

Personally familiar faces: Higher precision of memory for idiosyncratic than for categorical information

Isabelle Bühlhoff

Max Planck Institute for Biological Cybernetics

Mintao Zhao

Max Planck Institute for Biological Cybernetics &
University of East Anglia

Keywords: face recognition; familiarity; identity; caricature; race; gender; identification

Note: this manuscript is accepted for publication in Journal of Experimental Psychology: Learning, Memory, and Cognition on 15th September 2019.

Abstract

Many studies have demonstrated that we can identify a familiar face on an image much better than an unfamiliar one, especially when various degradations or changes (e. g. image distortions or blurring, new illuminations) have been applied, but few have asked how different types of facial information from familiar faces are stored in memory. Here we investigated how well we remember personally familiar faces in terms of their identity, gender, and race. In three experiments, based on the faces personally familiar to our participants, we created sets of face morphs that parametrically varied the faces in terms of identity, sex or race, using a 3-dimensional morphable face model. For each familiar face, we presented those face morphs together with the original face and asked participants to pick the correct “real” face among morph distracters in each set. They were instructed to pick the face that most closely resembled their memory of that familiar person. We found that participants excelled in retrieving the correct familiar faces among the distracters when the faces were manipulated in terms of their idiosyncratic features (their identity information), but they were less sensitive to changes that occurred along the gender and race continuum. Image similarity analyses indicate that the observed difference cannot be attributed to different levels of image similarity between manipulations. These findings demonstrate that idiosyncratic and categorical face information is represented differently in memory, even for the faces of people we are very familiar with. Implications to current models of face recognition are discussed.

Introduction

Many studies have demonstrated that we recognize familiar and unfamiliar faces differently (Armann, Jenkins, & Burton, 2016; Bruce, 1982; Burton, Wilson, Cowan, & Bruce, 1999; Megreya & Burton, 2006; Ramon & Van Belle, 2016). Recognition of familiar faces is very robust. We can easily recognize a familiar face, even after 20 years or despite dramatic appearance changes such as added glasses, moustache, or new hair color (Bahrnick, Bahrnick, & Wittlinger, 1975; Russell, Duchaine, & Nakayama, 2009). We can also recognize familiar famous faces even when facial images are severely distorted (Harmon, 1973; Sinha, Balas, Ostrovsky, & Russell, 2006; Yip & Sinha, 2002). In contrast, we are usually bad at pairing an unfamiliar face to another picture of the same face, even when they are shown under very similar viewing conditions (Bruce et al., 1999; Hancock, Bruce, & Burton, 2010; Megreya & Burton, 2006; Young & Burton, 2018). This is true even for professionals who have been working on recognizing faces for years (e.g., passport officers, White, Kemp, Jenkins, Matheson, & Burton, 2014). Recent studies suggest that for unfamiliar faces, machine vision systems even perform better than humans do (Dehon, Bredart, & Brédart, 2001; Phillips et al., 2007; Schroff, Kalenichenko, & Philbin, 2015; Taigman, Yang, Ranzato, & Wolf, 2014).

Differences between the processing of familiar and unfamiliar faces not only manifest in how well people can recognize them, but also lie in different features that people use for encoding and recognizing faces. For instance, it has been shown that people use primarily the inner facial features for recognizing familiar faces, whereas for recognizing unfamiliar faces they pay additional attention to extra-facial information like hairdo, glasses or beards (Ellis, Shepherd, & Davies, 1979; Hancock et al., 2010; Johnston & Edmonds, 2009; Osborne & Stevenage, 2008). The reliance of more superficial, pictorial features for recognizing unfamiliar than familiar faces sometimes can be an advantage for unfamiliar

Memory of personally familiar faces

face processing. For example, people are better at remembering specific images of unfamiliar faces than familiar faces (Armann, Jenkins, & Burton, 2016; Ramon & Van Belle, 2016; but see Dunn, Ritchie, Kemp, & White, 2019). Note that when facial features are altered by displacing them or by stretching the face images, people are not always better at detecting or correcting these changes in familiar faces than in unfamiliar faces (Burton, Schweinberger, Jenkins, & Kaufmann, 2015; Ramon & Van Belle, 2016; Sandford & Burton, 2014).

Familiar and unfamiliar faces are also processed differently in the brain (Gobbini & Haxby, 2006; Natu & O'Toole, 2011, 2015; Ramon, Vizioli, Liu-Shuang, & Rossion, 2015; Visconti Di Oleggio Castello, Halchenko, Guntupalli, Gors, & Gobbini, 2017). Recent studies found that core face-selective brain areas—the fusiform face area (FFA) and the occipital face area (OFA)—show different patterns of neural responses to familiar and unfamiliar faces, both temporally and spatially (Natu & O'Toole, 2015). Subcortical areas, such as the amygdala, also respond differently to familiar and unfamiliar faces (Ramon et al., 2015, see also Platek & Kemp, 2009). The hippocampus also tends to show stronger neural responses to familiar than unfamiliar faces (e.g., (Leveroni et al., 2000; Ramon et al., 2015). These results indicate that familiarity with faces changes face representation in both the core and the extended brain networks related to face processing.

Although there is a general consensus that people process familiar and unfamiliar faces differently (e.g. Bruce, 1982; Burton et al., 1999; Megreya & Burton, 2006; Young & Burton, 2018), how familiar faces are represented in memory remains unsettled. For instance, some studies have argued that the idiosyncratic features of familiar faces are exaggerated in memory (P. P. W. Chang, Levine, & Benson, 2002; Gillian Rhodes, Brennan, & Carey, 1987). These studies showed that children and adults tend to select caricaturized face images (i.e.,

Memory of personally familiar faces

augmenting identity-specific features) as the best likeness of familiar persons' faces. In contrast, other studies indicated that these identity-diagnostic features are weakened in the representation of familiar faces (Allen, Brady, & Tredoux, 2009; Benson & Perrett, 1991, 1994; Frowd, Bruce, Ross, McIntyre, & Hancock, 2007; Kaufmann & Schweinberger, 2008). This line of studies often found that people rate the anti-caricatures (i.e., faces with their idiosyncratic features transformed toward an average appearance) of familiar faces as their best likeness. Finally, there is also research suggesting that familiar faces are represented in a relatively veridical way (Lee & Perrett, 2000; Srismith, Zhao, & Bühlhoff, 2016) with participants giving the original faces the highest rating in terms of best likeness. Such discrepant results about the representation of familiar faces may be partially due to methodological differences [e.g., (anti-)caricaturization was applied to line drawings, 2D photographs, or 3D facial shapes; different ranges/extents of (anti-)caricaturization, etc.] and partially due to different ways of defining familiar faces (e.g., famous faces vs. personally familiar faces).

In contrast to extensive research on how identities of familiar faces are represented in memory, little is known about how we represent categorical information of familiar faces (i.e., information that is not specific to an identity, such as race, gender, or age).

Independently of how familiar they are to the observers, faces can be classified into general categories based, for example, on their unchanging characteristics (gender and race appearance) or their changing characteristic (age, expression). Such categorical facial information is often integrated with the identity information (Zhao & Hayward, 2013) and can be automatically and incidentally acquired during identity learning. Previous studies have suggested that faces are also represented along those categorical characteristics (Armann, Jeffery, Calder, & Rhodes, 2011; Jaquet, Rhodes, & Hayward, 2008; Rhodes et al.,

Memory of personally familiar faces

2011), however, whether the representation of categorical information differs between familiar and unfamiliar faces, and, if yes, in which manner, remains unaddressed.

Recently, Dobs, Isik, Pantazis, and Kanwisher (Dobs, Isik, Pantazis, & Kanwisher, 2019) have shown that familiarity with faces enhances the representation of gender and identity, but not the age of faces. Moreover, they found that the enhanced representation of gender and identity even occur at an early stage, that is before gender and identity information being extracted. Such a finding may help revealing the mechanism underlying faster gender categorization for familiar than unfamiliar faces (e.g., Rossion, 2002) and, when searching for a target face, why slower rejection of same-gender than different-gender foil faces (Baudouin & Tiberghien, 2002) occurs. Nonetheless, representational similarity analysis of Dobs and colleagues (Dobs et al., 2019) on MEG responses to faces cannot tell in which way the representation of gender and identity of familiar faces is enhanced (e.g., does an enhanced gender representation mean that a male face is represented as more masculine than it is really or as a more accurate representation of its gender along the feminine-masculine continuum?). In addition, it remains unspecified whether familiarity enhances similarly the representation of identity and gender.

