

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Vision Research

journal homepage: www.elsevier.com/locate/visres

Same critical features are used for identification of familiarized and unfamiliar faces

Naphtali Abudarham, Galit Yovel*

School of Psychological Sciences, Tel Aviv University, Israel

ARTICLE INFO

Number of Reviews = 1

Keywords:

Face processing
Face identification
Familiar and unfamiliar faces
Facial features face space

ABSTRACT

Many studies have shown better recognition for faces we have greater experience with, relative to unfamiliar faces. However, it is still not clear if and how the representation of faces changes during the process of familiarization. In a previous study, we discovered a subset of facial features, for which we have high perceptual sensitivity (PS), that were critical for determining the identity of unfamiliar faces. This was done by assigning values to 20 different facial features based on perceptual rating, converting faces into feature-vectors, and measuring the correlations between face similarity ratings and distances between feature-vectors. In the current study, we examined the contribution of high and low-PS features to face identity after familiarization. To familiarize participants with unfamiliar faces, we used an individuation training protocol that was found to be effective in previous studies, in which different names are assigned to different faces and participants are asked to learn the face-name association. Our findings show that even after repeated exposure to the same image of each identity, which allows close examination of all facial features, only high-PS features contributed to face identity, while low-PS features did not. This subset of high-PS features includes both internal and external features and part and configuration features. We therefore conclude that identification of familiarized and unfamiliar faces may rely on the same subset of critical features. These findings further support a new categorization of facial features according to their perceptual sensitivity.

1. Introduction

One of the most fundamental question in the study of face identification, is which facial features are used to define the identity of a face. Previous studies have considered different types of features including configural vs. part-based features (for reviews see (Maurer, Grand, & Mondloch, 2002; McKone & Yovel, 2009), or external vs. internal facial features (Ellis, Shepherd, & Davies, 1979; O'Donnell & Bruce, 2001; Young, Hay, McWeeny, Flude, & Ellis, 1985) that may play different roles in face recognition. In a previous study, we proposed a new categorization of facial features, based on the perceptual-sensitivity (PS) to detect differences in these features across different faces. We found that facial features for which we have high perceptual sensitivity (high-PS) were critical for identification of unfamiliar faces more than facial features for which we have low perceptual sensitivity (low-PS). To discover these critical features, we used a novel “reverse-engineering” approach. The main premise of our approach was that critical features for face identity are those features that when changed, change the identity of a face. The complementary assumption is that feature changes that do not change the identity of the face, are not critical for

face identity, and are therefore considered “allowable” changes in appearance (Abudarham & Yovel, 2016).

To reveal which features are critical for face identity we used the following procedure: First, we constructed a face-space, by perceptually assigning values to facial features, thereby representing faces as feature-vectors (Fig. 1A). This allowed us to calculate face-space distances between faces, i.e. mathematical distances between feature-vectors (Fig. 1B). We then computed the inter-rater agreement for each feature. This analysis revealed that features vary considerably in their level of inter-rater agreement, suggesting that participants have different perceptual sensitivity to detect feature differences across faces (Supp. Fig. 1). For example, we found higher perceptual sensitivity to detect eyebrow-thickness or lip-thickness differences across faces than for eye-distance or jaw-width. To test the hypothesis that features for which we have high perceptual sensitivity (PS) are critical for face identity, we changed facial features systematically, by either changing high-PS features or low-PS features, and measured identification. Our results showed that changing high-PS features changed the identity of a face, whereas changing low-PS features did not change the identity of the face. Accordingly, we concluded that high-PS features are critical for

* Corresponding author.

E-mail address: gality@post.tau.ac.il (G. Yovel).

<https://doi.org/10.1016/j.visres.2018.01.002>

Received 1 October 2017; Received in revised form 2 January 2018; Accepted 3 January 2018
0042-6989/ © 2018 Elsevier Ltd. All rights reserved.

