Cognition 182 (2019) 73–83

Contents lists available at ScienceDirect

Cognition

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

Original Articles

Critical features for face recognition

Naphtali Abudarham[a], Lior Shkiller[a], Galit Yovel[a,b,*]

[a] School of Psychological Sciences, Tel Aviv University, Tel Aviv, Israel
[b] Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel

ARTICLE INFO

Keywords:
Face recognition
Familiar faces
Face space
Deep neural network
Feature space

ABSTRACT

Face recognition is a computationally challenging task that humans perform effortlessly. Nonetheless, this remarkable ability is better for familiar faces than unfamiliar faces. To account for humans' superior ability to recognize familiar faces, current theories suggest that different features are used for the representation of familiar and unfamiliar faces. In the current study, we applied a reverse engineering approach to reveal which facial features are critical for familiar face recognition. In contrast to current views, we discovered that the same subset of features that are used for matching unfamiliar faces, are also used for matching as well as recognition of familiar faces. We further show that these features are also used by a deep neural network face recognition algorithm. We therefore propose a new framework that assumes similar perceptual representation for all faces and integrates cognition and perception to account for humans' superior recognition of familiar faces.

1. Introduction

Face recognition is a computationally challenging task that requires fine discrimination between similarly looking images of different identities, as well as generalization across different images of the same individual. Although humans are considered experts in face recognition, studies have shown that our face recognition abilities are superior to faces we are familiar with, whereas our ability to match unfamiliar faces is error-prone (Young & Burton, 2017b, 2017a). These findings led to the suggestion that familiar face recognition depends on a different set of facial features, based on the extensive experience that we have with them than those used for unfamiliar faces. For example, it has been suggested that familiar face recognition is primarily based on internal facial features, whereas unfamiliar face matching is primarily based on external facial features (Ellis, Shepherd, & Davies, 1979; Kramer, Towler, Reynolds, & Burton, 2017; O'Donnell & Bruce, 2001; Young, Hay, McWeeny, Flude, & Ellis, 1985). According to another view, the representation of familiar faces is based on the average of their different appearances, which excludes superficial image-based information that may dominate the representation of unfamiliar faces (Jenkins & Burton, 2011). This view further posits that throughout our experience with variable images of familiar faces, we learn the idiosyncratic features that remain invariant across their different appearances and are unique for each identity. This view therefore suggests that a different set of features is used to recognize different familiar identities (Burton, Kramer, Ritchie, & Jenkins, 2016).

In a recent study, we used a novel *reverse engineering* approach to reveal which facial features are critical for face identity. We found a subset of features for which humans have high perceptual sensitivity to detect differences between different identities (high-PS features) (Abudarham & Yovel, 2016) (see Fig. S2). We then showed that systematically changing high-PS features changes the identity of faces, whereas changing features for which humans have low perceptual sensitivity (low-PS features) did not change the identity of faces (see Fig. 1). Importantly, these high-PS features remain invariant across different head views (Abudarham & Yovel, 2016), making them useful not only for discrimination between identities but also for generalizing across different appearances of the same identity.

Nevertheless, this subset of features was shown to be critical for unfamiliar faces and may not generalize to familiar faces, with which we have much greater experience. Thus, the goal of the current study was to use the same reverse engineering approach to reveal which features are critical for familiar face recognition. This allowed us to test the common view that different facial features are used for the identification of familiar and unfamiliar faces.

To that end, in Experiment 1 we first examined the role of high-PS vs. low-PS features in a familiar face matching task, using the same matching task that was used for unfamiliar faces in our previous study (Abudarham & Yovel, 2016). An important difference between familiar and unfamiliar faces is that familiar faces are represented in memory. Features that are used for matching two faces presented simultaneously, may not be used for matching a familiar face to its representation in

* Corresponding author at: School of Psychological Sciences & Sagol School of Neuroscience, Tel Aviv University, Tel Aviv 69978, Israel.
E-mail address: gality@post.tau.ac.il (G. Yovel).

https://doi.org/10.1016/j.cognition.2018.09.002
Received 2 April 2018; Received in revised form 3 September 2018; Accepted 3 September 2018