In the present study, we investigated how information about familiar faces is kept in memory, and, more specifically, how well people store idiosyncratic (e.g., identity) and categorical (e.g., gender and race) information of very familiar faces in memory. While identity information refers to diagnostic information related to a single specific person, gender and race are prominent categories shared by different identities. We concentrated on those aspects, as they represent natural characteristics that we use most commonly to describe or classify faces. We tested both identity and gender (or race) memory to obtain comparable assessments of how a familiar person's face is represented in memory.

Memory of personally familiar faces

Therefore, our findings may not only help reconcile previous discrepant results about familiar identity representation, but also address the long overlooked question of whether idiosyncratic and categorical information of familiar faces are similarly well represented in memory.

We focused our study on personally familiar faces. While unfamiliar faces can be simply defined as faces that people have not seen before, familiar faces often refer to different types of faces in the literature (Natu & O'Toole, 2011; see also Gobbini & Haxby, 2007). In some earlier studies familiarization was based on learning one single image per face (e. g. Bühlhoff & Newell, 2004), whereas in more recent studies *visually familiar* faces were obtained by asking participants to learn several different pictures of a person (e. g. Bonner, Burton, & Bruce, 2003). While the newly acquired visual familiarity with faces can be well controlled, it does not resemble well how we usually get familiar with faces in daily life (Burton, 2013). *Famous faces* have been frequently used as familiar faces in the literature (Harmon, 1973; Sinha et al., 2006; Yip & Sinha, 2002). These faces often belong to well-known actors or politicians, thus participants have, prior to the experiment, already gathered some personal and visual information about them in addition to their pictures used in the lab. In comparison with famous faces, we have more interactions with *personally familiar faces* (e.g., family members, friends, and colleagues). In this study we used personally familiar faces because they represent the highest level of familiarity participants can achieve with other persons' faces (Taylor et al., 2009). Furthermore they are the faces that people remember best, which allows us to test very precisely how identity and categorical facial information is represented in memory.

To directly demonstrate how idiosyncratic and categorical information of familiar faces are represented in memory, we adopted the classical "line-up" paradigm (e. g. Bruce et

Memory of personally familiar faces

al., 1999; Hancock et al., 2000; Megreya & Burton, 2006) with newly created line-up face sets.

In the classical paradigm, mostly *unfamiliar* target faces were presented in a line-up in which distracters are faces from other identities. Participants' task was to choose the face with the correct identity in the line-up. With this paradigm, these studies demonstrated how difficult recognition of unfamiliar faces is.

In the present study, participants also had to find out the correct face in a line-up, but there were three important differences from the classical paradigm. First, our participants were all personally familiar with the persons whose faces we used. Second, instead of matching two images (a learned image and a different test image of the same person), participants' task was to find one face in the line-up that best matched their memory of a very familiar face. Third, each line-up consisted of faces that were variations around one identity, rather than faces of completely different identities. More specifically, the distractors in the line-up trials were created by parametrically modifying each familiar face along the natural perceptual dimensions of gender, race or identity. Therefore, our modified line-up task not only required participants to make a binary decision about whether or not a face image depicts a familiar person's face, but also asks for their fine-grained judgment about how well each face image in a set represents the familiar face. By looking at how often different face variations were chosen as the best match to their memory of familiar faces, we would be able to directly demonstrate how a familiar face stored in memory may look like, and how precise each dimension of familiar faces—identity, gender and race—is represented in memory.

General methods

Stimuli

Our face stimuli were derived from original faces using the multidimensional face space framework (Valentine, 1991). According to this framework, each face is represented as a point in a multidimensional space, with dimensions referring to specific characteristics along which faces may vary. The average of all faces is situated in the center of the face space (Figure 1a), and each face is coded in relation to the average face in this space (Leopold, Bondar, & Giese, 2006; Leopold, O'Toole, Vetter, & Blanz, 2001).

We used two ways to manipulate the idiosyncratic properties of faces: morphing between identities and (anti)caricaturizing faces. Figure 1a depicts a schematic illustration of these two manipulations in a simplified 2-dimensional face space. For morphing identity between two faces (here between face B and face C or between face B and face D), we used the vectors connecting these faces in face space. Every aspect that differs between B and C or B and D (e.g., shape of the nose, color of the skin, distance between the eyes, etc.) changes gradually along the vector, so that face morphs approaching one face display increasing similarity to that face identity and decreasing similarity to the other face. Similarly, we can create (anti)caricaturized face morphs by using the vector connecting the original face (e.g., face A in Figure 1a) and the average face (Jiang, Blanz, & O'Toole, 2007; Leopold et al., 2001). The more a face morph is located near the average face, the less its idiosyncratic properties (all aspects of a face that differ from the average face) are preserved (i.e., anti-caricaturization). Conversely, the further away a face morph is from the average face, the more the idiosyncratic properties of the original are amplified (i.e., caricaturization).

Figure 1b illustrates how we manipulated the categorical information of faces (e.g., gender or race). Here the gender vector connects the average of all male faces and the

Memory of personally familiar faces

average of all female faces; such a vector defines how the appearance of a face changes when its gender is manipulated. This gender vector can be applied in full or in part to any face to modify it (Bülthoff & Newell, 2004). Note that it allows us not only to make a male face more feminine, or a female face more masculine, but it can also be used to masculinize a male face or to make a female face more feminine. We manipulated the race of a face in the same manner, using a race vector connecting the average Asian and Caucasian faces.




Figure 1. Schematic illustration of stimuli creation. (a) A norm-based face space with examples of caricature and identity morphs. Faces are represented as black dots and the axes represent two exemplar dimensions of the space. The average of all faces is located at the center. The gray dashed line represents the axis of caricaturization for face A, with the pointing-up triangle representing an anti-caricature and the pointing-down triangle a caricature of face A. The solid lines between faces B, C and D represent axes of identity morphs with the 50% morph level located at the middle of each axis. (b) Representation of the gender dimension (black arrow) in face space. The vertical dashed line represents the gender boundary between female and male faces. The empty dots represent the average of all male and all female faces, and the line connecting them defines the gender vector. Stars are exemplar gender morphs of the male face C obtained by applying the gender vector. They are more male or female looking than the original face.

Our stimuli were generated using three-dimensional laser scans of our colleagues at the Max Planck Institute for Biological Cybernetics. A description of the scanning method has been given in previous studies (Blanz & Vetter, 1999; Bülthoff & Newell, 2004; O'Toole, Vetter, Volz, & Salter, 1997; Troje & Bülthoff, 1996). For each of these scanned faces (i.e.,

Memory of personally familiar faces

target face), we created different sets of face stimuli: identity set, caricature set, gender set, and race set (see details and examples below). Each face set consisted of 11 faces (one target and ten morph images as distracters). The target was always in the center of the morph sequence (see below). We used the 3D morphable model to generate variations of the target faces by gradually manipulating either their identity, gender, or race appearance (Blanz & Vetter, 1999; Troje & Bühlhoff, 1996; <http://faces.kyb.tuebingen.mpg.de>; see also O'Toole et al., 1997). This model allows to change the identity, gender, or race of faces by modifying both shape and texture information simultaneously. We name *original* or *target* faces, the images derived from the laser scans without any manipulations; there is one original face per identity. The others images are all obtained with manipulations along one of the four dimensions are called *distracters* or *morphs*.

Identity sets. Each target was morphed with two unfamiliar faces in 10% increments, up to 50%; so that the endpoint faces were a mixture of half target face and half-unfamiliar face (i.e., located midway between the target identity and one of the two unfamiliar face identities in face space; **Figure 1a and Figure 2a**). Except for two faces used twice, all other unfamiliar faces were used once. Those unfamiliar faces were of the same gender, had a similar age, and were visually similar to the target face they were paired with.




Figure 2. Examples of face stimuli used in Experiment 1a. We show only 7 of the 11 face morphs of each set. **(a)** Identity set; 50% morphs are halfway between the target and an unfamiliar face. **(b)** Caricature set; 50% anti-caricatures are halfway between the target and the male average face, 50% caricatures are at the same computational distance from the original face. **(c)** Gender set, 50% feminized morphs are halfway between the target and their female corresponding faces; 50% masculinized morphs are at the same computational distance from the target. The morph percentages refer to how much each morph differs computationally from the original face (target) along different dimensions.