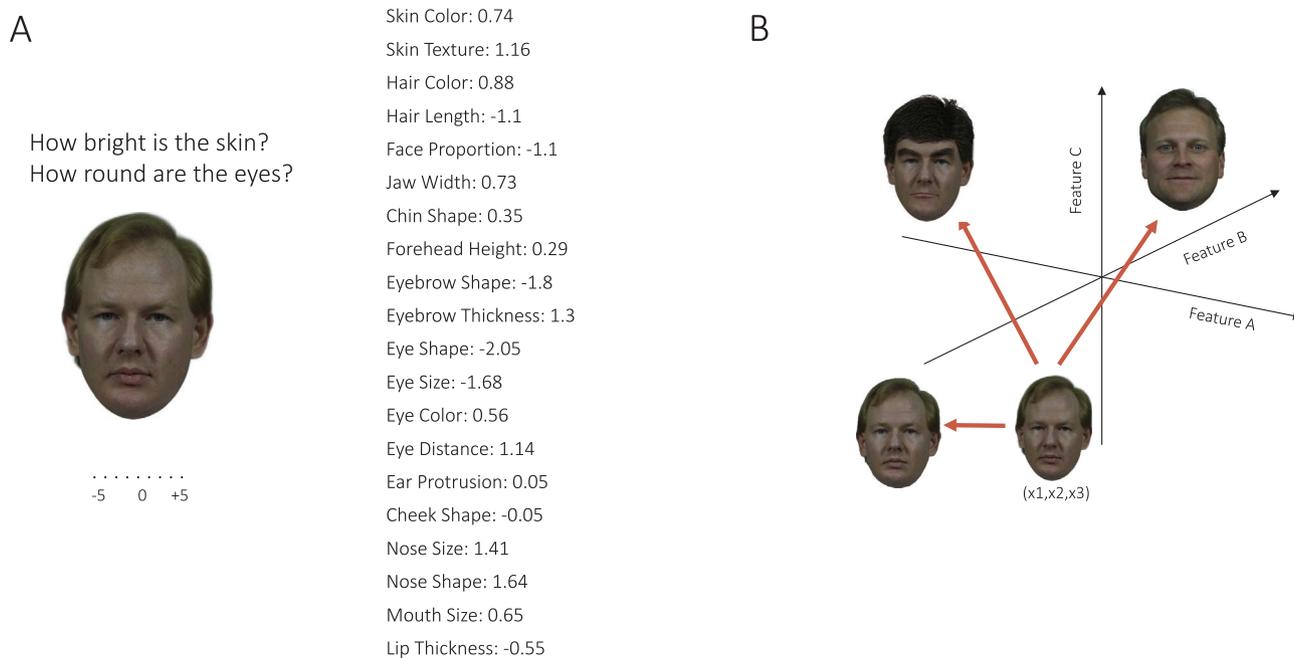


Fig. 1. A. Faces were converted to a feature representation by asking participants to perceptually rate each of 20 facial features. For example, how long is the hair, how wide is the face. B. Face space distances between faces were computed based on the mathematical distance between feature vectors.

face identity. Finally, we found that high-PS features tend to vary less under pose changes, compared with low-PS features. This suggests that high-PS features are more important for identification because they allow identification of the same person across different appearances.

The critical features that we revealed were based on matching of unfamiliar faces. However, it is possible that repeated exposure to the same face may allow participants to acquire information about face identity from additional facial features. Thus, our findings that high-PS features contribute more than low-PS features to face identity may be limited to faces that we have no prior exposure to, whereas all or additional features may become equally important once we get familiarized with a face. Alternatively, if following repeated exposure to the same face, participants still use the same subset of high-PS features, that would indicate that high-PS features dominates face identity even when low-PS features are available for repeated and close inspection.

To test this hypothesis we familiarized participants with unfamiliar faces, using an individuation training protocol (Tanaka & Pierce, 2009; Yovel et al., 2012), and then applied the same approach that we used to discover the critical features for unfamiliar faces, after familiarization. This familiarization protocol included repeated exposure to a single image of each familiarized face, allowing participants to study and memorize faces without variations. Although this protocol does not model how we naturally become familiar with faces in real life, a process which typically includes seeing the same face in varying appearances, it does enable us to test whether familiarization changes the weights of high-PS vs. low-PS features in a controlled manner. In addition, familiarization with multiple and variable images may encourage participants to encode those features that stay invariant across appearances (high-PS features), whereas using a single image is more likely to enable using also low-PS features.

To this end, we first familiarized participants with 30 different identities. We then changed facial features (Fig. 2), measured similarity between original and changed faces, and the feature vector distance (i.e., face space distance) between them. We then measured the correlation between perceptual similarity scores between faces, and the face-space distances, as a function of the type and number of features used to compute the face space distances. This allowed us to reveal which features are used to determine the identity of familiarized faces.

2. Method

2.1. Participants

Twenty-two Psychology students participated in the study in exchange for class credit. The study was approved by the ethics committee of Tel Aviv University.

The following methods of creating the stimuli, tagging faces, changing facial features and the face matching task are identical to those described in (Abudarham & Yovel, 2016), and are reproduced here for convenience.

2.2. Stimuli

A total of 100 Caucasian male faces, with no glasses or facial hair were taken from the Color FERET database. All faces had neutral expression. The pictures were frontal, and taken with adequate lighting. For each face, we took two different pictures that were physically different, i.e. taken at different times under similar lighting/pose/camera conditions (see Fig. 3 for an example of a Same pair). One of these pictures served as the “reference” picture used for comparison, and the other served as a “base” picture, and was later changed.

2.3. Face tagging: Converting faces into feature vectors, and measuring face-space distances

To assess face-space distances between faces we converted each of the 100 faces in our database into a vector representation, by assigning values to 20 features for each face (see Fig. 1A for an example of a feature vector of a face). These values were assigned by asking participants to rate each feature between -5 and $+5$ (for example: how bright is the skin? how large are the eyes?). (Fig. 1, see also (Abudarham & Yovel, 2016) This allowed us to measure face-space distances between faces, by calculating the L1-norm between feature-vectors, i.e. taking the sum of the absolute differences between the feature-vectors before and after change.

