.

of the typical (or atypical) parent and the morph faces were counterbalanced. Pairing each of the four typical faces with each of the four atypical faces produced 16 morphing combinations, and at four levels of morphing, this yielded a total of 64 *different* atypical trials and 64 *different* typical trials. Each typical and atypical study face was matched with an equivalent number of *same* trials. In total, there were 256 trials (128 *same* trials and 128 *different* trials) in Experiment 2. Trials were presented randomly on Macintosh computers using the RSVP (Tarr & Williams, 1996) experimental package.

Results and Discussion

The percentages of correct rejections for atypical morphs and typical morphs at four morph levels are shown in Figure 4. An ANOVA by participants, with face type (atypical, typical) and morph level (50%, 60%, 70%, 80%) as factors, showed reliable effects of face type [$F(1,32) = 11.202, p < .01$] and morph level [$F(3,32) = 139.586, p < .01$] and a reliable face type \times morph level interaction [$F(3,96) = 4.918, p < .01$]. Planned comparisons between the atypical and typical responses showed reliable differences at the 60% and 70% morphing levels ($p < .001$), but not at the 50% and 80% levels ($p > .10$), which might be attributable to ceiling and floor effects, respectively.

The ANOVA by items showed reliable main effects of face type [$F(1,30) = 7.990, p < .01$] and morph level [$F(3,30) = 286.424, p < .001$] and a reliable interaction between face type and morph level [$F(3,90) = 3.436, p > .05$]. Planned comparisons between the atypical and typical items showed reliable differences at the 80%, 70%, and 60% contribution levels ($p < .05$), but not at the 50% level ($p > .10$).

An A' score was derived for each participant by subtracting their hits from their false alarms for atypical and typical faces shown at the four levels of morphing.

Table 2
Mean A' Values for Atypical and Typical Face Discriminations in a Simultaneous Same/Different Matching Task in Experiment 2 at the Four Levels of Morphing

	Morph Level								Mean
	80%		70%		60%		50%		
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	
Atypical face	.671	.037	.755	.035	.881	.016	.929	.010	.809
Typical face	.727	.025	.834	.025	.911	.014	.938	.012	.852
Mean	.699		.794		.896		.933		

The A' analysis showed reliable main effects of face type [$F(1,32) = 9.283, p < .01$] and morph level [$F(3,96) = 67.370, p < .001$], but no reliable interaction of morph level \times face type [$F(3,96) = 2.286, p > .10$]. Direct comparisons of A' values between typical and atypical morphs were reliable at the 80% and 70% levels of morphing ($p < .05$), but not at the easier 60% and 50% levels. (See Table 2.)

In Experiment 2, the discrimination of typical and atypical faces was tested in a matching task that reduced the participant's memory load and removed potentially confounding aftereffects. The main finding of this experiment was that under direct viewing conditions, the participants were more sensitive to changes in typical morphs than to changes in atypical morphs. Combined with previous morphing findings (Tanaka et al., 1998), the results from Experiments 1 and 2 demonstrate the complementary aspects of typicality effects in face perception. Although atypical faces are perceived to be more similar to their morphs than are typical faces (Tanaka et al., 1998), the present findings show that they are also more difficult to discriminate. Collectively, these results are consistent with an attractor field interpretation in which similarity and discrimination judgments are determined by the location of the representation in face space and its attendant attractor field size. The robustness of the typicality effect for faces raises the question of whether an analogous typicality effect can be shown for the perception of other, nonface objects. This question is addressed in Experiment 3.

EXPERIMENT 3

Atypicality Bias for Nonface Object Categories of Birds and Cars

Recently, there has been considerable discussion as to whether face recognition processes recruit cognitive operations and neural substrates that are specific to faces or whether these processes are employed in the recognition of other expert objects (Kanwisher, 2000; Tarr & Gauthier, 2000). Faces share the same internal features (i.e., eyes, nose, and mouth) arranged in a similar spatial configuration (i.e., the eyes are above the nose, which is above the mouth). Therefore, recognition of an individual face depends on the accurate discrimination of its features and their spatial configuration—what has been referred to as second-order relational properties (Diamond & Carey, 1986). If the typicality effect is dependent on the capacity to quickly individuate objects according to their second-

order relational properties, this effect may be exclusive to the domain of faces and face recognition processes.