Caricatures sets. Each original face was morphed toward the average face to create anti-caricature morphs and was morphed away from the average face to create caricatures. Specifically, we first calculated the shape and texture differences between the target face and the average face. We then reduced these differences in steps of 10%, up to 50%, to create anti-caricatures and amplified them the same way to create caricatures (see **Figure 1a** and **Figure 2b**). We used the average of male or female faces for male and female target faces, respectively.

Memory of personally familiar faces

Gender sets. The gender sets were created by varying the facial masculinity/femininity (Bülthoff & Newell, 2004). We first defined the gender vector as the differences between average male and female faces. This vector specifies how the shape and texture values at each vertex of a scanned face vary when the gender is modified. We then applied this vector to each face (male or female) in steps of 10%, up to 50%, to create increasingly more male and female looking face morphs (**Figure 1b and Figure 2c**).

Race sets. The race sets were created in the same way as the gender sets. We first defined a race vector as the difference between averaged Asian and Caucasian faces. We then used the race vector to modify the racial appearance of each original face in 10 percent steps up to 50 percent, making five morphs that look more Asian and five more Caucasian than the original face (**Figure 7d, e, f**).

Test trial. All target and morph images were presented in color and turned 20 degrees to allow a better view of the shape of the nose and jaw. Each trial depicted all 11 images of one set derived from one target face and varied in only one dimension (e.g., identity, race, caricature or gender). The 11 face images were randomly arranged on a 4 by 3 numbered grids (**Figure 3**). The grid number 8 was always used for numbering the trials. The alignment of the faces within each grid was jittered to hinder participants in using simple heuristics like comparing face sizes.



Figure 3 Example of a test trial for an identity set. The target is face 3; faces 6 and 1 differ most (i.e., 50%) from the target. The large number 4 refers to trial number.

Design and Procedure

We ran four experiments over several years with our colleagues at the Max Planck Institute for Biological Cybernetics. Participants completed the experiment as a group in a lecture hall. Trials in each experiment were blocked by facial manipulations (e.g. gender or identity), but participants were not informed about this. The order of blocks and trials were fixed as all participants performed the test together. Each trial was shown only once.

Before the experiment, all participants were given a response sheet to fill in their responses. The trial numbers with the corresponding names of the original faces were printed on the response sheet. For each trial, participants were required to write down which face they believed to be the original face of their colleagues, and to guess if they were not sure. They also indicated how familiar they were with the person by ticking one of

Memory of personally familiar faces

three possible choices (I know him/her very well; I know who the person is; I don't know him/her). Participants did not need to enter any answer for faces they did not know or their answers were excluded from analysis.

All trials were displayed on the projection screen of a lecture hall using a PowerPoint presentation. Participants seated at various distances from the screen ranging between 4 and 8 m. The width of each face image was about 40 cm. Each trial was shown for 25s followed by a blank screen, during which participants entered their answers. Participation was voluntary and fun.

Experiment 1a: Memory of identity and masculinity of familiar male faces

To investigate how well people represent very familiar faces in memory, we tested participants' recognition performance on identifying personally familiar faces (i.e., faces of their colleagues) among finely modified distracter faces. These distracters were created by morphing the target faces along different dimensions¹ in a multidimensional face space (Valentine, 1991, see Figure 1): identity, caricature, and gender. The identity morphs are also called oblique or lateral caricatures (Rhodes & Tremewan, 1994) because, compared to morphs created by caricaturing, they are located along a trajectory that does not pass through the norm (i.e., the centre of the face space). Both the identity and the caricature morphs modifies the idiosyncratic information of a face (Lewis & Johnston, 1998; Rhodes, 1996; Rhodes, Hickford, & Jeffery, 2000; Rhodes & Tremewan, 1994), whereas enhancing/weakening the femininity/masculinity modifies categorical information of facial gender (see Armann & Bülthoff, 2009; Bülthoff & Newell, 2004).

¹ This is a simplified view, we are aware that gender and race are semantic concepts that are underlined by multiple dimensions.

Memory of personally familiar faces

In Experiment 1a, we tested memory of identity and masculinity for personally familiar male faces. We created identity set, caricatures set, and gender set based on 3D face scans of ten male colleagues, resulting in three test trials for each target face and 30 trials in total (see **Figure 2** for examples). Thirty-eight colleagues (4 female) volunteered to participate in Experiment 1a. They first completed the identity sets, then the caricature sets and lastly the gender sets.

Results

For 1140 trials across all participants and all face sets, there were 105 non-responses (9.2%; 37, 33, and 35 non-responses for the identity, caricature, and gender sets respectively). For each face manipulation, we calculated how often (in percentage) each morph level was chosen as the target person and tested whether the response differed from chance level (i.e., 1/11) using a Bonferroni-corrected one-sample t-test (i.e., $\alpha = .005$). The results are shown in **Figure 4** (left column).

Identity sets

When target identity was morphed toward unfamiliar face identities, participants chose the target faces more often than any face morphs (**Figure 4a**). The more a morph was close to an unfamiliar face, the less it was selected as the target person, resulting in a symmetrical choice distribution peaked at the target face. Consistent with this observation, one-sample t-tests revealed that the target and its closest neighbours (10%) were chosen significantly more often than the chance level (chance level: 1/11; $ts(37) > 5.21$, $ps \leq .001$), whereas the 3 most extreme morph levels on each side were chosen significantly under chance level ($ts(37) > 4.91$, $p < .001$). These results indicate that participants form an accurate identity

Memory of personally familiar faces

representation of personally familiar faces and are sensitive to identity changes when familiar faces are morphed with unfamiliar faces.

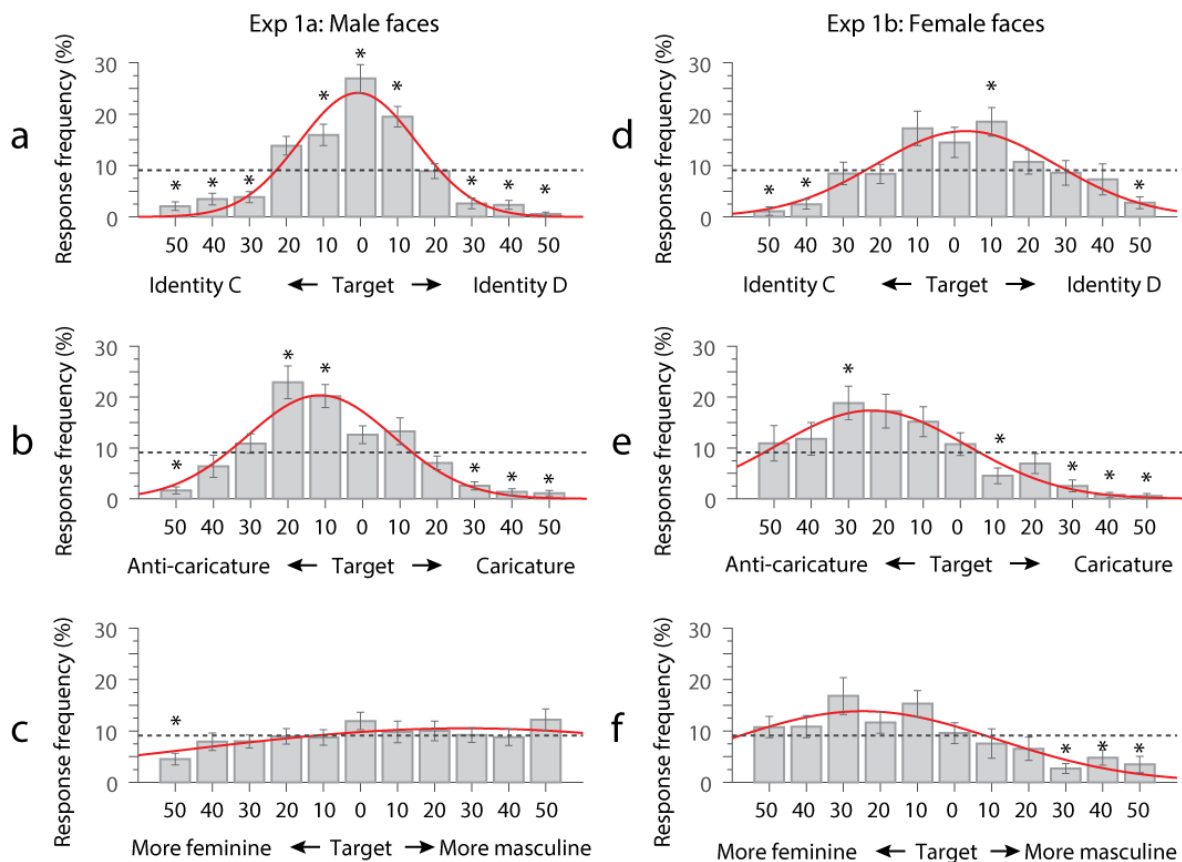


Figure 4. Results of Experiment 1a and 1b. Mean frequency of choices to male (left column) and female (right column) faces plotted as a function of morph level. Top row (a, d) shows responses to identity sets, identities C and D refer to two unfamiliar faces used to create the morph, which are corresponding to the labels in Figure 1a. Middle row (b, e) shows responses to caricature sets. Bottom row (c, f) shows response to gender sets. Asterisks indicate that the responses to the specific morph levels differed significantly from chance level (dashed line). Curved solid lines represent best Gaussian distribution fit.