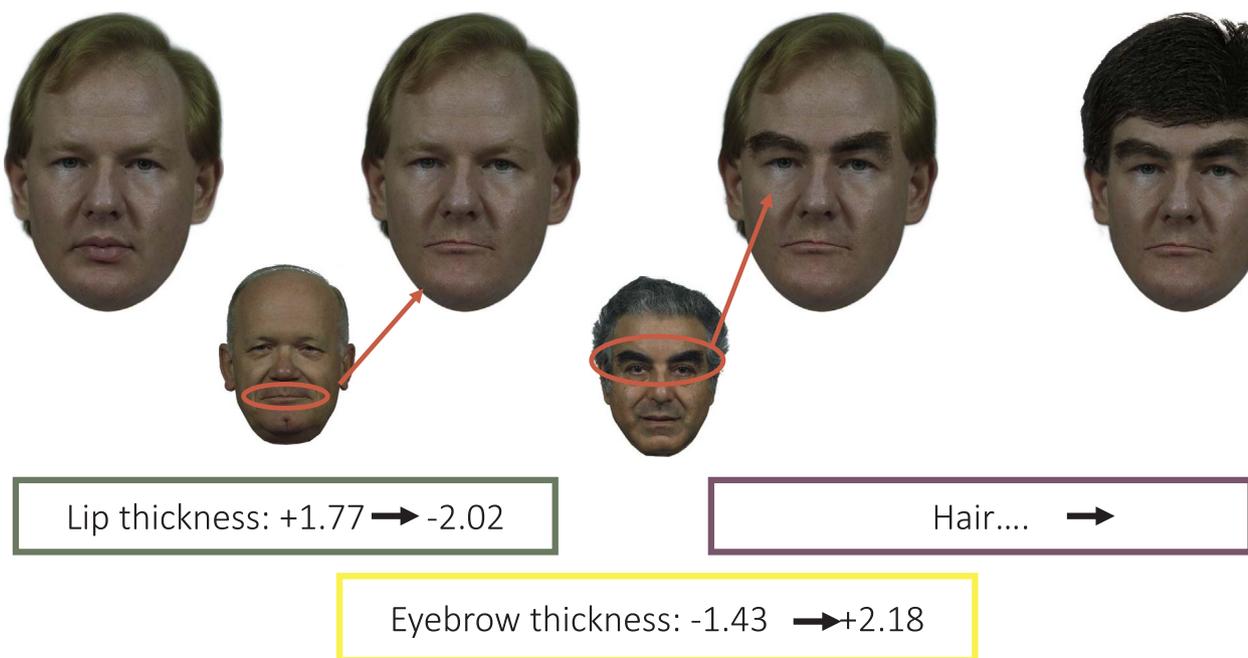


Fig. 2. Feature changes was done by changing the most distinctive features (based on their perceptual scores obtained in the tagging stage) in each of 30 faces. This resulted in 30 pairs of original and changed faces.

2.4. Changing facial features

We randomly selected 30 out of the 100 faces to test the effect of different facial-feature changes. To perform the feature changes, we developed a method using Adobe Photoshop ©, which included copying

features from donor faces, based on feature values obtained during the tagging procedure (see Fig. 2 for the face changing procedure). To decide which features to change we sorted the facial features of each face by the absolute magnitude of their values, and started replacing features from the largest values, and downwards to smaller values. Each