However, members of other object categories are similarly differentiable on the basis of their second-order relational properties, and therefore might show a similar atypicality bias effect. That is, within any object category, there is sufficient structural variation so that some members are regarded as bearing a stronger resemblance to the prototypical category shape than other members (Jolicœur, Gluck, & Kosslyn, 1984; Murphy & Brownell, 1985). The structural typicality is reflected in category responses in which the typical exemplars (e.g., robin, sparrow) can be verified as category members (e.g., bird) faster than less typical exemplars (e.g., penguins, ostriches) (Jolicœur et al., 1984; Murphy & Brownell, 1985).

In Experiment 3, the atypicality bias was tested for two nonface object categories: birds and cars. We collected normative ratings to identify the car and bird exemplars that were judged to be typical and atypical category members. The typical and atypical exemplars were then morphed together (see Figure 5). Then, a new group of participants judged whether the morphed object more closely resembled the atypical object parent or the typical object parent. The prediction was that if the typicality effect is specific to faces, participants should be just as likely to judge the morph as resembling the typical object (bird or automobile) as the atypical object (bird or automobile). On the other hand, if the perception of objects—like the perception of faces—is influenced by structural typicality, participants should demonstrate an atypicality bias for the birds and automobiles.

Method

Participants. Thirty undergraduate students from Oberlin College participated in the rating study. A separate group of 35 undergraduate students from Oberlin College and the University of Victoria participated in the perceptual similarity experiment. The participants received course credit for their participation, and all of the participants had normal or corrected-to-normal vision.

Materials. The 40 car stimuli were taken from the October 1994 issue of *Road & Track* magazine and the 40 bird stimuli came from the *Peterson First Guides to Birds*. In selecting both sets of stimuli, the goal was to have a wide variety of cars and birds. The selected photographs depicted cars from oblique viewing angles, ranging from 20° to 40° from profile, whereas birds were shown in full profile. The images were scanned in grayscale at 348 dpi using a Microtek ScanMaker 600ZS scanner and reduced to 100 dpi in Adobe Photoshop. The car and bird stimuli were reflected about their vertical axes so that the car stimuli were right facing and the bird stimuli were left facing.




Figure 5. Examples of birds and car exemplars rated high or low in typicality and their 50/50 morphs.

The bird and car images were formatted onto 8.5×11 in. sheets of paper containing six bird images per page and eight car images per page. The pages were assembled into separate bird and car test booklets. A group of 30 participants from Oberlin College participated in the normative rating phase of the study. After receiving their test booklets, the participants were instructed to rate the birds (cars) according to how distinctive looking they appeared relative to the other birds (cars), 1 being *nondistinctive in appearance* and 5 being *very distinctive in appearance*. The presentation order of the booklets was counterbalanced across object type so that half of the participants rated the cars first and the other half rated the birds first. The participants had unlimited time to complete the experiment; most of them finished within 15 min.

From this set of 40 car and 40 bird images, 8 car and bird stimuli were selected on the basis of being either high or low in rated typicality. The 4 atypical birds had a mean distinctiveness rating of 4.38 and the 4 typical birds had a mean distinctiveness of 1.78. The 4 atypical cars had a mean distinctiveness rating of 4.22 and the typical cars had a mean distinctiveness rating of 1.77. The car and bird pictures were digitized using a MicroTek Z flatbed scanner and Adobe Photoshop. Using the Morph 2.5 program, we created atypical–typical morph images by averaging each typical car (bird) with each atypical car (bird), using the same morphing procedure described in the previous experiments. Approximately 30 control points were used for the bird stimuli and 25 control points were used for the car stimuli.

Morphing each atypical bird (car) with each typical bird (car) yielded a total of 16 atypical–typical car morph images and 16 bird morph images. The 16 face morph images were derived from a set of typical and atypical male faces that were used in a previous study (Tanaka et al., 1998). The face stimuli were presented on a computer monitor with a resolution of 72 dpi. Face images subtended visual angles of 2.5° and 4° in the vertical and horizontal dimensions, respectively. The bird stimuli subtended visual angles of 5° and 5° in the vertical and horizontal dimensions, whereas the car stimuli subtended a visual angle of 4° in the vertical dimension and 5° in the horizontal dimension.

Procedure. Each trial consisted of the participant seeing pairs of atypical and typical birds, cars, or faces. The images were presented on the screen for 2.5 sec and were replaced by their morph image for 1 sec. Following the morph image, the words “Left or Right?” appeared on the screen. The participants were told that their task was to decide whether the morph image more strongly resembled the typical or the atypical parent image. The participants indicated their decisions by pressing the key labeled *left* or the key labeled

right on the keyboard. The participants were given unlimited time to make their responses. Trials were separated by a 2-sec ITI. Each of the 16 bird, car, and face pairs and morph faces were presented randomly, twice, for a total of 96 experimental trials. The left and right positions of the typical and atypical bird, car, and face images were counterbalanced. The participants were tested on Macintosh computers using the SuperLab experimental package.

