http://www.journalofvision.org/content/11/8/15

The structure of face–space is tolerant to lighting and viewpoint transformations

Idan Blank Department of Psychology, Tel Aviv University, Israel

Galit Yovel Department of Psychology, Tel Aviv University, Israel

According to the face–space framework, faces are represented as locations in a multidimensional space, where the distance separating representations is proportional to the degree of dissimilarity between faces. The present study tested whether similarities between faces, and thus the structure of face–space, were tolerant to ("invariant" under) identity-preserving transformations such as changes in lighting or view. To examine the correspondence between the configurations of face–space under different transformations, perceived similarity was rated for two variants of a set of faces, differing either in illumination (Experiment 1) or viewpoint (Experiment 2). We found that similarity ratings within the first variant were highly correlated with ratings within the second variant. In addition, based on these ratings, a separate face–space was constructed for each variant using multidimensional scaling. Procrustean analysis revealed that the different spaces shared comparable structures. This correspondence serves as a face–space manifestation of the tolerance of identity representations. Accordingly, we suggest that tolerance may rely on the fact that similarities between faces under one transformation are isomorphic to similarity patterns under a different transformation. Thus, recognizing faces under varying viewing conditions may only require similarity evaluations within—rather than across—different transformations.

Introduction

The face–space framework is a theoretical account of the cognitive representation of faces (Valentine, 1991). According to this model, face representations are iso-morphic to locations in a multidimensional psychological space. The dimensions spanning this space are assumed to encode physical or abstract attributes that render different faces discriminable from one another (e.g., Busey, 1998; Hancock, Burton, & Bruce, 1996). Hence, face–space is naturally endowed with a dissimilarity-based metric (at least locally; see Craw, 1995). In other words, the distance separating representations in face–space is proportional to the degree of dissimilarity between faces. This property serves as a structural embodiment of the fundamental notion of the framework: face processing is established upon similarity evaluations (also see Nosofsky, 1992a, 1992b; Shepard, 1987).

Impressively, this similarity-based framework proposes a unified account for a range of face-recognition phenomena, such as the effects of distinctiveness (Tanaka, Giles, Kremen, & Simon, 1998; Valentine & Endo, 1992), caricaturing (Lee, Byatt, & Rhodes, 2000; Lewis & Johnston, 1999), race (Byatt & Rhodes, 2004; Chiroro & Valentine, 1995; but also see Levin, 1996; Lewis, 2004), gender (Campanella, Chrysochoos, & Bruyer, 2001), age (Johnston, Kanazawa, Kato, & Oda, 1997), and perceptual

distortion following adaptation (Leopold, O'Toole, Vetter, & Blanz, 2001; Rhodes & Jeffery, 2006) (for a review, see Valentine, 1995, 2001). This ability to account for such diverse effects, otherwise accounted for by unrelated mechanisms, renders the face–space one of the most widely acknowledged and comprehensive models of face processing.

However, both the framework's theoretical basis and its empirical support have largely overlooked one crucial property of face processing: its tolerance to ("invariance" under) varying viewing conditions, such as changes in the illumination, viewpoint, size, or position of a face (Edelman, 1999; Moses, Ullman, & Edelman, 1996; Rolls, 2000; Zoccolan, Kouh, Poggio, & DiCarlo, 2007). This property refers to the visual system's ability to recognize the same face in different images, i.e., to relatively compensate for "identity-preserving transformations" so as to extract the identity of a face, ignoring its "accidental" changes in appearance (Hasselmo, Rolls, Baylis, & Nalwa, 1989; Rolls & Baylis, 1986; Tovee, Rolls, & Azzopardi, 1994), as well as to generalize to novel viewing conditions (e.g., Moses et al., 1996). Being fundamental to face (and object) representation, the little attention given to the tolerance property in the face–space context is surprising; whereas Newell, Chiroro, and Valentine (1999) have suggested that tolerance was somehow embedded in the structure of the space (also see Eifuku, De Souza, Tamura, Nishijo, & Ono, 2004), this hypothesis remains an implicit

assumption (e.g., Leopold, Bondar, & Giese, 2006; Valentine, 2001) that, to the best of our knowledge, has not been tested empirically.

Indeed, the association of a similarity space with tolerance to identity-preserving transformations is not elementary. To illustrate this, think of representing two individuals, Jim and Dan, in face–space: Intuitively, one might reason that different images of Jim ought to be represented as more similar to each other than to images of Dan. This is, presumably, the condition for a similarity-based representation to be tolerant to changes in viewing conditions. Nevertheless, from an image-based perspective, two images of Jim under different viewing conditions might be less similar than the images of Jim and Dan viewed under the same illumination or viewpoint (Adini, Moses, & Ullman, 1997). In this case, some representations of Jim in face–space might be closer to some representations of Dan than to each other. Under such circumstances, it seems that similarity evaluations—and the framework itself—lose their aptness as an account of face processing, of which tolerance is a main characteristic.

Fortunately, this is but an illusory problem. One can reconstruct the theoretical link between face–space and tolerance to achieve a straightforward solution. Specifically, demanding that "images of Jim are more similar to each other than to images of Dan" is an unnecessary condition. In fact, we might need to abandon it altogether. To understand this, we follow Shepard's notion of "second-order isomorphism" (Shepard, 1968; Shepard & Chipman, 1970): Shepard claims that a correspondence should exist not between an object and its cognitive representation but between similarities of objects and similarities of representations. Thus, a representation of Jim's face need not relate in any way to Jim's actual face—rather, the similarity between Jim and Dan's representations should relate to the similarities between their actual faces. Note that, given the variability in Jim's images induced by varying viewing conditions, Shepard's argument does not imply different images of Jim having similar representations. However, it does allow the following: The similarity between Jim and Dan's representations under one viewing condition resembles the similarity between their representations under a different viewing condition.

By extension, if, under one viewing condition, Jim and Dan's representations are more similar than Jim and Joe's, this pattern of similarities should also be observed under a different viewing condition. This common similarity pattern across viewing conditions can serve as a face–space correlate of the fact that all of Jim's images belong to the same individual, even when their locations are very distant from each other (i.e., Jim's images are themselves very dissimilar). In other words, the structure of face–space may itself exhibit tolerance to identity-preserving transformations (also see Newell et al., 1999): If representation is representation of similarities (Edelman, 1998),
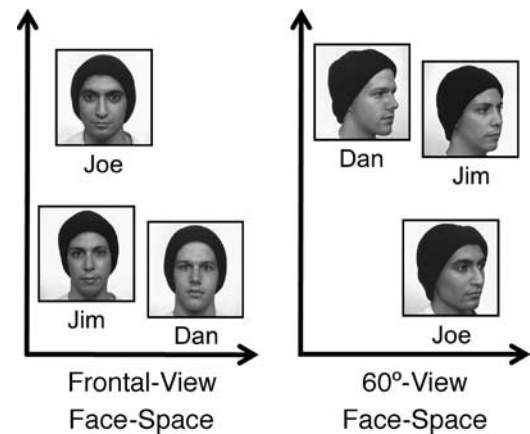

Figure 1. An illustration of our hypothesis that similarity patterns in face–space exhibit tolerance to identity-preserving transformations. Two spaces are shown, each representing a different transformation: (left) frontal viewpoint and (right) 60° viewpoint. In both spaces, Jim and Dan are more similar to each other than Jim and Joe (names are for illustration purposes). This common pattern of similarities preserves the structure of face–space across the two viewpoints. Note that the dimensions of the two spaces need not necessarily correspond—only the relative location of a representation with respect to other representations is preserved.

we hypothesize that "invariance" is "invariance of similarities" (Figure 1).