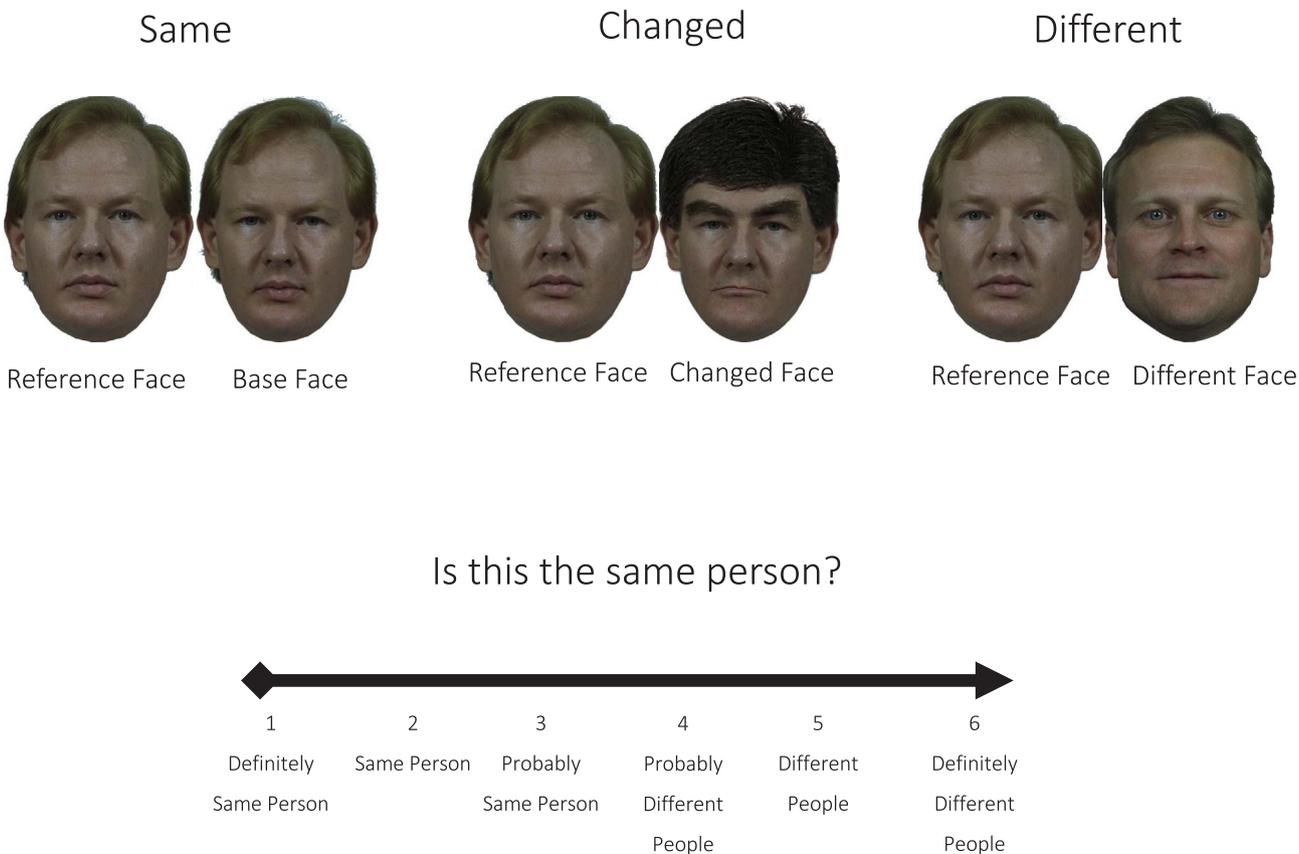


Fig. 3. Participants were asked to rate how similar are three types of face pairs: Same pairs: two different images of the same person, a base image and a reference image. Changed pairs: a modified base image and the reference image. Different pairs: a different individual and the reference face.

feature was replaced with a feature with as far away value as possible (from the existing database of tagged faces), thus making the largest possible change in the feature-vector. Our goal was to create a variety of changes, with a variety of face-space distances that can be correlated with perceptual similarity scores obtained from another group of participants (see below). We changed each of the 30 faces by replacing a different number of features. The stopping criterion was the number of features at which a face started to look un-natural, or “Photoshopped”, according to judgments made by a group of participants not participating in the main experiment. This created a set of 30 pairs of original and changed faces that varied in the number of features that were changed (ranging between 2 and 11) and in the resulting face space distances among them (ranging between 15 and 37). We could therefore measure if face-space distances co-varied with perceptual differences between the original and changed face across the 30 face pairs.

2.5. Face matching task

To measure perceptual differences between the original and changed face, based on these 30 changed faces we created 90 face-pairs, in three conditions: Same, Changed and Different face-pairs (Fig. 3). In the Same condition, we compared the two different pictures of the same face (the reference picture and base picture), of the 30 faces that we manipulated. In the Changed condition, we used the reference pictures used in the Same condition, and the changed picture, which was created from the base picture. Thus, Changed face pairs did not differ only in the features that were changed in our manipulation, but also in the low-level pixel information. This was done to ensure that participants would not rely on low-level information when they made similarity judgments, and will perform face-matching rather than image-matching. Finally, in the Different condition, we used two pictures of different identities, including either original or changed pictures.

2.6. Procedure

The experiment included 2 sessions: a familiarization session, and a face-matching test.

2.7. Familiarization procedure

We used individuation training in which participants learned to associate 30 faces with names to familiarize participants with the unfamiliar faces. The training included 3 sessions, each with 10 different identities. Each session included 5 phases: Phase 1: the 10 faces were presented sequentially, in a random order, for 2 s each, and a name was displayed under each face. The names were selected such that each face was assigned a name that started with a different letter of the alphabet. This was repeated 5 times. Phase 2: The 10 faces with the names below them were again displayed sequentially (in a random order), but this time each face appeared on the screen until the participant pressed the first letter of the face name. Phase 3: The 10 faces were displayed sequentially, this time without their names, and participants were asked to press the first letter of the name of each face. If the participant pressed the correct letter, the next face was displayed. If not, the correct name was displayed for 2 s, and then the next face was displayed. This phase was repeated 5 times. Phase 4: This phase was the same as Phase 3, only this time no feedback was given if the participant pressed the wrong letter. This phase repeated 6 times. Phase 5: The last phase was a face name matching test. A name was displayed for 1 s, followed by a 1 s interval, and then a face was displayed for 1 s. The participant was instructed to press 1 if they thought that the name matched the face, based on previous stages, or 0 if it was the wrong name for the face. This phase repeated 6 times.

2.8. Face-matching experiment

Following the familiarization session, each participant performed a face matching task. We presented participants with 30 face-pairs of each the three conditions (Same, Changed and Different), and asked them to mark, on a scale of 1 to 6, whether the two pictures belong to the same person or not (1 being “definitely the same person” and 6 being “definitely different people”) (Fig. 3). The two pictures in each pair were presented simultaneously till response. The order of the face-pairs and the right-left position in each pair were randomized across participants.