Caricature sets

When caricatures and anti-caricatures were placed as distracters alongside the targets, participants showed a skewed pattern of response distribution that peaked at the 20% anti-caricatures. Further, participants were biased to choose more often anti-caricatures than

Memory of personally familiar faces

caricatures (**Figure 4b**). One-sample t-tests revealed that the 20% and 10% anti-caricature levels, but not the targets, were chosen significantly more often than chance ($t(37) > 4.33$, $p < .001$). Responses at the 50% anti-caricature and the three most extreme caricature levels were significantly under chance level ($t(37) > 8.45$, $p < .001$). These results show that the representation of familiar face identity is shifted toward the average faces, and that participants are similarly sensitive to the changes of identity strength (as suggested by the peaked distribution of responses).

Gender sets

When the gender of the faces was manipulated, participant showed no significant preference to select any face as the target person, though they showed a slight bias to choose masculine faces over feminine faces (**Figure 4c**). One-sample t-tests revealed that response choices were not different from chance except for the most feminine level (50%, $t(37) = 4.03$, $p < .001$). These results indicate that participants may not form an accurate representation of facial gender, and are therefore not very sensitive to the manipulation of face gender, even for personally familiar faces.

Comparison of the response patterns for identity, caricature, and gender sets

To compare patterns of response across the three face manipulations, we performed three tests. Firstly, we tested whether the distribution of responses was uniform or not. Secondly, we tested whether there was a robust peak response. A robust peak was defined as a morph level that received more choices than both the averaged choices for all the morphs on its left and right sides. Lastly, we tested whether responses were skewed toward one side of morph continuum over the other. The results are summarized in **Table 1**.

Table 1 Summary of the analyses of all three experiments

<i>Manipulation</i>	Experiment 1a: Male faces			Experiment 1b: Female faces		
	<i>Identity</i>	<i>Caricature</i>	<i>Sex</i>	<i>Identity</i>	<i>Caricature</i>	<i>Sex</i>
Chi square: unequal distribution	Yes	Yes	No	Yes	Yes	Yes
T tests: robust peak (location)	Yes (target)	Yes (anti-caricature)	No	Yes (both flanking morphs)	Yes (anti-caricature)	Yes (feminine)
T test: Biased choice distribution	No	No	No	No	Yes, toward anti-caricature	Yes, toward same sex

<i>Manipulation</i>	Experiment 2: Inverted faces	Experiment 3: Identity and race			
	<i>Identity</i>	<i>Identity*</i>	<i>Race male</i>	<i>Race female</i>	<i>Race Asian</i>
Chi square: unequal distribution	No	Yes	No	Yes	No
T tests: robust peak (location)	No	Yes (target)	No	No	No
T test: Biased choice distribution	No	No	Yes, toward same race	Yes, toward same race	Yes, toward same race

* the same results were obtained for male Caucasian, female Caucasian and male Asian faces

Responses to the identity sets showed a peaked and symmetrical pattern (Figure 4a). A chi-square test revealed a non-uniform distribution of responses ($\chi^2 = 136.9$, $df = 10$, $p < .001$). Participants showed a peaked response at the target face, which was chosen more often than the average choices of the morphs on either sides (one-sample t-tests, both $ts(4) > 5.19$, both $ps \leq .004$). Moreover, morphs on both sides of the targets were chosen equally often on average ($t(4) = .82$, $p = .46$). The averaged morph level of all responses (2.2%) was very close

Memory of personally familiar faces

to the targets (i.e., 0%, the middle of the morph continuum). This provides further support for a symmetrical pattern of responses.

Responses to the caricature sets showed a peaked but skewed pattern (Figure 4b). Responses were not equally distributed across all morph levels ($\chi^2 = 120.386$, $df = 10$, $p < .001$). Response distribution peaked at the 20% anti-caricature level, which was chosen more often than the average of all caricature or anti-caricature choices (both $ps \leq 0.019$). The average over all chosen morph levels is 8.6% toward anti-caricatures, nonetheless, the bias to choose anti-caricatures over caricatures was only marginally significant in a paired t-test ($t(4) = 2.67$, $p = .056$).

Responses to the gender sets showed a flat distribution without a clear peak (Figure 4c). A chi square test revealed a uniformly distributed responses across all morph levels ($\chi^2 = 7.700$, $df = 10$, $p = .66$). The average of all chosen morph levels is 4.9% more masculine than the target. Again, this preference to choose masculine over feminine morphs did not attain statistical significance ($t(4) = 1.90$, $p = .13$).

Discussion

Experiment 1a shows two main findings. First, participants seemed to form a more precise representation for face identity than for facial gender. Participants were sensitive to the amount of manipulations applied to idiosyncratic information, resulting in a peaked choice distribution for both the identity and the caricature sets. In contrast, participants showed no peak of choices around the targets or at any morph level in the gender sets. This flat distribution of the responses across gender morphs indicates that participants did not develop a precise representation of the femininity/masculinity of very familiar male faces. Second, while participants showed a general agreement about which face best resembles their memory of a familiar identity, their memory representation is very accurate when the

Memory of personally familiar faces

distractors are created by morphing with other identities but it is biased toward anti-caricatures when the distractors vary in identity strength (i.e., we observe a bias in caricatures sets).

Our findings for caricature sets are consistent with previous studies (Allen, Brady, & Tredoux, 2009; Benson & Perrett, 1991, 1994; Frowd, Bruce, Ross, McIntyre, & Hancock, 2007), which also report that participants tend to judge anti-caricatures to be more similar to the face in memory than the target or any caricatures. It is worth noting that we manipulated both facial shape and texture information in 3D spaces to create (anti)caricatures of target faces whereas the earlier study modified shape only. These findings are in contrast to the idea that caricatures are better representation of an identity (Benson & Perrett, 1994; Rhodes, 1996; Rhodes, Byatt, Tremewan, & Kennedy, 1997). Overall the general consensus is that the use of line representations in those older studies might have resulted in this caricature advantage as newer studies using photographs which offer more texture and shape information report in general an advantage for anti-caricatures (Benson & Perrett, 1991; Chang et al., 2002; Lee & Perrett, 2000). We would like to emphasize that we have caricaturized the texture (skin appearance) along with the shape. This process has created an increasingly unnatural blotchy appearance of the skin with increasing levels of caricaturing while anti-caricatures displayed a smoother appearance. This might have pushed participants' choices toward anti-caricatures over choosing the target or any caricatures. This hypothesis is substantiated by the results of a follow-up study, (Srismith et al., 2016) in which we caricatured only the shape of personally familiar faces, not their texture. Participants in a similar task chose mostly the veridical faces over anti-caricatures and caricatures.

Memory of personally familiar faces

Experiment 1b: Memory of identity and femininity for familiar female faces

Experiment 1b tested whether different representations of face identity and face gender can be generalized to female familiar faces. We did not include female faces in Experiment 1a because we only have very limited number of female colleagues known to most colleagues at that time. In Experiment 1b, we conducted the same experiment as Experiment 1a, but with stimuli derived from the faces of seven female colleagues². This resulted in 3 test trials for each target face and 21 trials in total. Thirty-eight colleagues (10 female) participated in the experiment. They saw first the caricature sets, followed by the identity sets and last the gender sets.