The goal of the present study was, therefore, to examine the tolerance of similarity patterns in face–space to identity-preserving transformations. To quantify the structure of the human face–space, we first collected subjects' inter-item perceived similarity ratings for a set of facial stimuli. These ratings were then converted to a spatial arrangement of the images in a concrete face–space, via multidimensional scaling (MDS; Shepard, 1957, 1980). To only evaluate similarities *within* different viewing conditions, ratings were made separately for two variants of the same stimuli set differing either in illumination (Experiment 1) or viewpoint (Experiment 2). In each experiment, configurations in face–space were separately generated for each variant of the stimuli set, and the correspondence between these configurations was evaluated.

Experiment 1

Experiment 1 tested the tolerance of similarity patterns in face–space to illumination changes. If similarity patterns are tolerant to lighting transformations, then faces that are relatively similar under one lighting condition should remain so under a different lighting condition, and

faces that are relatively dissimilar should thus remain dissimilar under a change in illumination. Therefore, we examined whether the similarity pattern (i.e., relative location) of a face in the space of frontally-lit faces corresponded to its location in the space of top-lit faces.

Methods

Participants

Twenty-four Caucasian subjects (10 males) volunteered to participate in the study. Age ranged from 18 to 24 years ($M = 20.8$, $SD = 3.8$). All reported normal or corrected-to-normal vision and were not familiar with the face stimuli used in the study.

Stimuli

Digitized photographs (300 × 300 pixels, 256 gray levels) of 36 Caucasian adult males from the Harvard Face Database were used as stimuli. Two photographs of each face were included, creating two sets: frontally-lit (FL) and top-lit (TL). All faces were of frontal view, neutral expression, free from external features (e.g., facial hair, glasses), and had their hair covered with a ski hat.

Stimuli were presented centered on a 17″ computer screen and subtended a visual angle of 7.36° × 7.36° (width by height) at a viewing distance of 60 cm. Stimulus presentation and response recording were controlled by MATLAB (The MathWorks), using the Psychophysics Toolbox (Brainard, 1997).

Design and procedure

Each subject was randomly assigned to either the Frontal Lighting or Top Lighting condition and was presented accordingly with faces from either the FL or TL variants. In each trial, a centered fixation point appeared for 750 ms, followed by a sequential presentation of two faces. Each face appeared for 1 s, with an inter-stimulus interval of 500 ms. After the second face disappeared, the subject was asked to press a key corresponding to the level of perceived similarity between the two faces, on a scale ranging from 1 (identical) to 7 (extremely different). The next trial started 1 s after a response was made (Figure 2). Subjects were encouraged to give their first impression, but the duration of trials was not limited. In addition to similarity ratings, reaction time (RT) was also recorded. Prior to the experiment, subjects were presented with a 25-trial practice session, including 5 faces not used in the experiment itself.

Each subject rated a total of randomly presented 1332 pairs, which included every possible pairing of faces in both presentation orders and two pairings of every face with itself. The session was split into 5 blocks, each lasting approximately 20 min. Upon completion of the
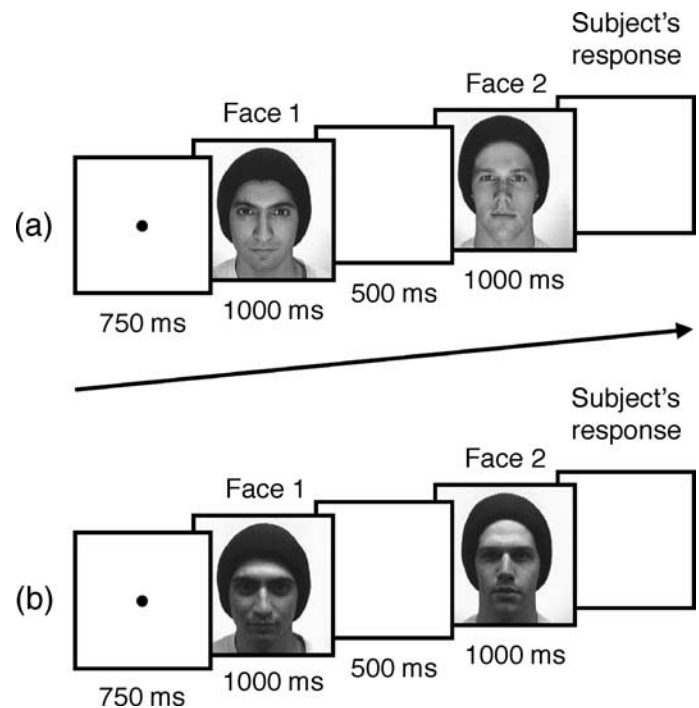

Figure 2. The perceived similarity rating task used in Experiment 1. (a) The frontal lighting (FL) condition. (b) The top lighting (TL) condition.

experiment, subjects were debriefed by the experimenter as to the purpose of the study.

Results

Correspondence of perceived similarity ratings across lighting transformations

Prior to data analysis, outliers—identified as trials in which RT exceeded the subject's mean RT ± 2.5 SD—were removed (2% of the trials in the FL condition and 1.9% of the trials in the TL condition). In each experimental condition, the remaining perceived similarity ratings given to each pairing of different faces were averaged across subjects and order of presentation to produce a set of 630 Mean Similarity Ratings (an MSR matrix). The ratings given to identical pairs were not considered further.

To roughly evaluate whether similarity relations were preserved under the two lighting conditions, we first tested the correlation between corresponding inter-item similarities across the FL and TL MSR matrices. This analysis was repeated for 36 data samples, each including 35 independent pairs of corresponding similarity ratings, revealing significant correlations (mean Spearman's $r = 0.60$, $SD = 0.16$, $t_{(33)} = 4.3$, $p < 10^{-4}$).