Same and Different conditions served as a baseline, for calculating the similarity between changed pairs. This baseline was used in three ways: first, it provided the participants with reference conditions to tune themselves and to understand how same and different pairs look. Second, it enabled us to measure the baseline of identification in natural faces, in both Same and Different conditions, on this particular set of pictures. Previous studies (Bruce et al., 1999; Burton, White, & McNeill, 2010) have shown that human performance is far from perfect on this task, therefore we needed to tune our scale. Third, we used the baseline score for each Same pair for calculating the score for the Changed pair. Same pairs consisted of two pictures, and one of these was then modified to create the Changed pair. Therefore, we needed to subtract the score for the Same pair from the score of the Changed pair, to compensate for the basic differences between the two pictures. Accordingly, the formula for calculating the perceptual-distance between a Changed pair was: (mean score for Changed pair - mean score for the Same pair of that face)/(mean score for all Different pairs - mean score for the Same pair of that face). The resulting perceptual-distance scores ranged between ~ 0 , and ~ 1 , where 0 indicates that the Changed pair received the same score as the Same pair of that face, and 1 indicated the Changed pair got a score that is close to the mean score for Different pairs. Thus, perceptual-distances closer to 1 indicate that these faces were perceived as different people meaning, a change in identity, and scores that are significantly lower than 1 are feature changes that did not change the identity of a face.

2.9. Measuring the face-space distances between original and changed faces

To measure the face-space distances, i.e. the distances between the feature-vectors of faces before and after change. The changed faces were tagged, using the same tagging procedure explained above. 44 participants tagged the 30 changed faces (each participant tagging 7 out of the 20 features, to avoid fatigue, resulting in an average of 15 raters per feature), to obtain feature-vectors for the changed faces. These participants did not participate in the face-matching procedure. Face-space distances were calculated as the L1-norm between the feature-vectors before and after change.

3. Results

The familiarization procedure was overall successful, and participants were able to correctly match the faces with their assigned names ($\sim 97\%$ correct responses).

Fig. 4 shows the correlation between face-space distances and the perceptual-distances as measured using all 20 features. The correlation is high and significant ($r = 0.69$, $p < .01$) and very similar to what was found for unfamiliar faces ($r = 0.72$) (Abudarham & Yovel, 2016). Fig. 5 shows the correlation between the perceptual similarity scores obtained for the same faces when they were unfamiliar (data taken from (Abudarham & Yovel, 2016)), and after familiarization. Perceptual similarity scores were very similar and highly correlated, suggesting that the features that were changed had the same effect on identification for both novel and familiarized faces. Note that the unfamiliar and familiarized faces were rated by different groups of participants and therefore the high correlation between them does not reflect effect of

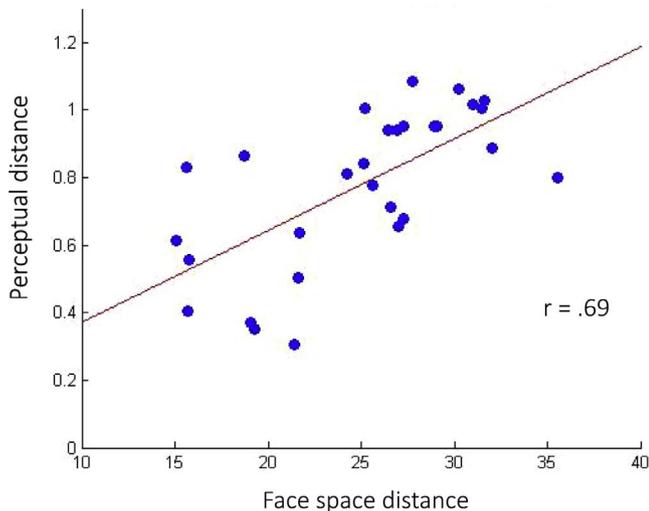


Fig. 4. The correlation between face space distance (distance between feature vectors) based on all 20 facial features, and perceptual similarity scores for the same pairs of faces. The high correlations indicate that the selected features are used to determine the similarity between two familiarized faces.

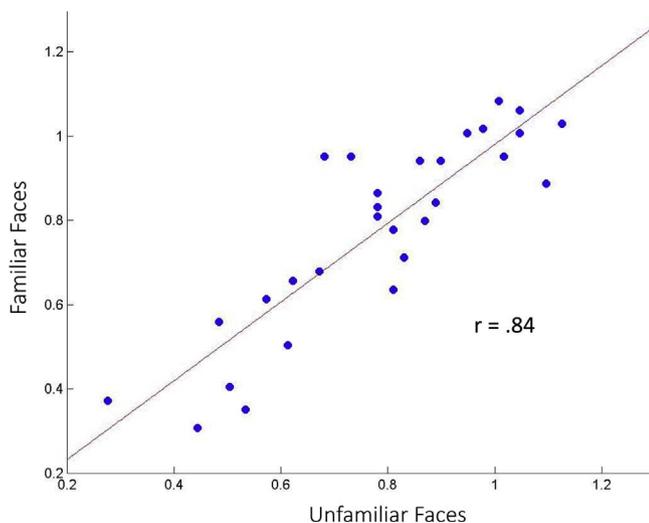


Fig. 5. High correlations between perceptual similarity scores for the same face pairs with and without familiarization. Similarity ratings for unfamiliar and familiarized faces were obtained from different groups of participants.

previous experience with the matching task.