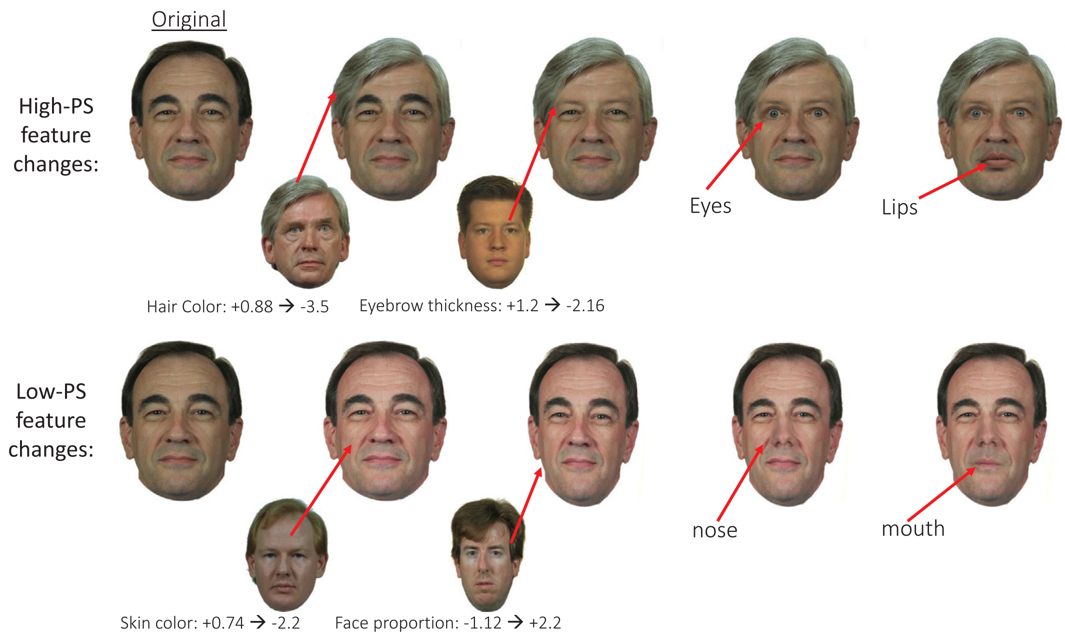

Fig. 1. To reveal which features are critical for the identity of the face we replaced 5 features for which we have high perceptual sensitivity (high-PS features) (top) or 5 features for which we have low perceptual sensitivity (low-PS features) (bottom) (Abudarham & Yovel, 2016). Features were taken from faces of different identities (see methods for more information about the feature changing procedure). Similarity matching between the original face and the changed face on the far right showed that changing high-PS features changed the identity of the face whereas changes in low-PS features did not change the identity of the face (see Fig. 3 for similarity matching results).

memory. Therefore, in Experiments 2 and 3 we examined whether these features are also used for face recognition. Finally, the features that we found correspond to semantic descriptions of facial features (e.g., eyes, mouth), and may therefore overlook visual information that cannot be described by these labels. We therefore examined whether these features are also used by a face recognition algorithm, that is not bound to these semantic meanings. Recently, Deep Neural Network (DNN) algorithms have reached human level performance on unconstrained ("wild") facial images, in which faces appear in various poses, expressions and illuminations (Schroff & Philbin, 2015). These advances are the result of the capability of DNNs to extract the invariant information through supervised learning with many different images of the same identity (O'Toole, Castillo, Parde, Hill, & Chellappa, 2018). We therefore hypothesized that a DNN may be tuned to the same invariant, high-PS features that humans use for face recognition (Experiment 4).

2. Experiment 1 – Critical features for matching familiar faces

To determine whether changing high-PS features, but not low-PS features, changes the identity of a familiar face, we used a matching task similar to the one we used in a previous study with unfamiliar faces (Abudarham & Yovel, 2016). Familiar faces were modified by either changing five high-PS features or five low-PS features (Fig. 2, Fig. S3). We presented participants with pairs of celebrity faces, before and after feature changes, and asked them to rate whether the two pictures belong to the same person or to different people (Fig. 3A, top). Pairs of same identity and different identity faces were also presented to obtain baseline performance to which matching abilities for low-PS and high-PS pairs can be compared.

2.1. Method

2.1.1. Participants

All participants were Amazon-Mechanical-Turk workers, participating in the experiment for payment (approximately 1$ per 15 min of work). A total of 38 participants (American residence, 18 females, 28 Caucasians, 6 East-Asians, 2 African-American, 1 Hispanic/Latin and 1 Middle-Eastern, ages 23–66 (mean 39.4, standard deviation (SD) 13.6) performed the experiment.

2.1.2. Stimuli

Ten American celebrities – all adult Caucasian males – were selected for the experiment. For each identity, we downloaded from the internet two frontal neutral expression images, with no glasses, hat or facial hair, and with adequate lighting and quality. All pictures were cropped from the background, and cut below the chin, leaving just the face, including the hair and ears. One of these images was selected as a "base" picture, a picture that was later modified, and the other designated as a "reference" picture, which was left unchanged. Additional 100 frontal pictures of Caucasian male faces, with no glasses or facial hair, were taken from the Color FERET database (Phillips, Wechsler, Huang, & Rauss, 1998) and cropped in the same way.

2.1.3. Face tagging: converting faces into feature vectors, and measuring face-space distances

In our previous study we described faces as feature-vectors embedded in a multidimensional feature space. We showed that by perceptually assigning values to a set of 20 features, we can measure distances between faces, and these distances were correlated with perceptual face similarity scores. In this study, we repeated this procedure with familiar faces and converted each one of the faces in our database into a feature-vector representation (see Fig. S1 for an example of feature-vectors of two celebrity faces). For the 100 faces from the color-FERET dataset we used the feature-vector representations obtained in our previous study (Abudarham & Yovel, 2016). For tagging the ten celebrity faces we ran a face-tagging procedure. To provide participants with a large enough dataset for tagging, allowing them to judge facial features with respect to a variance of feature sizes and shapes, we created a dataset of 60 face images. These 60 images included the selected 10 celebrity faces, 20 pictures of other celebrities of similar characteristics as the original 10, and 30 randomly selected pictures from the 100 color-FERET dataset. In the tagging procedure, participants were asked to rank each of the 20 features for each of the sixty faces on a scale of −5 and +5 (for example: how thick are the