Next, tolerance was more accurately assessed by examining the correspondence between similarity patterns of single faces across illumination changes. To this end,
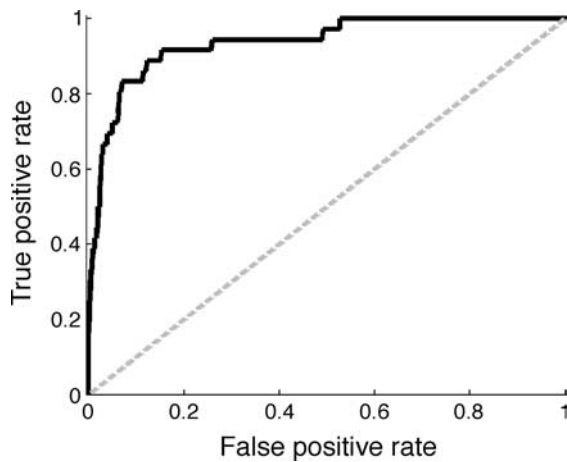
Figure 3. ROC curve for Spearman's rank correlation coefficients of FL and TL similarity patterns (pairs of the "same" patterns vs. pairs of "different" patterns). Area under the curve (AUC) = 0.94. Dashed gray line indicates chance performance.

we extracted from the FL and TL MSR matrices: (1) the similarity pattern (i.e., similarity vector) of every FL face to the remaining 35 FL faces and (2) the similarity pattern of every TL face to the remaining 35 TL faces. Spearman's rank correlation coefficient was then calculated for every pair of FL and TL similarity patterns: this included 36 "same" pairs (similarity patterns of the same face across lighting conditions) and 1260 "different" pairs (an FL similarity pattern of one face paired with a TL similarity pattern of a different face). To test whether the "same" patterns corresponded more than "different" patterns, we plotted the Receiver Operating Characteristic (ROC) curve (Figure 3) for the Spearman correlation coefficients. The area under the curve (AUC) was 0.94, indicating an almost perfect ability to infer whether two FL and TL faces were the "same" or "different" based on their similarity patterns.

Multidimensional scaling and Procrustes analysis

In order to derive a face–space configuration corresponding to the perceived similarity ratings, the MSR matrix of each experimental condition was submitted to MATLAB's non-metric (Kruskal, 1964; Shepard, 1966) multidimensional scaling procedure MDSCALE. The metric parameter was specified as Euclidian (e.g., Johnston et al., 1997; Lee et al., 2000; if invariance is intrinsic to face–space, it should be evident with any metric: Craw, 1995). Since the dimensionality of the spaces underlying the observed similarity ratings is a mere speculation, FL and TL MDS solutions were generated in 2 to 15 dimensions, showing a decrement in stress (Shepard, 1980) with increasing dimensionality. Each pair of configurations was analyzed for structural correspondence, which enabled us to exclude the possibility that our
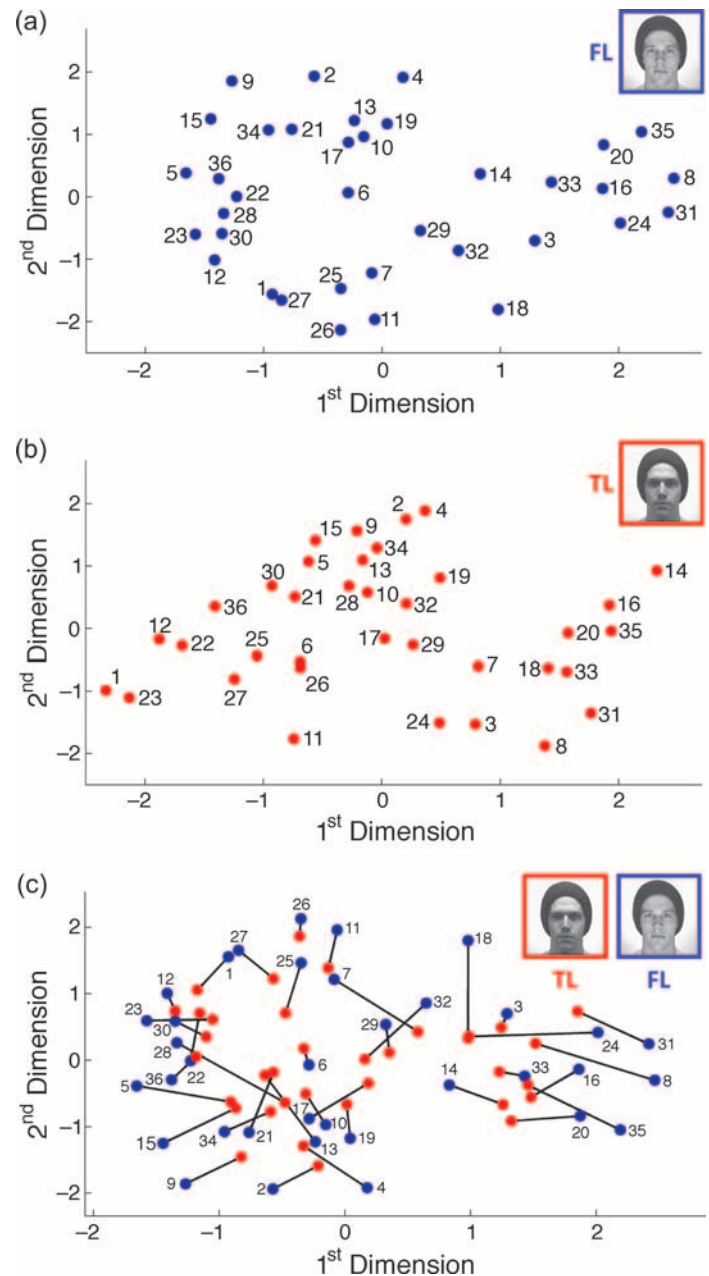

Figure 4. (a) Stimuli locations on the first two dimensions of the 10-dimensional space of frontally-lit (FL) images (stress = 0.049). (b) Stimuli locations on the first two dimensions of the 10-dimensional space of top-lit (TL) images (stress = 0.047). Since the choice of dimensions for the MDS solution is random, the configurations in (a) and (b) are plotted following principal component analysis (PCA). (c) The configurations of FL images (blue) and the Procrustes-transformed configuration of TL images (red) superimposed. Locations connected by a line correspond to FL and TL variants of the same facial identity and are marked by a number. Note that TL and FL variants of the same face have relatively similar locations, as indicated by the short connecting lines.

results were dependent upon the specific dimensionality of the MDS solutions.

The correspondence between the TL and FL configurations was quantified using Procrustes analysis (Borg & Lingeos, 1987; Gower, 1975) as implemented in MATLAB. This descriptive analysis minimizes the sum of squared residuals between the point values of the two spaces, by transforming one configuration (e.g., the TL space) to optimally fit the other (FL) configuration. The transformation is a combination of scaling, orthogonal rotation, reflection, and translation and, hence, does not affect the shape of the transformed configuration. Figure 4 shows the distribution of faces on the first two dimensions of the 10-dimensional TL and FL spaces prior to, and following, this transformation (the 10-dimensional spaces replicated well the similarity data, having stress values of 0.049 (FL) and 0.047 (TL), exceeding Kruskal and Wish's (1978) criterion of 0.05 for a good fit).

We tested the significance of the correspondence between the spaces with a PROcrustean randomization TEST (PROTEST; Jackson, 1995), using 9999 permuted spaces based on resamplings of point values in the TL space. Each permuted space was Procrustes-transformed to fit the original FL space, and badness of fit was measured using the d statistic—a standardized sum of squared residuals between the two spaces. The resulting d values (Figure 5) indicated that the fit between our two original spaces was highly significant ($p = 10^{-4}$ for all dimensionalities). Comparable results were found using 9999 permuted MSR matrices, instead of permuted spaces (following Cutzu & Edelman, 1998).