Results and Discussion

The 50/50 morph was judged to be more similar to the atypical parent than to the typical parent for birds, cars, and faces on 66%, 58%, and 56% of the trials, respectively. An ANOVA showed that the magnitude of the atypicality bias differed between the bird, car, and face categories [$F(2,33) = 5.395, p < .001$]. The bird stimuli showed a greater atypicality bias than the cars and faces ($p < .01$). The atypical bird preference was reliably above the chance level of 50% by participants [$t(34) = 8.08, p < .0001$] and by items [$t(15) = 10.52, p < .0001$]. The preference for the atypical car was significantly above the level of 50% by participants [$t(34) = 3.50, p < .01$] and by items [$t(15) = 2.56, p < .05$]. For faces, the atypical face was reliably preferred according to participants [$t(34) = 2.94, p < .01$], but not by items [$t(15) = 1.57, p > .10$].

In Experiment 3, we found that participants judged 50/50 morphs of cars and birds as showing a stronger resemblance to their atypical than to their typical parent object. The bird and car results showed that the atypical bias is not exclusive to faces but generalizes to other object categories. Results from Experiment 3 also replicated previous findings (Tanaka et al., 1998), demonstrating that the 50/50 face morph was judged as more similar to the atypical than to the typical parent face. Although an atypicality bias was not demonstrated by items, this may be due to the smaller number of faces tested in this experiment in comparison with the number tested in previous studies (Tanaka et al., 1998). An important follow-up study to Experiment 3 would be to test whether atypical birds and cars exhibit a disadvantage in perceptual discrimination similar to the effect shown for atypical faces.

GENERAL DISCUSSION

In Experiments 1 and 2, typicality effects in face perception were investigated in a same/different discrimination task. On the basis of the attractor field model, it was predicted that if face morphs more strongly resemble the atypical than the typical parent image, it should be more difficult to detect differences between atypical morphs than differences between typical morphs. Consistent with the attractor field prediction, participants were less sensitive to differences in atypical morphs than to differences in typical morphs, as tested in a sequential and simultaneous same/different discrimination task. In Experiment 3, tests of typicality effects were extended to the perception of nonface objects. The participants were presented with morphs generated from pairs of atypical and typical cars, birds, and faces. The main result was that bird, car, and face morphs were judged to be more similar to the atypical than to the typical parent image. Hence, the atypicality bias was not specific to the faces, but extended to the perception of objects that are from structurally homogenous categories (e.g., cars and birds).

These results were predicted by Krumhansl's distance-density theory (1978), in which she claimed that perceptual discrimination was influenced not only by the physical distance between exemplars, but by the location of the exemplars in the similarity space. In the attractor field model, the distance-density relationship is represented in terms of the size of an exemplar's attractor field. Typical representations located in a dense subregion have smaller attractor fields and are therefore more sensitive to differences in their inputs than are atypical representations that lie in a sparser subregion with larger attractor fields (as shown in Figure 1). The attractor model would predict that the slope of the attractor field will be steeper for typical representations than for atypical representations, assuming that the depth of the attractor basins are equal and the basins are shaped like Gaussian bowls.¹ However, simulations of the typicality effect have shown that the attractor basins are shaped more like flat-bottomed bowls, and the settling rates of inputs into typical and atypical attractors at the boundaries of the basins are similar (M. S. Bartlett & Tanaka, 2007). Behaviorally, attractor field slopes could be determined more precisely by using finer gradations of morph levels than the 10% increments used in the present experiments.

In contrast to the typicality advantage reported in the present experiments, Kuhl (1991) and colleagues (Iverson & Kuhl, 1995) demonstrated an opposite bias in speech perception. When adult listeners identified and rated the similarity of stimulus pairs of vowel sounds, they found evidence of what they referred to as a *perceptual magnet effect*, in which there was reduced sensitivity between good exemplars of a vowel category and better differentiation between poorer vowel exemplars. What might account for the *enhanced* discrimination of typical exemplars in our studies and the *reduced* discrimination of typical vowel exemplars in the speech perception experiments? An important distinction between vision and speech is that in visual recognition, the goal is to individuate a given face (e.g., Brad Pitt) or object (e.g., a Honda Civic) from other

category representations. The demands of individuation draw the observer's attention to the second-order relational properties that distinguish exemplars (Tanaka, 2001). The goal of speech perception, on the other hand, is to categorize slightly different speech tokens into the same phonetic category and to ignore phonetically irrelevant variations in the speech signal. The contrasting goals of individuation and categorization might help to account for the enhanced sensitivity to typical representations in vision and reduced sensitivity to typical representations in speech.