² We had only so many female coworkers that were known to most participants

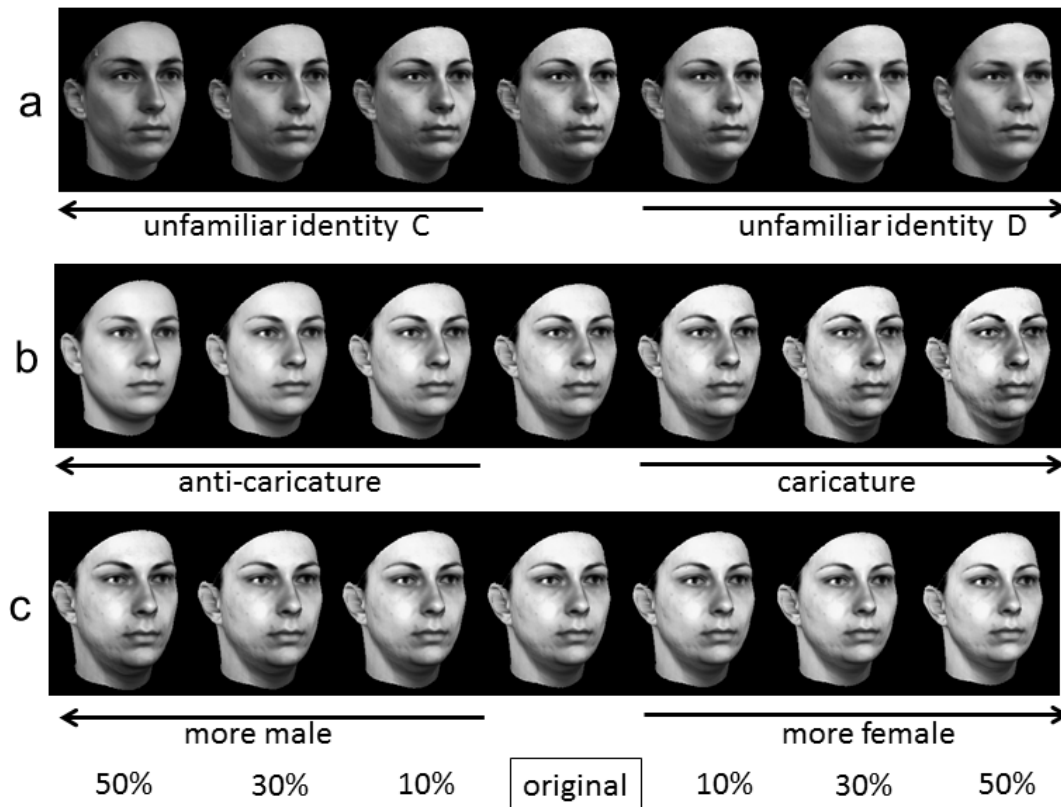


Figure 5. Examples of face stimuli used in Experiment 1b. (a) Identity set; (b) Caricature set; (c) Gender set. The morph percentages refer to how much each morph differs computationally from the original face (target) along the specific dimension.

Results

About 20% of 798 trials across participants and face sets received no response (163 trials, 51 for each of the three face manipulations). The results are similar to those observed in Experiment 1a (Figure 4, right column).

Identity sets

Responses to identity sets showed a symmetric distribution centred on the targets and their nearest neighbours (i.e., 10% morphs, Figure 4d). The targets were chosen more often than all morph levels except for their two flanking 10% morph levels. One-sample t -tests revealed that the targets were not chosen significantly more often than chance ($t(37) = 1.87, p = 0.07$), instead, one of the two flanking morph levels was chosen significantly more

Memory of personally familiar faces

often than chance ($t(37) = 3.14, p < .001$)³. Three of the extreme morph levels were significantly under chance level ($ts(37) > 5.65, ps < .001$).

Caricature sets

Participants were generally biased to choose anti-caricature distracters over the targets, with a peak response at the 30% anti-caricature level (**Figure 4e**). The representation of familiar face identities appears to be shifted towards the average face. One-sample t-tests revealed that only this 30% anti-caricature level was chosen more often than chance ($t(37) = 2.97, p = .005$). Four of the five caricatures were chosen significantly less often than chance ($ts(37) > 5.78, ps < .001$).

Gender sets

Participants displayed a preference for more feminine looking face morphs than the targets and the choice distribution appeared to peak at the 30% more feminine morph level (**Figure 4f**). Nonetheless, one-sample t-tests revealed that no morph level was chosen significantly more often than chance (all $ts(37) < 2.19$, all $ps > .035$). The three most masculine morph levels were chosen significantly less often than by chance (all $ts(37) > 3.06$, all $ps \leq .004$).

Comparison of the response patterns for identity, caricature, and gender sets

None of the three response patterns showed a uniform distribution (identity sets, $\chi^2 = 136.9$, $df = 10, p < .001$; caricature sets, $\chi^2 = 425.0$, $df = 10, p < .001$; gender sets, $\chi^2 = 29.01$, $df = 10, p = .001$). However, they differed in terms of how each morph level was chosen as the target face.

³ The other flanking morph level was the only other morph level that differed significantly from chance when no Bonferroni correction was applied: $p = .02$

Memory of personally familiar faces

As summarized in Table 1, responses to the identity sets showed a symmetric pattern with a peaked plateau around the target face. The morphs next to the target faces on each side (i.e., 10% levels on both side) were chosen more often than the average choices of all other morphs towards either unfamiliar faces (all $ps \leq .023$), though the targets themselves were not (both $t(4) \leq 2.10$, $ps \geq .103$). Morphs on both sides of the targets were chosen equally often ($t(4) = 0.12$, $p = .91$), indicating a symmetrical distribution. This also concurs with the average morph level of 2.2% over all choices—a value very close to the targets (i.e., 0%).

Responses to the caricatures sets showed a peaked but asymmetrical distribution. The choices for the 30% anti-caricature was above the average choices for both sides (both $ts \geq 5.04$, both $ps \leq .038$). Participants showed a clear bias to choose anti-caricatures over caricatures ($t(4) = 3.15$, $p = .03$). This is in line with the finding that the averaged morph level of all responses was at the anti-caricature morph level of 18%.

Responses to the gender sets were also asymmetrical with a peak response at the 30% feminine level. Participants chose the 30% feminine level more often than the average choices on both the more masculine and more feminine sides (both $ps \leq .004$), though this peak performance itself was not significantly different from chance level (see above). Participants were significantly biased toward choosing more feminine than more masculine morphs ($t(4) = 5.22$, $p = .006$). Correspondingly, the average morph level of all choices is 12.1% more feminine than the target.

Discussion

Responses to familiar female faces are similar to the findings obtained with male familiar faces in Experiment 1a. Participants showed again a relative precise representation of idiosyncratic information, and their responses were sensitive to the changes in the identity

Memory of personally familiar faces

and caricature sets. For identity sets, participants chose female morphs next to the targets more often than all other morph levels in the identity sets, resulting in a peaked and symmetrical distribution around target faces. Surprisingly, in contrast to the male sets, the targets themselves were not the most often chosen morphs. For caricature sets, response distribution peaked around an anti-caricature level (i.e., 30% anti-caricature) a finding similar to what we observed for male sets. These peaked distributions suggest that all participants were sensitive to the amount of caricaturization applied to both male and female faces. In addition, representation of familiar female faces showed a stronger bias toward anti-caricatures, as indicated by both a peaked response at an anti-caricature level and an overall asymmetrical distribution biased to anti-caricature side (Figure 1, **Table 1**).

In the gender sets, in contrast to male faces for which we only observed a trend to choose more masculine faces, participants significantly preferred female faces with a more feminine appearance than the targets. The distribution of choices peaks at the more feminine morph 30% whereas there was no peak with the male faces. Importantly, for both male and female faces, the choice distribution for gender sets was more widely distributed across all morph levels than for the identity and caricature sets (see **Figure 4**). We observed larger number of morph levels showing chance level performance for the gender sets (10 and 8 out of 11 for male and female faces respectively) than for the identity and caricatures sets (identity sets: 2 and 7 out of 11 levels; caricatures sets: 5 and 6 out of 11 levels for male and female faces, respectively). Overall, we observed a trend of enhanced gender representation for both familiar male and female faces: male faces look more masculine in memory whereas female faces look more feminine in memory (**Figure 4c** and **4f**).