Finally, in order to assure that differences between subjects' individual rating tendencies did not influence our results, all analyses were performed on perceived similarity ratings that were subjected to intra-subject ranking. Analysis of the ranked data yielded comparable results to those described above (Figure 5).

Discussion

Our findings indicate that the face–space configuration of frontally-lit face images is similar to that of the corresponding top-lit image variants. These results suggest that "invariance" under some illumination transformations is reflected in the structure of face–space, resulting from the tolerance of similarity relations to changes in lighting conditions. It should be noted, however, that the two configurations were each based on ratings obtained from a different group of individuals and were thus also averaged across subjects. This design has not allowed us to evaluate the within-subject correspondence of two individual spaces. In addition, having examined only one class of transformations, the ability to generalize our results to other identity-preserving transformations was limited. Experiment 2 was designed to address these issues directly.

Experiment 2

Experiment 2 employed a within-subject design, whereby every subject rated two variants of a stimulus set, each in a separate session. Furthermore, tolerance was studied not under illumination changes but rather under changes in viewpoint.
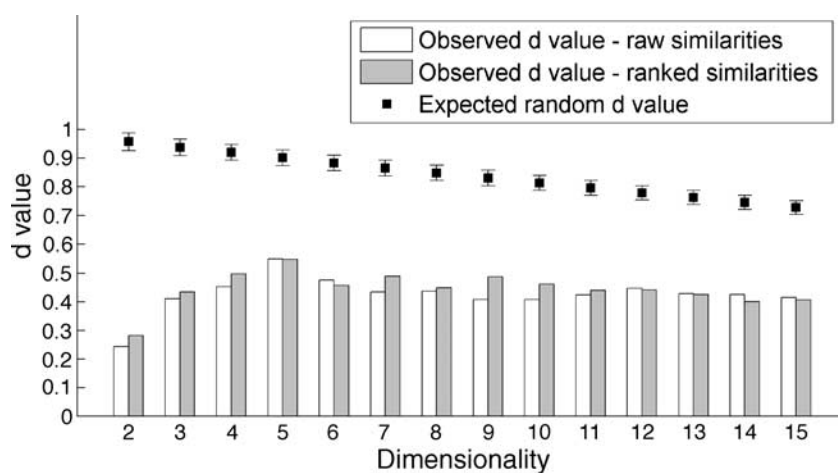

Figure 5. Results of Procrustean analysis for comparing FL and TL configurations in 2 to 15 dimensions. The observed d values, indicating the badness of fit of the TL space to the FL space, are plotted both for the raw similarities (white bars) and for similarities following intra-subject ranking (gray bars). The plot also shows the expected d values under the null hypothesis (mean ± SE), generated by PROTEST with 9999 random data permutations (these values were almost identical for the raw data and ranked data).

Methods

Participants

Twelve undergraduate students in the Department of Psychology, Tel-Aviv University (2 males) who did not participate in Experiment 1 participated in the study, receiving credit toward a course requirement. One subject was removed from the analysis due to long RTs and apparent misunderstanding of instructions. The age of the remaining 11 subjects ranged from 22 to 25 years ($M = 23.7$, $SD = 1$). All reported normal or corrected-to-normal vision and were not familiar with the faces used as stimuli.

Stimuli

To avoid the laborious rating session of 1332 trials, we used a subset containing 24 of the 36 faces presented in Experiment 1 (analysis of the data collected in Experiment 1 for this subset revealed comparable findings to those found for the entire set of 36 faces). Two frontally-lit photographs of each of these 24 faces were now included, one presented from a frontal viewing angle (V0) and another presented from a right 60° viewing angle (V60; see Figure 1). All other characteristics of the stimuli and experimental parameters were identical to those in Experiment 1, except the location of the second stimulus that was displaced relative to the first stimulus (175 pixels lower and 350 pixels to the right). This was done to ensure that subjects compared high-level face representations rather than rated iconic picture similarity.

Design and procedure

The procedure was similar to that of Experiment 1, with the exception that each subject rated perceived similarities both within the V0 variant and within the V60 variant. To minimize effects of familiarity, these variants were viewed in two different sessions, 3 weeks apart: half of the subjects were presented with the V0 variant in the first session followed by the V60 variant in the second session (V0-1st/V60-2nd), and the other half were presented with the variants in the opposite order (V60-1st/V0-2nd). Each session consisted of a total of 600 pairs of faces and lasted 45 min.

Results

Group analysis

Outlier removal was carried out in the same fashion as in Experiment 1 (2.4% of the trials in the V0 condition and 1.8% of the trials in the V60 condition). Before averaging data across sessions to create one MSR matrix for each variant, we sought to confirm that familiarity with the stimuli had not significantly affected subjects' ratings in the 2nd session. Therefore, we tested the correspondence between the average V0-1st and V0-2nd MSR matrices, as well as between the average V60-1st and V60-2nd MSR matrices: If familiarity had a negligible effect, subjects who had been exposed to one of the variants during the 1st session (when the facial identities were still unfamiliar) should have rated perceived similarity in concordance with subjects who were exposed to the same variant during the 2nd session (when the identities might have been familiar). MDS solutions in 2 to 15 dimensions were generated based on the MSR matrices, following the considerations outlined in Experiment 1. Procrustean analyses revealed that across the two sessions both the V0 and V60 variants shared similar face–space configurations (Table 1). Thus, data were collapsed across sessions (and subjects), resulting in a single V0 MSR matrix and a single V60 MSR matrix.

We found that similarity relations were highly preserved across the two viewpoints: The correlation between corresponding inter-item similarities across the two MSR matrices, calculated for 24 data samples each including 23 independent pairs of corresponding similarity ratings, was highly significant (mean *Spearman's* $r = 0.63$, $SD = 0.14$, $t_{(21)} = 3.75$, $p < 0.001$). In addition, as in Experiment 1, an ROC curve was plotted for the Spearman correlations between pairs of V0 and V60 similarity patterns of single faces. We found that the correspondence between pairs of the "same" similarity patterns was higher than that of "different" similarity patterns (AUC = 0.91).