In our previous study, we measured inter-rater agreement for feature ratings and interpreted these scores as a measure of perceptual sensitivity to detect differences in features across faces (see [supp Fig. 1](#) for perceptual-sensitivity scores). Our findings showed that these perceptual sensitivity (PS) scores predicted feature importance for identification. To examine the importance of the different features for familiarized faces, we computed the face space distance with an increasing number of features, either starting from features that have high-PS scores or features that have low-PS scores. [Fig. 6](#) shows the values of the Spearman correlation between face-space distances and perceptual distances, as a function of the number of features used to calculate the face-space distance. The blue (solid) line shows the values of the correlations as we added features, starting from the feature with the highest PS score (i.e., lip thickness), and adding features in the order of their PS scores, up to the 20th feature with the lowest PS score (i.e. mouth size). The green (dashed) line shows the opposite – starting from the lowest PS score and going up to the highest PS score.

When starting with high PS features, the correlation increases with

the number of features until it plateaus at 6–7 features, where it reaches the numerical correlation value that was found when all 20 features are used to compute face space distance. Correlations between perceptual similarity scores and face space distance that was based on 1 or 2 features were significantly different from the correlation obtained using all 20 features. At 3 features ($r = 0.37$) and higher, the correlations were not significantly different from 20 features. In contrast, when starting with low-PS features the correlations start increasing gradually only when 8 features are used to compute the face space distance, and are significantly lower than the correlations with 20 features up to 14 features ($r = 0.36$). At 15 features ($r = 0.41$) and higher the correlations are not significantly different than at 20 features.

For unfamiliar faces, we find a similar pattern such that correlations between perceptual similarity scores and face space distance that were based on 1 or 2 features were significantly different from the correlation when the face space distance was computed using all 20 features ($r = 0.72$). When using 3 features to compute face space distance the correlation ($r = 0.43$) is not significantly different than the correlation at 20 features. When starting to add low PS features to compute the face space distance, the correlations are significantly different from 20 features up to 11 features, whereas the correlations based on 12 features ($r = 0.45$) is not significantly different from 20 features.

In addition, we compared the difference between high and low-PS features as the number of features increases and found that for familiarized faces the correlations based on 4 through 15 features, when starting with high-PS features, were significantly higher than when using the same number of features but starting with low-PS features (See [Fig. 6](#), left pane, the differences between the green (dashed) and blue (solid) lines). The same analysis for unfamiliar faces ([Fig. 6](#), right pane) reveals that the differences between correlations starting with high or low PS features are significantly higher when 4 through 11 features are used to compute the face space distance.

Finally, we ran a mixed-effect model with Feature order and Familiarization as fixed effects and correlation values as a random effect. The analysis revealed a significant interaction ($t(57) = 2.09$, $p < .05$) indicating a larger difference between correlations when starting with high than low-PS features for familiarized than unfamiliar faces (see [Fig. 6](#)). Thus, low-PS features contribute even less to face identity after than before familiarization.

4. Discussion

The goal of the current study was to assess whether features that were found to be critical for matching unfamiliar faces may be also critical for matching familiarized faces. Our results show that, similar to unfamiliar faces, perceptual similarity rating is highly correlated with face-space distances based on 20 features. In addition, a subset of 3–4 features, for which we have high perceptual sensitivity, accounts for this high correlation ([Fig. 6](#)). These results indicate that despite repeated exposure to all facial features during familiarization, high-PS features are still more important for identification than low-PS features.

Previous studies have suggested that internal features are more critical for identification of familiar faces but not for unfamiliar faces due to increased attention to internal facial features during the interaction with familiar faces ([Ellis et al., 1979](#); [Young et al., 1985](#)). This division into internal and external features is coarser than the detailed analysis of features we present in the current study. More importantly, our selection of features was based on an empirical measure of perceptual sensitivity rather than an arbitrary division of the face to external and internal features. The 20 features that we used included both external features such as hair, face-proportion, jaw-width, as well as internal features, such as eye-size, eye-distance, lip-thickness or nose-size. Out of the top 5 high-PS features, 4 are internal features (lip-thickness, eye-color, eye-shape and eyebrow-thickness), and the fifth is the hair, which was found to be important for familiar faces in previous studies ([Sinha & Poggio, 1996](#)). However, eye-distance and mouth-size