Experiment 2: Role of low-level information: Familiar faces presented upside down

Experiment 2 investigated whether the peaked responses observed for identity and caricature sets were due to a precise representation of face identity or was caused by low-level cues in the face images or by some participants' response strategies. To this end, we presented all test faces upside-down. Face inversion has been well documented to disrupt face recognition (among others: Hochberg & Galper, 1967; Rossion & Gauthier, 2002; Yin, 1969), but it does not influence low-level image-based face information. We tested the effect of face inversion with identity sets because they show the clearest peaked distribution around targets for upright faces in Experiment 1. If the peaked responses had resulted from the precise representation of face identity, it should be less visible for inverted faces.

However, if the peaked response patterns were rooted in certain low-level image features or strategies based on those features, they should remain irrespective of face orientation at test.

We used 7 male and 5 female identity sets from those used in Experiment 1, so that all target identity were familiar to most participants at test. A total of 38 colleagues (9 female) participated in the study. The procedure was the same as in Experiment 1 except for a longer presentation time of 30s for processing inverted faces. The 12 identity sets were blocked by gender during test, with male faces being tested before female faces.

Results

Sixty of 456 trials received no response (13.16%). The results for all male and female faces together are shown in Figure 6. Neither the targets nor any morph levels were chosen significantly more or less often than chance for both sexes analysed together or separately. A chi square test also revealed a uniform distribution ($\chi^2 = 2.55$, $df = 10$, $p = .99$). We performed the same analyses as in Experiment 1, none of them yielded any significant findings (Table

Memory of personally familiar faces

1). Correspondingly the average morph level of all choices was very close to the middle of the morph continuum (1.7 % away from the target).




Figure 6 Results of Experiment 2. Mean frequency of response choices to inverted faces in identity sets plotted as a function of morph level. Dashed line represents chance level performance. Curved solid lines represent best Gaussian distribution fit.

Discussion

Presenting the faces upside down resulted in a completely different distribution of responses compared to Experiment 1 where similar identity sets were presented in upright position. When the faces were presented upside down, participants distributed their responses equally across all morph levels, even though they were given more time to answer. This result indicates that the peaked responses obtained in Experiment 1 were driven by a precise representation of face identity rather than by low-level image similarity. When perception of identity was disrupted by face inversion, participants lost their ability to differentiate the targets from their identity variations. They did not rely on low-level (inversion-insensitive) cues in the images to perform the identity task more easily or precisely than the other tasks in Experiment 1. These findings extend the face inversion effect reported in numerous previous studies (e.g. Freire, Lee, & Symons, 2000; Yin, 1969). While previous studies have demonstrated that our ability to discriminate different individuals based on their faces is dramatically *reduced*, our study here shows that our ability to finely

Memory of personally familiar faces

differentiate two similar faces (e.g., identity morphs, or faces of identical twins) is totally wiped out by inversion (see also Robbins & McKone, 2003).

Experiment 3: Memory of identity and race of familiar faces

Experiment 3 tested whether the difference between representations of face identity and a face category (e.g., gender) applies to other categorical facial information (e.g., race). Experiment 1 showed that people have a relatively precise representation of familiar face identity with a fuzzy representation along the feminine/masculine dimension (i.e., gender). If familiarity with faces mainly sharpens the representation of idiosyncratic but not categorical facial information, other facial categories (e.g., race) should also exhibit such fuzzy categorical representation in face memory. To test this hypothesis, we manipulated personally familiar faces along the dimensions of identity and race (Asian vs. Caucasian), and tested whether the race of very familiar faces is also represented differently from face identity.

In Experiment 3, we created identity and race sets for ten Caucasian female, ten Caucasian male, and three Asian male colleagues⁴. All identities used were new except for one female face. The illumination and colour rendering of the faces was slightly different from the previous experiments. There were 46 test trials in total, two for each identity (**Figure 7**). A total of 35 colleagues (10 female) participated in the study. The procedure was similar to that of Experiment 1. Test trials were blocked by face manipulation (identity sets presented before race sets). For each block, participants saw first the male Caucasian faces followed by the female Caucasian faces and last the three Asian male faces.

⁴ We had only so many Asian colleagues that were known to most participants

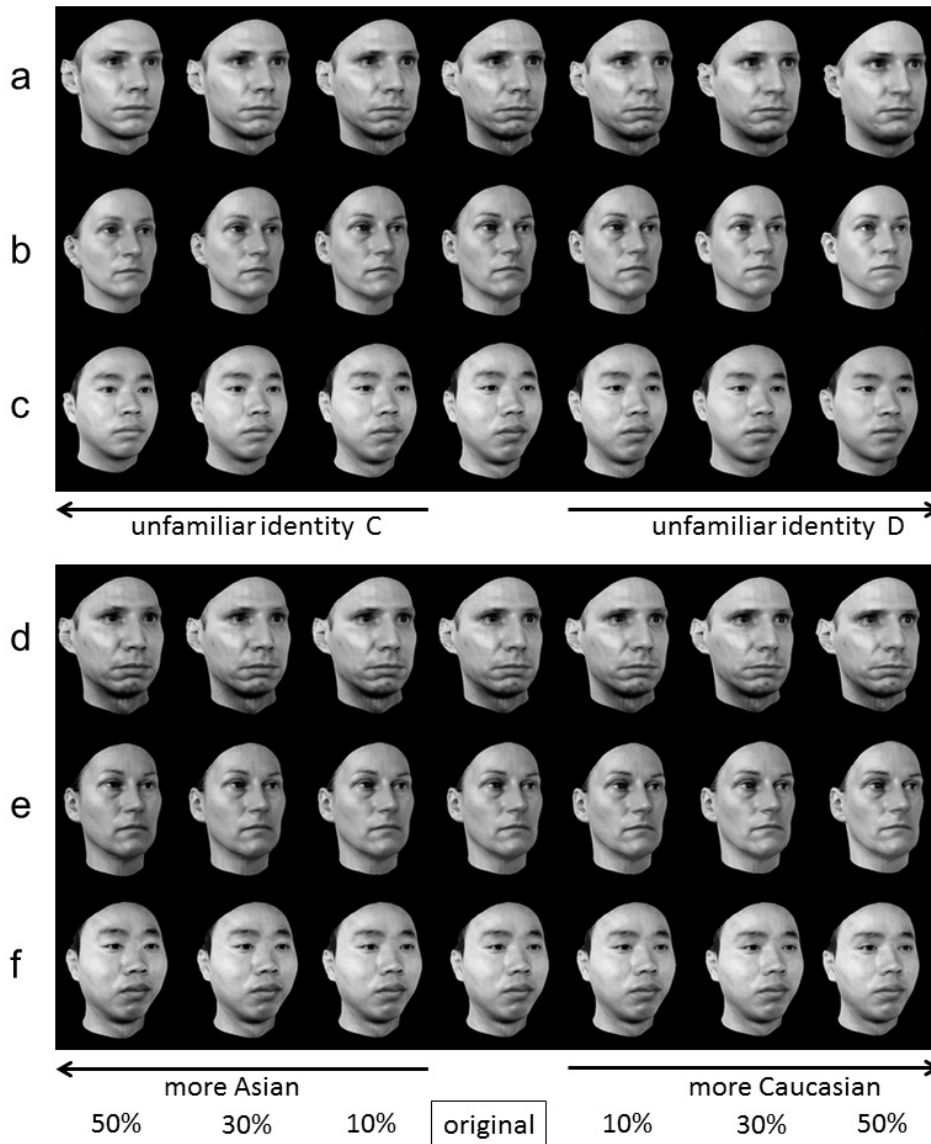


Figure 7. Examples of face stimuli used in Experiment 3. Identity sets for male Caucasian, female Caucasian, and male Asian faces are shown in rows a, b, and c, respectively. Corresponding race sets for these faces are shown in rows d, e and f, respectively. The morph percentages refer to how much each morph differs computationally from the original face (target) along a specific dimension.

Results

Overall, 90 out of 1610 trials received no response (5.6%; 42 for identity sets and 48 for race sets). The results are shown in **Figure 8**.

Identity sets

Responses to the identity sets showed a peaked and symmetrical distribution for both Caucasian and Asian faces (**Figure 8**, left column).

For male Caucasian faces (Figure 8a), all but both 20% morph levels were chosen significantly above or below chance level (all $t_{s(34)} > 4.37$, $p_s < .001$). For female Caucasian faces (Figure 8c), we found similar results except that the 30% morph level also showed chance level response (all other $t_{s(34)} > 3.98$, $p_s < .001$). For Asian male faces, the targets and the four most extreme morphs were chosen significantly above or below chance level ($t_{s(34)} \geq 2.64$, $p_s < .001$). Therefore, participants' choices were well concentrated on the targets for all three types of faces tested.