Condition	Dimensionality	MDS stress[a] (1st/2nd session)	Observed Procrustean *d* value	PROTEST expected random *d* value (mean $\pm$ SE)
V0	2	0.227/0.236	0.543**	0.935 ± 0.046
	8	0.053/0.049	0.463**	0.768 ± 0.038
	15	0.007/0.005	0.333**	0.581 ± 0.031
V60	2	0.233/0.241	0.837*	0.936 ± 0.045
	8	0.049/0.042	0.512**	0.768 ± 0.038
	15	0.007/0.004	0.343**	0.587 ± 0.032

Table 1. Comparison of ratings from different sessions in Experiment 2. *Note:* [a]A stress value smaller than 0.05 indicates a good fit (Kruskal & Wish, 1978). *$p < 0.05$. **$p < 10^{-4}$.
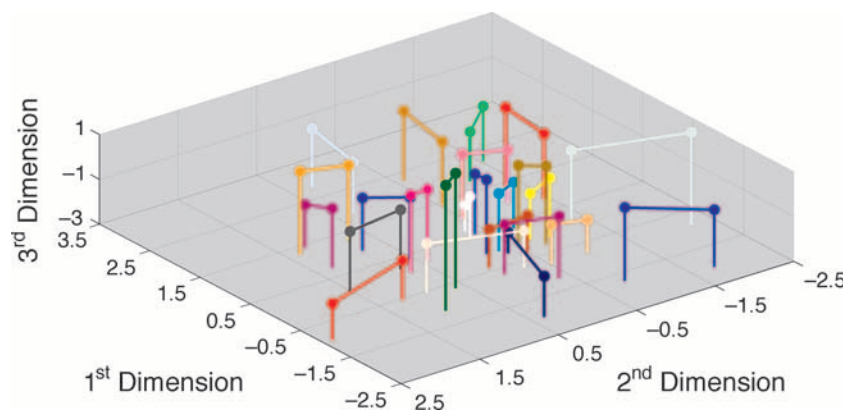
Figure 6. A scatterplot superimposing the 8-dimensional configuration of frontal view (V0) images and the Procrustes-transformed 8-dimensional configuration of 60° view (V60) images. Only the first three dimensions are shown. Locations connected by a horizontal line correspond to V0 and V60 variants of the same facial identity.

Next, V0 and V60 MSR matrices were converted to MDS configurations in 2 to 15 dimensions, and each pair of V0 and V60 spaces was submitted to Procrustes analysis. Figure 6 shows the distribution of faces on the first three dimensions of the 8-dimensional V0 (stress = 0.045) and V60 (stress = 0.041) spaces following this analysis. PROTEST confirmed that the degree of correspondence between the V0 and V60 configurations was highly significant ($p = 0.0092$ for the 2-dimensional configurations; $p = 10^{-4}$ for all other dimensionalities; Figure 7).

Within-subject analysis

Comparison of V0 and V60 similarity ratings was also carried out for each subject separately, using the same analyses as described above. When testing the 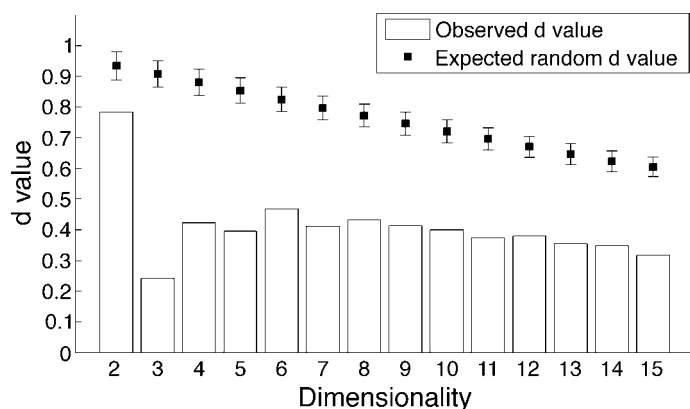overall correspondence of individual MSR matrices across experimental conditions, it seemed that similarities did not exhibit much tolerance to viewpoint changes: The correlation between corresponding V0 and V60 inter-item similarities was relatively weak, with a mean *Spearman's r* of 0.30 ($SD = 0.08$) across subjects (according to Fisher's combination of experiments rule: $\chi^2_{(22)} = 56.53$, $p < 10^{-4}$). However, an ROC curve based on subjects' individual data revealed that the similarity patterns of single faces exhibited a high degree of tolerance: Pairs of the "same" V0 and V60 similarity patterns were more correlated than pairs of "different" V0 and V60 similarity patterns (AUC = 0.72; Figure 8). This result indicated that, based on similarity patterns obtained from an individual subject, it was possible to discriminate between the "same" the "different" faces across viewpoints.

Subjects' individual MSR matrices were also converted to MDS configurations in 2 to 15 dimensions and analyzed as previously described. PROTEST confirmed that the

Figure 7. Results of Procrustean analysis for comparing V0 and V60 configurations in 2 to 15 dimensions. Plotted are observed *d* values, indicating the badness of fit of the V0 space to the V60 space (white bars). The plot also shows the expected *d* values under the null hypothesis (mean ± *SE*), generated by PROTEST with 9999 random data permutations.
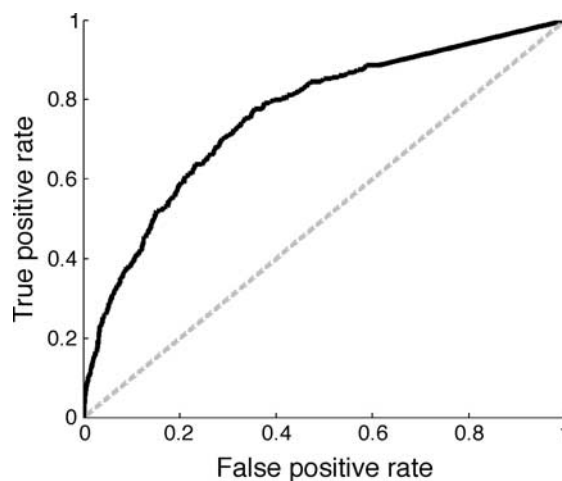

Figure 8. ROC curve for Spearman's rank correlation coefficients of V0 and V60 similarity patterns (pairs of the "same" patterns vs. pairs of "different" patterns). Area under the curve (AUC) = 0.72. Dashed gray line indicates chance performance.
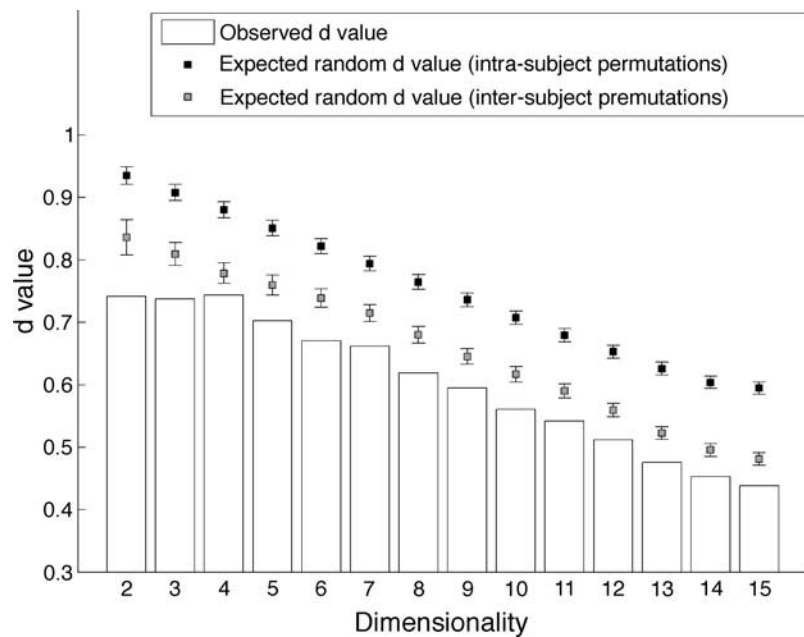
Figure 9. Results of Procrustean analysis for comparing individual V0 and V60 configurations in 2 to 15 dimensions. The plot shows observed *d* values, indicating the badness of fit of the individual V0 spaces to the V60 spaces (white bars); expected *d* values (mean ± *SE*), for testing whether the fit is better than expected by chance (black), generated by PROTEST with 9999 random intra-subject permutations of the data; and expected *d* values (mean ± *SE*), for testing whether the fit between spaces of the same subject is better than the fit between spaces of two different subjects (gray), generated by PROTEST with 9999 random inter-subject permutations of the data.