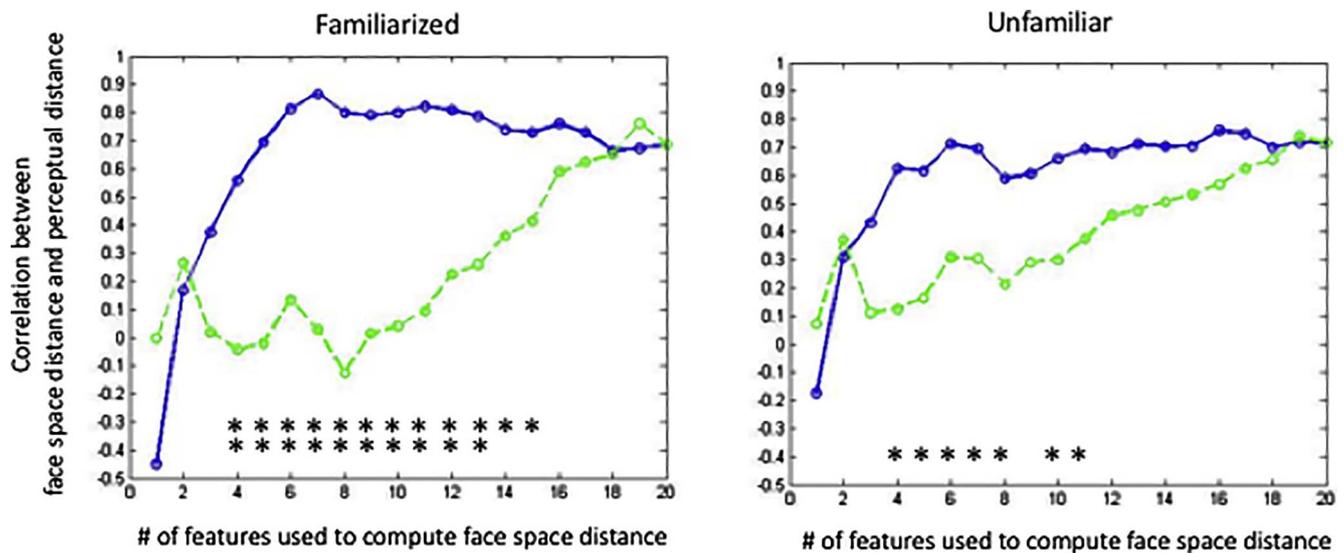


Fig. 6. Left: The correlations between face space distance and perceptual similarity for familiarized faces, as a function of the number of features used to compute face space distance. Blue (solid) line: when starting from high PS features the correlations reaches the peak value after 6–7 features and then plateaued. Green (dashed) line: when starting from low PS features, the correlations reaches the peak value only when high PS features are also used. Right: the same plots for unfamiliar faces (taken from (Abudarham & Yovel, 2016)). The asterisks indicate where the differences between the correlation values when starting with high PS or low PS for the same number of features are significant (*: $p < .05$, **: $p < .01$).

are also internal features but they got the lowest PS scores, therefore our fine distinction between features allows us to differentiate between internal features that are critical for identification and those that are not. These findings suggest that a gross categorization to internal and external features is not sensitive for revealing which features are used for face recognition.

Another frequently used categorization of facial features is the part vs. configuration (i.e., spacing) distinction. This dichotomy has been commonly used in the face processing literature and previous studies have reported that both types of facial information were important for familiar and unfamiliar faces (Collishaw & Hole, 2000; Schwaninger, Lobmaier, & Collishaw, 2002). O'Donnell and Bruce, however, found that eye-distance, which is considered a configural feature, was more important for identity of familiarized faces (O'Donnell & Bruce, 2001). In contrast, our results suggest that configural features are less important for face recognition. For example, distances and general shapes like eye-distance and face-proportion were found to be less important for identification, consistent with studies that showed intact face recognition of compressed familiar faces (Hole, George, Eaves, & Rasek, 2002; Sandford & Burton, 2014). In contrast, eye-color and eyebrow-thickness, which correspond to face parts, were found to be important for face identity. The reason for the inconsistency between our results with previous data, is that most studies that compared configural and part-based facial information, tested participants ability to detect differences between the same image of a face that differ only in a configural or part-based feature (Le Grand, Mondloch, Maurer, & Brent, 2001; McKone & Yovel, 2009; Yovel & Kanwisher, 2004). Our method for assessing feature importance is based on measuring features or comparing features across different faces, a method that is more relevant for real-life face identification.

The distinction between part-based and configural information originated from studies that compared between identification of upright and inverted faces (for review see McKone & Yovel, 2009), under the assumption that features that are used for identification of upright faces but not inverted faces are critical for face recognition. Critical features found in our study included the hair and eye color, can be easily detected in inverted faces (Yovel & Duchaine, 2006). A similar assumption was that features that are used by prosopagnosic individuals for face recognition, such as the hair, are not critical for face recognition. However, the fact that a feature such as the hair or eye color can be used also when face recognition is impaired (when faces are inverted or

in prosopagnosia), does not mean that it is not used for intact face recognition. Hair, which is often excluded in studies of face recognition, has been shown both in our study and in previous studies to be critical for face recognition (Abudarham & Yovel, 2016, Sinha & Poggio, 1996, 2002). Our method therefore offers a direct way to determine the importance of different facial features for face recognition, which relies on changing features and measuring the effect of different changes on face similarity. We found that the importance of different features for perceptual similarity is correlated with perceptual sensitivity to detect differences in these features.

As discussed in our previous study, high-PS features may be critical for face identity because they tend to vary less under different appearances, such as lighting or pose. Previous studies have suggested that in the process of familiarization we learn what information varies or stays constant for each face, based on the specific exposure or experience with that face (Burton, Kramer, Ritchie, & Jenkins, 2016; Jenkins, White, Van Montfort, & Mike Burton, 2011). Accordingly, it has been suggested that exposure to high variability of images is critical for obtaining the view-invariant representation that we have for familiar faces. In our study, we used an individuation training, in which participants were exposed to only one image of each individual person associated with a name. Thus, our findings may not generalize to familiarization that is based on different images of the same individual (Burton et al., 2016), as well as to personally familiar or famous faces for which we have rich perceptual experience.