None of these three distributions were flat (all $\chi^2 \geq 35.72$, $df = 10$, $p_s < .001$). As summarized in Table 1, for both Caucasian and Asian faces, the targets were chosen more often than the average choices on both sides (all $t_{s(4)} \geq 4.79$, $p_s \leq .006$). Participants chose equally often the distracters on both sides of the targets (all $t_{s(4)} < 2.58$, all $p_s \geq .06$). The symmetric choice patterns were also suggested by the average morph level of all selected faces (0.1%, 1.0% and 0.3% for male Caucasian, female Caucasian, and male Asian faces respectively).

Memory of personally familiar faces




Figure 8. Results of Experiment 3. Mean frequency of response to male and female Caucasian faces and to male Asian faces in identity sets (left column) and race sets (right column) plotted as a function of morph level. Identities C and D refer to the unfamiliar faces used to create the morph. Asterisks indicate responses to morph levels (calculated across participants) that differ significantly from chance level (dashed line). Curved solid lines represent best Gaussian distribution fit.

Race sets

Responses to the race sets show a less peaked and more asymmetric distribution than observed with identity sets (Figure 8, right column). Participants favoured faces that displayed amplified own-race information, resulting in an asymmetrical distribution of responses. Participants chose most often the 20% and 30% more Caucasian morph levels for the male and female Caucasian sets respectively; whereas they chose the 20% more Asian morph level for male Asian faces.

Memory of personally familiar faces

For male Caucasian faces (Figure 8b), only the three most Asian-looking morphs were chosen significantly less often than chance ($t(34) \geq 2.35$, $ps \leq .001$), no other values diverged significantly from chance. For the female Caucasian faces (Figure 8d), the two most Asian-looking morphs were chosen significantly less often than chance ($t(34) \geq 9.98$, $ps \leq .001$), whereas the 30% more Caucasian morph was chosen significantly more often than chance ($t(34) = 3.87$, $p < .001$). For Asian faces (Figure 8f), only the 20% more Asian morph received significantly higher-than-chance level response ($t(34) = 3.04$, $p = .005$).

As summarized in Table 1, the distributions of choices for Caucasian and Asian males were flat (all $\chi^2 \leq 15.53$, $df = 10$, $p \geq .114$) whereas it was not flat for the Caucasian females ($\chi^2 = 18.79$, $df = 10$, $p = .043$). Although the peak faces in all groups were chosen significantly more often than the average of distracters toward the other-race end (all $ps \leq .006$), this was not the case for the comparison with other distracters displaying more same-race characteristics (all $ps \geq .093$). As such, these response maxima do not fulfil our criteria to be evaluated as robust distribution peaks. Participants showed a consistent response bias toward distracters that enhance the original race information for all three race sets (all $ps \leq .05$). For male Caucasian, female Caucasian, and male Asian target faces, the averaged morph levels of all choices were 13.1% more Caucasian, 13.4% more Caucasian and 6.5% more Asian, respectively.

Discussion

Experiment 3 reveals again target-centred patterns of responses to the identity sets as in Experiment 1. Note that the identity sets in both experiments differed in terms of the rendering techniques, targets identities, and were tested with different groups of participants. Experiment 3 revealed a robust response peak for both female and male targets

Memory of personally familiar faces

in the identity sets. The peaked distributions centred at the targets indicate that participants could identify very well the correct face of their colleagues among slight identity variations thereof and they were similarly sensitive to the amount of variation applied.

Response to the race sets is somewhat similar to the responses to gender sets observed in Experiment 1. The weakly peaked patterns of the choice distribution indicate that participants' choices were not centred on the targets but around distracters with enhanced own-race information compared to the targets. That is, familiar Caucasian faces are represented as more Caucasian-looking in memory, whereas familiar Asian faces are stored in memory as more Asian looking than the targets (**Figure 8**, right column). Again, the choice distribution for race sets was more distributed across all morph levels than that for the identity sets (see **Figure 8**). This is suggested by the larger number of morph levels showing chance level performance for the race sets (8/8/10 out of 11 for male Caucasian/female Caucasian/Asian male faces, respectively) than for the identity sets (2/3/6 out of 11 for male Caucasian/female Caucasian/Asian male faces, respectively). Furthermore, creating morphs along another category dimension than sex allows us to test whether the enhanced categorical information stored in memory is specific to the representation of gender information or is a more general principle that applies to other categorical facial information. The responses to race sets suggest that categorical facial information, but not the identity information of familiar faces is exaggerated in memory.

The results of Experiment 1 and 3 suggest that the representations of face identity and face categories differ. Peaks (if any) for the four category sets (male and female sex sets, male and female race sets) were less predominant than for identity sets. Furthermore, participants tended to choose faces with enhanced own-gender or own-race information. Interestingly, this bias occurred in both male and female faces although the participants were

Memory of personally familiar faces

predominantly male, and it occurred in both Asian and Caucasian faces although the participants were predominantly Caucasian. These observations suggest that such a category-enhancement bias in face representation is independent of face expertise (e.g., strong expertise of most participants with Caucasian faces).

Experiments 1-3: Role of image similarity

One may argue that the different patterns of responses to identity and gender or race changes may be due to different levels of image similarity between the targets and the morphs. For instance, the non-peaked response pattern that we found for gender sets might have resulted from all images (morphs and target) being very similar to each other within each set. To address this concern, we performed image-based similarity analyses and then examined whether our findings could be explained by different degree of image similarity between face manipulations.

We calculated the physical similarity between two facial images using a Gabor similarity analysis method (Lades et al., 1993). Prior studies have shown that Gabor similarity between facial images is correlated with perceptual similarity of facial identity (Yue, Biederman, Mangini, Malsburg, & Amir, 2012), expression (Xu & Biederman, 2010), and facial movements (Dobs et al., 2014). Gabor similarity between two facial images was computed with the following steps (see Dobs, Bühlhoff, & Schultz, 2016; Yue et al., 2012 for more details). First, all images were converted into grayscale images (256 by 256 pixels). Second, we filtered each image using a Gabor jet [5 scales \times 8 orientations \times 2 phases (sine and cosine) = 80 filters or kernels, centered at each of the intersections of a uniform 10 \times 10 grid covering the whole image], resulting in two high-dimensional feature vectors (one for each image). Finally, Gabor similarity between the two images was calculated as the

Memory of personally familiar faces

Euclidean distance between the above feature vectors. We performed Gabor similarity analyses for faces in each condition and for each experiment reported here. Mean image similarity between target faces and different levels of facial morphs in Experiments 1 to 3 are shown in **Figures 9-11** respectively (the sign of all Gabor similarity values was reversed for illustration purpose).

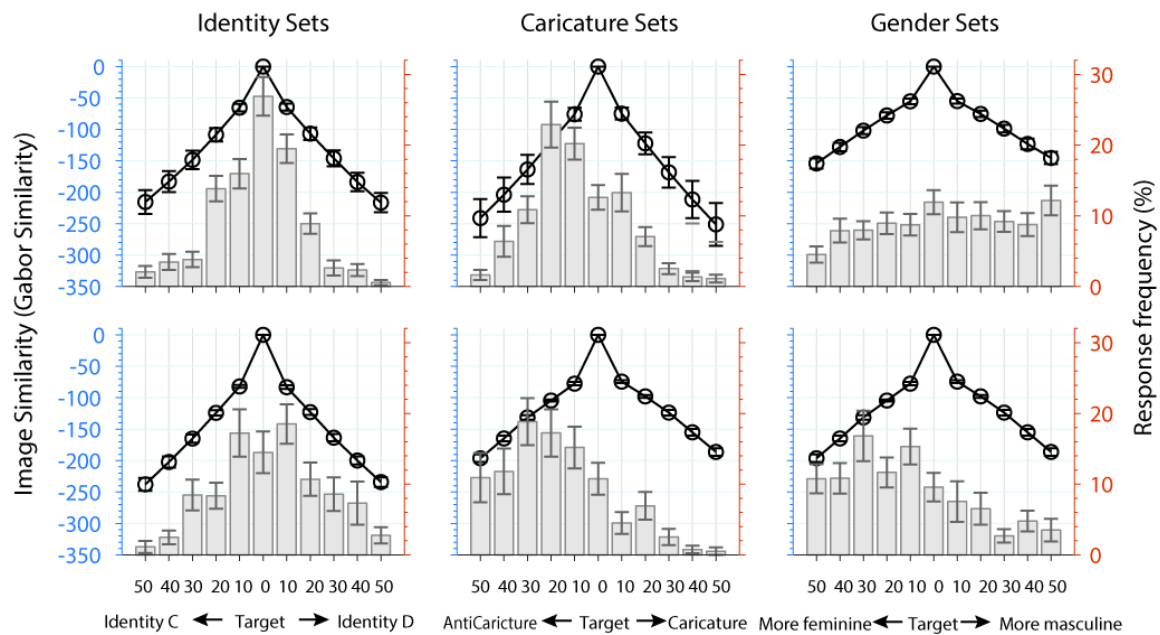


Figure 9. Gabor similarity between the face morphs and the target face used in Experiment 1. Black lines (left Y-axis) represent the image similarity and the grey bars (right Y-axis) represent participants' choice frequency. The upper row is based on data and stimuli from Experiment 1a (i.e., male faces), the lower row is based on Experiment 1b (i.e., female faces).