degree of correspondence between individual V0 and V60 configurations was highly significant ($p = 10^{-4}$ for all dimensionalities; Figure 9). In addition, a separate analysis was performed to test whether the V0 and V60 spaces of the same subject were more concordant than V0 and V60 spaces of two different subjects. This was carried out with PROTEST using 9999 "inter-subject" permutations, assigning to each individual V0 space a random individual V60 space. Although only marginally significant for lower dimensionalities, we found that spaces of the same subject indeed had a higher correspondence than inter-subject spaces ($p < 0.001$ for dimensionalities higher than 4; Figure 9).

Comparing the effects of illumination and view

To further appreciate the pattern of tolerance reflected in the structure of face–space, we used Procrustean analysis to compare each of the V0 and V60 spaces with each of the FL and TL spaces (constructed for the subset of 24 FL or TL stimuli presented in Experiment 1). The results, which are purely descriptive, suggest that the FL and TL spaces—both representing configurations of frontal pose stimuli—are more similar to the V0 than to the V60 space. In addition, the V0 and V60 spaces—both representing configurations of relatively frontally-lit stimuli—are more similar to the FL than to the TL space. These results are in line with the physical changes induced by the illumination and viewpoint transformations, as

measured by the Euclidean distance separating the stimuli in pixel space (Figure 10).

Discussion

The results of Experiment 2 show that the face–space configuration of frontal view face images is similar to that of the corresponding 60° view image variants. Thus, they extend the tolerance of similarity relations in face–space, previously observed for illumination changes, to another identity-preserving transformation—that of viewpoint. Such high degree of tolerance is evident when comparing both group-averaged configurations and individual (i.e., within-subject) configurations, thereby confirming that the results of Experiment 1 cannot be entirely explained in terms of overestimation caused by averaging.

General discussion

The findings of the current study demonstrate that two variants of a set of faces—differing either in illumination (frontally-lit vs. top-lit) or viewpoint (frontal view vs. 60° view)—share corresponding configurations in face–space. Correspondence between space configurations was found irrespective of dimensionality, both at the group-averaged and individual levels, reflecting the tolerance of
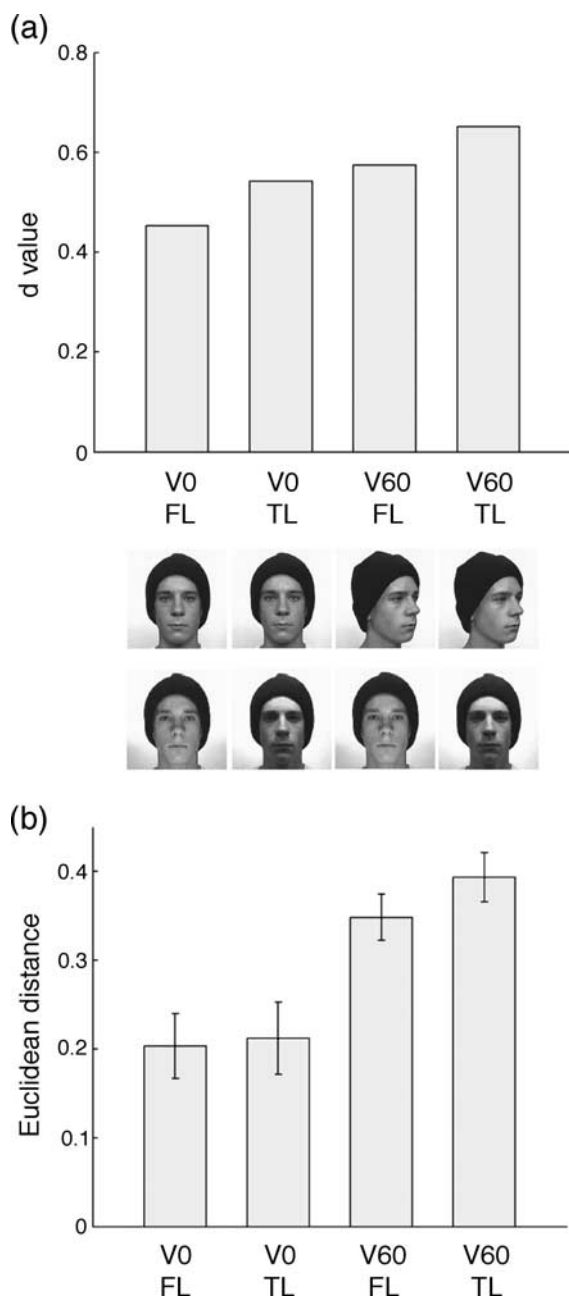
Figure 10. (a) Procrustean *d* values for comparisons of spaces from Experiment 2 (frontal view/60° view configurations) with spaces from Experiment 1 (frontally-lit/top-lit configurations). Data are presented for 9-dimensional configurations and are representative of the pattern of results for the other dimensionalities. Below each comparison of two spaces, one stimulus from each space is shown for illustration purposes. (b) Euclidean distance (mean ± *SD* across 24 stimuli) between images from the experimental conditions compared in (a).

similarity relations to the transformations used. These findings therefore suggest that the "invariance" of identity representations under illumination and viewpoint changes is echoed in the similarity structure underlying face–space configurations.

These results are in line with previous studies showing that representations of visual information in similarity spaces preserve their relative locations across different transformations (e.g., Cutzu & Edelman, 1998). However, the common approach in such studies has been to examine similarities both *within* and *across* viewing conditions: Similarity has usually been measured not only for different objects under the same view but also for different views of the same object. Interpreting the results of such designs is somewhat problematic, since similarity across viewing conditions only measures the behavioral manifestation of tolerance; it is thus not suited for tapping an underlying representation that may, or may not, reflect that tolerance. In other words, if subjects are presented with two different views of Jim, measuring similarity only inform us of their behavioral, surface ability to attribute these images to the same person. Incorporating such similarities into an MDS solution therefore enforces the space to structurally conform to the expectation for tolerance: Different views of the same face will be more closely located than different faces under the same view, giving rise to a space organized by "identity clusters." As the current study aimed to bypass the behavioral manifestation of tolerance, similarity was measured only *within* viewing conditions: subjects never compared two images of the same face across transformations. Hence, the observed tolerance of similarities could not have been directly affected by subjects' knowledge that the frontally-lit and top-lit Jim were indeed the same person.