In summary, the present findings add to a growing literature demonstrating the effects of category learning on visual perception (Corneille, Goldstone, Queller, & Potter, 2006; Goldstone, 1994, 1998). The present work emphasizes the influence that category structure exerts on the perception of category exemplars. These studies show that the perception of a face or an object depends not only on its fit to an underlying representation, but on the representation's neighbors in the similarity space.

AUTHOR NOTE

This research was supported by grants from the James S. McDonnell Foundation (Perceptual Expertise Network) and the National Science and Engineering Research Councils of Canada to J.W.T. and grants from the Swiss National Science Foundation and the European Science Foundation to O.C. We thank Jesse Lanz and Dan Weiskopf for their assistance with the research. We are grateful to Justin Kantner, Marni Bartlett, and four anonymous reviewers for providing thoughtful comments on the manuscript. Correspondence concerning this article should be addressed to J.W. Tanaka, Department of Psychology, University of Victoria, P.O. Box 3050, STN CSC, Victoria, BC, V8W 3P5 Canada (e-mail: jtanaka@uvic.ca).

NOTE

1. We thank an anonymous reviewer for making this observation.

REFERENCES

- BARTLETT, J. C., HURRY, S., & THORLEY, W. (1984). Typicality and familiarity of faces. *Memory & Cognition*, *12*, 219-228.
- BARTLETT, M. S. (in press). Information maximization in face processing. *Neurocomputing*.
- BARTLETT, M. S., & TANAKA, J. W. (2007). *An attractor field model of face representation: Effects of typicality and image morphing*. Manuscript submitted for publication.
- BUSEY, T. A. (1998). Physical and psychological representations of faces: Evidence from morphing. *Psychological Science*, *9*, 476-483.
- CORNEILLE, O., GOLDSTONE, R. L., QUELLER, S., & POTTER, T. (2006). Asymmetries in categorization, perceptual discrimination, and visual search for reference and nonreference exemplars. *Memory & Cognition*, *34*, 556-567.
- DIAMOND, R., & CAREY, S. (1986). Why faces are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107-117.
- GOLDSTONE, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178-200.
- GOLDSTONE, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585-612.
- IVERSON, P., & KUHL, P. K. (1995). Mapping the perceptual magnet effect for speech using signal detection theory and multidimensional scaling. *Journal of the Acoustical Society of America*, *97*, 553-562.
- JOLICŒUR, P., GLUCK, M. A., & KOSSLYN, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, *16*, 243-275.
- KANWISHER, N. (2000). Domain specificity in face perception. *Nature Neuroscience*, *3*, 759-763.
- KRUMHANSL, C. L. (1978). Concerning the applicability of geometric

- models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, **85**, 445-463.
- KUHL, P. K. (1991). Human adults and human infants show a "perceptual magnet effect" for the prototypes of speech categories, monkeys do not. *Perception & Psychophysics*, **50**, 93-107.
- LIGHT, L. L., KAYRA-STUART, F., & HOLLANDER, S. (1979). Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning & Memory*, **5**, 212-228.
- MURPHY, G., & BROWNELL, H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **11**, 70-84.
- TANAKA, J. W. (2001). The entry point of face recognition: Evidence for face expertise. *Journal of Experimental Psychology: General*, **130**, 534-543.
- 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.
- TANK, D. W., & HOPFIELD, J. J. (1987). Collective computation in neuronlike circuits. *Scientific American*, **257**, 104-114.
- TARR, M. J., & GAUTHIER, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience*, **3**, 764-769.
- TARR, M. J., & WILLIAMS, P. (1996). RSVP: A software package for experimental psychology using the Apple Macintosh OS. Providence, RI: Brown University, Department of Cognitive and Linguistic Sciences.
- 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.
- VALENTINE, T., & BRUCE, V. (1986). The effect of race inversion and encoding activity upon face recognition. *Acta Psychologica*, **61**, 259-273.
- VOKEY, J. R., & READ, J. D. (1992). Familiarity, memorability, and the effect of typicality on the recognition of faces. *Memory & Cognition*, **20**, 291-302.
- WOLBERG, G. (1990). *Digital image warping*. Los Alamitos, CA: IEEE Computer Science Society Press.

(Manuscript received January 26, 2006;
revision accepted for publication October 13, 2006.)