However, the type of familiarization used in the current study, of training on a single image, makes our findings even stronger. While training on a single image, participants may become more acquainted with that single appearance of a face, and perhaps learn to pay more attention to more features in the face including low PS features, giving them equal weight following familiarization. Our findings suggest the opposite – participants still attribute more weight to high-PS features than to low-PS features. Furthermore, our previous findings suggest that the critical features people use to identify faces are those that vary less under different appearances. Thus, exposure to variable images of the same face would in fact encourage participants to focus on the invariant features, which are the high-PS features, whereas exposing participants repeatedly to a single image allows them to study all features of that face. Our findings show that even though participants saw the same face image repeatedly for each identity, they attribute less importance to the low-PS features, just as they do for faces they see only

once (unfamiliar faces).

In summary, we showed that the same critical features are used to determine the identity of familiarized and unfamiliar faces. These critical features tend to vary less under different appearances and may be therefore useful for recognition of faces across different appearances. Future studies are needed to determine whether faces for which we have more variable experience, either through training with highly variable faces, or real-life familiar faces, may rely on additional or different sets of features.

Acknowledgments

Portions of the research in this manuscript use the FERET database of facial images collected under the FERET program, sponsored by the DOD Counterdrug Technology Development Program Office. The said images were processed by the author for this specific experiment. You may not use any of the images in this experiment without written permission from NIST and from the author.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.visres.2018.01.002>.

References

- Abudarham, N., & Yovel, G. (2016). Reverse engineering the face space: Discovering the critical features for face identification. *Journal of Vision*, *16*(3), 40. <http://dx.doi.org/10.1167/16.3.40>.
- Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, A. M., & Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied*, *5*(4), 339.
- Burton, A. M., Kramer, R. S. S., Ritchie, K. L., & Jenkins, R. (2016). Identity from variation: Representations of faces derived from multiple instances. *Cognitive Science*, *40*(1), 202–223. <http://dx.doi.org/10.1111/cogs.12231>.
- Burton, A. M., White, D., & McNeill, A. (2010). The Glasgow face matching test. *Behavior Research Methods*, *42*(1), 286–291.
- Collishaw, S. M., & Hole, G. J. (2000). Featural and configurational processes in the recognition of faces of different familiarity. *Perception*, *29*(8), 893–909. <http://dx.doi.org/10.1068/p2949>.
- Ellis, H. D., Shepherd, J. W., & Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features: Some implications for theories of face recognition. *Perception*, *8*(4), 431–439.
- Hole, G. J., George, P. A., Eaves, K., & Rasek, A. (2002). Effects of geometric distortions on face-recognition performance. *PERCEPTION-LONDON*, *31*(10), 1221–1240.
- Jenkins, R., White, D., Van Montfort, X., & Mike Burton, A. (2011). Variability in photos of the same face. *Cognition*, *121*(3), 313–323. <http://dx.doi.org/10.1016/j.cognition.2011.08.001>.
- Le Grand, R., Mondloch, C. J., Maurer, D., & Brent, H. P. (2001). Neuroperception: Early visual experience and face processing. *Nature*, *410*(6831), 890.
- Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, *6*(6), 255–260.
- McKone, E., & Yovel, G. (2009). Why does picture-plane inversion sometimes dissociate perception of features and spacing in faces, and sometimes not? Toward a new theory of holistic processing. *Psychonomic Bulletin & Review*, *16*(5), 778–797. <http://dx.doi.org/10.3758/PBR.16.5.778>.
- O'Donnell, C., & Bruce, V. (2001). Familiarisation with faces selectively enhances sensitivity to changes made to the eyes. *Perception*, *30*(6), 755–764. <http://dx.doi.org/10.1068/p3027>.
- Sandford, A., & Burton, A. M. (2014). Tolerance for distorted faces: Challenges to a configural processing account of familiar face recognition. *Cognition*, *132*(3), <http://dx.doi.org/10.1016/j.cognition.2014.04.005>.
- Schwaninger, A., Lobmaier, J. S., & Collishaw, S. M. (2002). Role of featural and configural information in familiar and unfamiliar face recognition. *Biologically motivated computer vision* (pp. 643–650). Springer.
- Sinha, P., & Poggio, T. (1996). I think I know that face. *Nature*, *384*(6608), 460.
- Sinha, P., & Poggio, T. (2002). 'United' we stand. *Perception*, *31*(1), 133.
- Tanaka, J. W., & Pierce, L. J. (2009). The neural plasticity of other-race face recognition. *Cognitive, Affective & Behavioral Neuroscience*, *9*(1), 122–131. <http://dx.doi.org/10.3758/CABN.9.1.122>.
- Young, A. W., Hay, D. C., McWeeny, K. H., Flude, B. M., & Ellis, A. W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception*, *14*(6), 737–746.
- Yovel, G., & Duchaine, B. (2006). Specialized face perception mechanisms extract both part and spacing information: Evidence from developmental prosopagnosia. *Journal of Cognitive Neuroscience*, *18*(4), 580–593. <http://dx.doi.org/10.1162/jocn.2006.18.4.580>.
- Yovel, G., Halsband, K., Pelleg, M., Farkash, N., Gal, B., & Goshen-Gottstein, Y. (2012). Can massive but passive exposure to faces contribute to face recognition abilities? *Journal of Experimental Psychology: Human Perception and Performance*, *38*(2), 285.
- Yovel, G., & Kanwisher, N. (2004). Face perception: Domain specific, not process specific. *Neuron*, *44*(5), 889–898.