Moreover, unlike the previous studies concerned with similarity-based representations of non-face (e.g., animal-like) objects, the current study addressed the degree of tolerance specifically evident in face–space. Thus, our finding that similarities *within* one viewing condition correspond to similarities *within* another viewing condition remains to be tested with regard to objects' shape-spaces in future studies. Even though we believe that such tolerance of similarities is fundamentally not face-specific and could be generalized to other object classes, this assumption is not self-evident: Since the processing of objects and faces recruit, to some extent, different cognitive mechanisms (Duchaine, Yovel, Butterworth, & Nakayama, 2006; Robbins & McKone, 2007) and cortical regions (e.g., Kanwisher & Yovel, 2006; also see McKone, Crookes, & Kanwisher, 2009), the principles underlying their organization in similarity spaces might also diverge.

While we interpret the tolerance of similarities revealed in the current study as indicating correspondence between space configurations under different viewing conditions, an alternative account must also be considered. It might be possible that all the spaces constructed in our experiments (FL, TL, V0, and V60) are in fact the same single, "abstract" space of highly tolerant identity representations. A representation based on the extraction of "invariant features," for instance, if available prior to similarity evaluations, would produce the same findings

revealed in our study. In such a case, the different experimental conditions would all tap the same illumination/viewpoint-independent representation, thus resulting in corresponding similarity patterns.

While the current study cannot exclude this alternative account, there is evidence that faces are encoded first using transformation-dependent schemes, e.g., by view-selective neurons (Freiwald & Tsao, 2010; Logothetis & Pauls, 1995; Perrett et al., 1985). Each of these schemes results, by definition, in a separate space of transformation-dependent representations. It might be computationally possible that, across different viewing conditions, such transformation-selective schemes give rise to comparable similarity patterns. These comparable patterns, in turn, induce in face–space a structural correlate of tolerance to identity-preserving transformations.

The observed correspondence of similarities is also evident when examining the similarity patterns of single faces. As revealed by the ROC analysis, it is possible to infer whether two faces under transformations T1 and T2 are the "same" or "different" based on their similarity patterns to other T1 and T2 faces, correspondingly. Thus, although our findings only present a face–space correlate of tolerance, we further draw upon them to theoretically speculate about a causal relationship: It is possible that the configuration of face–space does not only reflect but also gives rise to tolerance with regard to illumination and viewpoint changes. Such emergence of tolerance as a structural property of face–space could be construed through an extension of Edelman's "Chorus of Prototypes" model (Edelman, 1995; Edelman & Duvdevani-Bar, 1997a, 1997b): First, we require the storage of a set of prototypical faces under a variety of transformations, such that face–space is divided into subspaces corresponding each to some transformation (such as our FL, TL, V0, and V60 configurations; see Hasselmo et al., 1989; Newell et al., 1999; this model is plausible at least for view transformations: Perrett et al., 1985; for a discussion, see Rolls, Cowey, & Bruce, 1992). Next, we postulate that a face located in a specific subspace is compared only to the prototypical variants within that subspace (a frontal view face is compared only to frontal view prototypes, whereas a 60° view face is compared only to 60° view prototypes).

Our results imply that such a model may be able to account for the ability to match two unfamiliar faces across transformations, as well as for the recognition of a familiar face under a novel transformation. Specifically, two new faces could be matched for their identity if there were comparable similarity patterns between each of them and its corresponding prototypical variants. Similarly, if the similarity pattern of a new face image (within the appropriate subspace) resembled an existing similarity pattern in a different subspace, then face recognition could exhibit generalization to novel transformations.

Unlike previous implementations of the "Chorus of Prototypes" model (Cutzu & Edelman, 1998; Duvdevani-Bar, Edelman, Howell, & Buxton, 1998; Edelman &

Duvdevani-Bar, 1997a), our theoretical extension—following our ROC analysis—proposes to evaluate similarity patterns of faces *within* a specific transformation, instead of *across* different transformations. The traditional model, by contrast, required the existence of abstract, "invariant" prototypes (Edelman & Duvdevani-Bar, 1997b). Its implementing of an "abstract" representation—i.e., having tolerance established through a preceding processing stage—meant that the model's account of tolerance was external to its own fundamental notions. This was true even though some of the algorithms exploited the finding that a transformation (e.g., rotating a face from a frontal view to a 60° view) induced similar changes in the representations of different faces, i.e., regardless of identity (Edelman & O'toole, 2001; Lando & Edelman, 1995). Our suggestion that tolerance may be established through the evaluation of similarities to prototypes, and may therefore be incorporated into the model as an emergent property of its computational mechanism, further extends the parsimony of the Chorus of Prototypes. Future studies are needed to test the biological plausibility of the storage of prototypes under a variety of transformations, as well as the computational viability of our refined model.

In summary, we suggest here that the tolerance of identity representations to identity-preserving transformations is reflected by, and perhaps causally explained by, the tolerance of similarity patterns in face–space. Tolerance is observed based on similarities between different identities *within* each transformation rather than similarities between faces of the same identity *across* transformations. Consequently, our proposed account of tolerance relies on an indirect comparison of faces, by evaluating the similarity pattern of each face to its corresponding prototypical variants. Such a process does not necessarily require the extraction of "invariant features" from face images: only the differences between faces should exhibit tolerance, giving rise to similar configurations across different subspaces. This implies that facial identity is cognitively a negative entity, lacking a positive intrinsic essence (also see de Saussure, 1983): Faces sharing an identity need not have an inherent invariable quality but must only differ in the same fashion from other faces. This idea of negative identities defined by differences (not by "what they are" but by "what they are not") is evident in many interpretations of visual representation and recognition, yet it is the face–space framework that can most creatively celebrate its strengths.

Acknowledgments

References

Adini, Y., Moses, Y., & Ullman, S. (1997). Face recognition: The problem of compensating for changes in illumination direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 19,* 721–732.

Borg, I., & Lingeos, J. C. (1987). *Multidimensional similarity structure analysis.* Berlin, Germany: Springer.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436.

Busey, T. A. (1998). Physical and psychological representations of faces: Evidence from morphing. *Psychological Science, 9,* 476–483.

Byatt, G., & Rhodes, G. (2004). Identification of own-race and other-race faces: Implications for the representation of race in face space. *Psychonomic Bulletin & Review, 11,* 735–741.

Campanella, S., Chrysochoos, A., & Bruyer, R. (2001). Categorical perception of facial gender information: Behavioural evidence and the face–space metaphor. *Visual Cognition, 8,* 237–262.

Chiroro, P., & Valentine, T. (1995). An investigation of the contact hypothesis of the own-race bias in face recognition. *Quarterly Journal of Experimental Psychology A, 48,* 879–894.

Craw, I. (1995). A manifold model of face and object recognition. In T. Valentine (Ed.), *Cognitive and computational aspects of face recognition: Explorations in face–space* (pp. 183–203). London: Routledge.

Cutzu, F., & Edelman, S. (1998). Representation of object similarity in human vision: Psychophysics and a computational model. *Vision Research, 38,* 2229–2257.

de Saussure, F. (1983). Linguistic value (R. Harris, Trans.). In C. Bally and A. Sechehaye (Eds.), *Course in general linguistics* (pp. 118–121). London: Duckworth.

Duchaine, B. C., Yovel, G., Butterworth, E. J., & Nakayama, K. (2006). Prosopagnosia as an impairment to face-specific mechanisms: Elimination of the alternative hypotheses in a developmental case. *Cognitive Neuropsychology, 23,* 714–747.

Duvdevani-Bar, S., Edelman, S., Howell, A. J., & Buxton, H. (1998). *A similarity-based method for generalization of face recognition over pose and expression.* Paper presented at the 3rd International Symposium on Face and Gesture Recognition (FG98), Washington, DC.

Edelman, S. (1995). Representation, similarity and the chorus of prototypes. *Minds and Machines, 5,* 45–68.

Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Sciences, 21,* 449–498.

Edelman, S. (Ed.) (1999). *Representation and recognition in vision.* Cambridge, MA: MIT Press.

Edelman, S., & Duvdevani-Bar, S. (1997a). A model of visual recognition and categorization. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 352,* 1191–1202.

Edelman, S., & Duvdevani-Bar, S. (1997b). Similarity, connectionism, and the problem of representation in vision. *Neural Computation, 9,* 701–720.

Edelman, S., & O'toole, A. J. (2001). Viewpoint generalization in face recognition: The role of category-specific processes. In M. J. Wegner & J. T. Townsend (Eds.), *Computational, geometric, and process perspectives on facial cognition: Contexts and challenges* (pp. 297–428). Mahwah, NJ: Lawrence Erlbaum Associates.

Eifuku, S., De Souza, W. C., Tamura, R., Nishijo, H., & Ono, T. (2004). Neuronal correlates of face identification in the monkey anterior temporal cortical areas. *Journal of Neurophysiology, 91,* 358–371.

Freiwald, W. A., & Tsao, D. Y. (2010). Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science, 330,* 845–851.

Gower, J. C. (1975). Generalized Procrustes analysis. *Psychometrika, 40,* 33–51.

Hancock, P. J., Burton, A. M., & Bruce, V. (1996). Face processing: Human perception and principal components analysis. *Memory and Cognition, 24,* 21–40.

Hasselmo, M. E., Rolls, E. T., Baylis, G. C., & Nalwa, V. (1989). Object-centered encoding by face-selective neurons in the cortex in the superior temporal sulcus of the monkey. *Experimental Brain Research, 75,* 417–429.

Jackson, D. A. (1995). PROTEST: A PROcrustean randomization TEST of community environment concordance. *Ecoscience, 2,* 297–303.

Johnston, R. A., Kanazawa, M., Kato, T., & Oda, M. (1997). Exploring the structure of multidimensional face–space: The effects of age and gender. *Visual Cognition, 4,* 39–57.

Blank & Yovel

Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 361,* 2109–2128.

Kruskal, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika, 29,* 1–27.

Kruskal, J. B., & Wish, M. (1978). *Multidimensional scaling.* Beverly Hills, CA: Sage.

Lando, M., & Edelman, S. (1995). Receptive-field spaces and class-based generalization from a single view in face recognition. *Network: Computation in Neural Systems, 6,* 551–576.

Lee, K., Byatt, G., & Rhodes, G. (2000). Caricature effects, distinctiveness, and identification: Testing the face–space framework. *Psychological Science, 11,* 379–385.

Leopold, D. A., Bondar, I. V., & Giese, M. A. (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature, 442,* 572–575.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience, 4,* 89–94.

Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1364–1382.

Lewis, M. B. (2004). Face–space R: Towards a unified account of face recognition. *Visual Cognition, 11,* 29–69.

Lewis, M. B., & Johnston, R. A. (1999). A unified account of the effects of caricaturing faces. *Visual Cognition, 6,* 1–42.

Logothetis, N. K., & Pauls, J. (1995). Psychophysical and physiological evidence for viewer-centered object representations in the primate. *Cerebral Cortex, 5,* 270–288.

McKone, E., Crookes, K., & Kanwisher, N. (2009). The cognitive and neural development of face recognition in humans. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences IV* (pp. 467–482). Cambridge, MA: MIT Press.

Moses, Y., Ullman, S., & Edelman, S. (1996). Generalization to novel images in upright and inverted faces. *Perception, 25,* 443–461.

Newell, F. N., Chiroro, P., & Valentine, T. (1999). Recognizing unfamiliar faces: The effects of distinctiveness and view. *Quarterly Journal of Experimental Psychology A, 52,* 509–534.

Nosofsky, R. M. (1992a). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–393). Hillsdale, NJ: Lawrence Erlbaum Associates.

Nosofsky, R. M. (1992b). Similarity scaling and cognitive process models. *Annual Review of Psychology, 43,* 25–53.

Perrett, D. I., Smith, P. A., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., et al. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society of London B: Biological Sciences, 223,* 293–317.

Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision Research, 46,* 2977–2987.

Robbins, R., & McKone, E. (2007). No face-like processing for objects-of-expertise in three behavioural tasks. *Cognition, 103,* 34–79.

Rolls, E. T. (2000). Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. *Neuron, 27,* 205–218.

Rolls, E. T., & Baylis, G. C. (1986). Size and contrast have only small effects on the responses to faces of neurons in the cortex of the superior temporal sulcus of the monkey. *Experimental Brain Research, 65,* 38–48.

Rolls, E. T., Cowey, A., & Bruce, V. (1992). Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 335,* 11–20; discussion 20-11.

Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika, 22,* 325–345.

Shepard, R. N. (1966). Metric structures in ordinal data. *Journal of Mathematical Psychology, 3,* 287–315.

Shepard, R. N. (1968). Cognitive psychology: A review of the book by U. Neisser. *American Journal of Psychology, 81,* 285–289.

Shepard, R. N. (1980). Multidimensional scaling, tree-fitting, and clustering. *Science, 210,* 390–398.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science, 237,* 1317–1323.

Shepard, R. N., & Chipman, S. (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology, 1,* 1–17.

Tanaka, J., Giles, M., Kremen, S., & Simon, V. (1998). Mapping attractor fields in face space: The atypicality bias in face recognition. *Cognition, 68,* 199–220.

Tovee, M. J., Rolls, E. T., & Azzopardi, P. (1994). Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. *Journal of Neurophysiology, 72,* 1049–1060.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology A, 43,* 161–204.

Valentine, T. (1995). *Cognitive and computational aspects of face recognition: Explorations in face–space.* London: Routledge.

Valentine, T. (2001). Face–space models of face recognition. In M. J. Wegner & J. T. Townsend (Eds.), *Computational, geometric, and process perspectives on facial cognition: Contexts and challenges* (pp. 83–113). Mahwah, NJ: Lawrence Erlbaum Associates.

Valentine, T., & Endo, M. (1992). Towards an exemplar model of face processing: The effects of race and distinctiveness. *Quarterly Journal of Experimental Psychology A, 44,* 671–703.

Zoccolan, D., Kouh, M., Poggio, T., & DiCarlo, J. J. (2007). Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. *Journal of Neuroscience, 27,* 12292–12307.