The analyses of image similarity revealed three main findings. First, at increasing morph level we found a reduced image-based similarity between the morph and the target face. The further away a morph level is from the target face, the less similar it is to the target face—in terms of physical measures. This is true for all of our identity, caricature, gender, and race manipulations (see black lines in **Figures 9-11** that show peaked and symmetrical shape of image similarity to the target faces). The results also indicate that the “feeling” of image similarity by eyeballing does not necessarily correspond to its true image-based similarity.

Memory of personally familiar faces

For instance, the male faces in the identity sets in Figure 2 look more different from each other than in the caricature sets, however, the computed similarity between images of the caricature sets ($M \pm SEM = -148.38 \pm 15.44$) is in fact numerically lower than in the identity sets (-130.73 ± 15.44 ; $F(1,18) = 0.65$, $p = .43$).

Second, responses to identity sets, but not gender or race sets, mirrored image-based similarity between the targets and different levels of morphs. Response frequency to identity morphs decreases with decreasing image-based similarity with the target faces (Figures 9 and 11, left column). In contrast, responses to gender and race sets showed no such patterns anchored to the target faces, and participants' response distributions were either flat across morphs levels or biased toward morphs at one side of the target faces (Figures 9 and 11, right column).

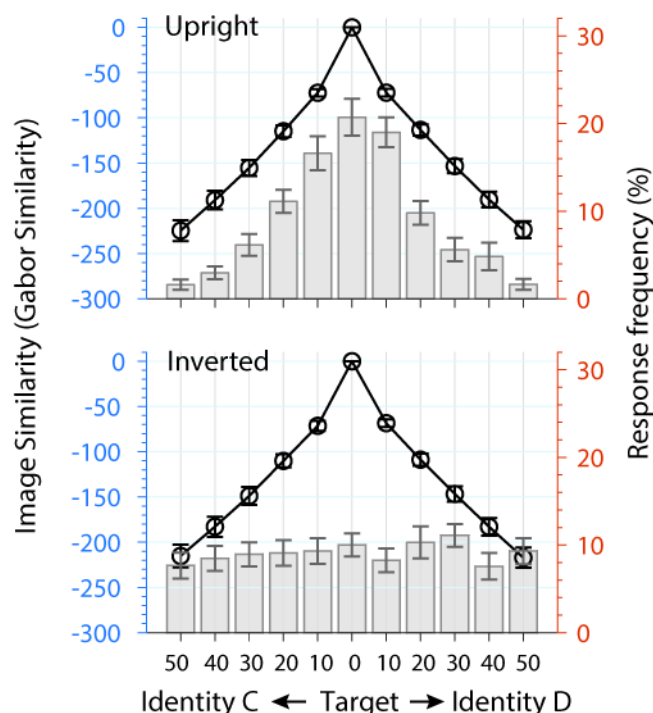


Figure 10. Gabor similarity between face morphs and target faces used in Experiments 1 and 2. The black lines (left Y-axis) represent the image similarity and the grey bars (right Y-axis) represent participants' choice frequency. Note that the inverted faces shared 12 out of 17 identities used as upright faces in Experiment 1.

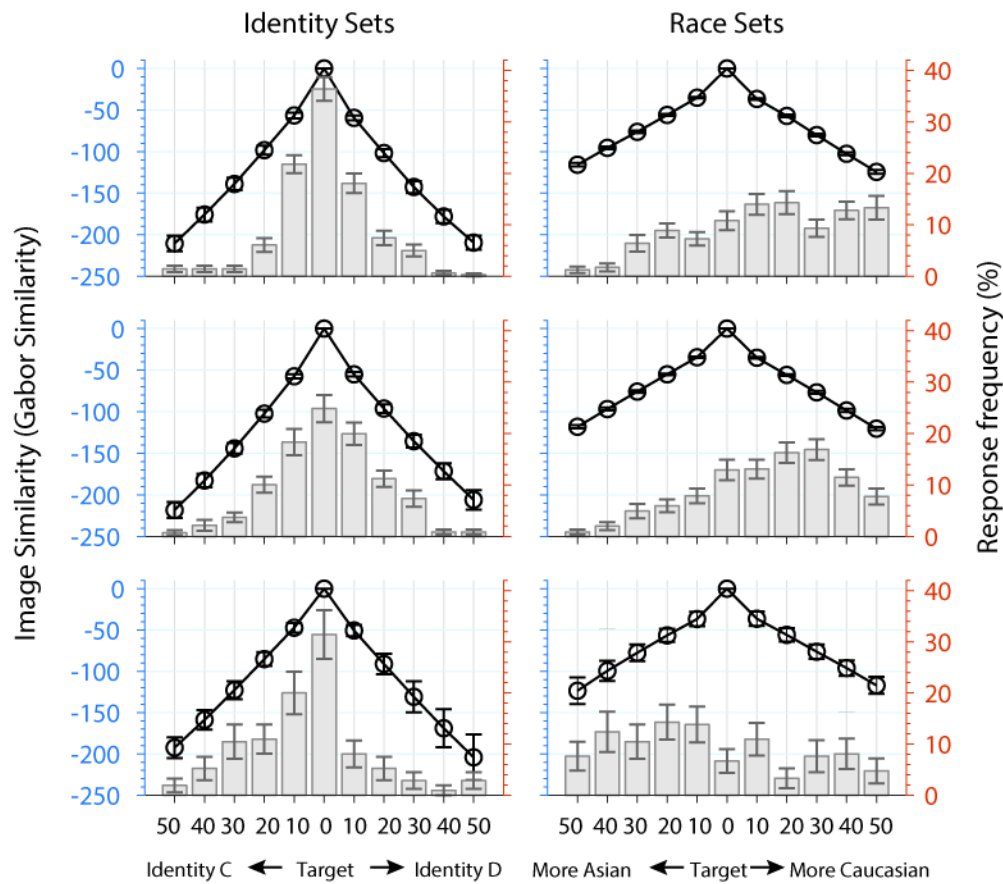


Figure 11. Gabor similarity between the face morphs and the target face used in Experiment 3. The black lines (left Y-axis) represent the image similarity and the grey bars (right Y-axis) represent participants' choice frequency. Top, middle, and bottom rows are based on the data and stimuli from Caucasian male faces, Caucasian female faces, and Asian Male faces, respectively.

Last but not the least, the overall image similarity is higher between gender/race morphs than between identity morphs (e.g., **Figure 11**, black lines in the left vs. in the right column), however, image similarity alone cannot account for the different patterns of responses to identity and gender/race manipulations. Even for the identity sets (**Figure 10**), the same level of image similarity between upright and inverted faces (137.36 ± 6.40 vs. 132.06 ± 7.62 , $F(1,27) = 0.28$, $p = .60$) can produce completely different patterns of response (interaction between orientation and morph levels: $F(10,1090) = 8.00$, $p < .001$). To directly test whether similarity between images underlies the different responses to identity and

Memory of personally familiar faces

gender/race sets, we performed a further test. To make idiosyncratic and categorical face manipulations have comparable image similarity to the target faces, we manually selected 5 morphs in each of the identity, gender, and race set in Experiment 1 (a and b) and Experiment 3 that displayed equivalent level of similarity to their target face, independent of their morph level. As shown in **Figure 12**, different patterns of responses to identity changes (black bars) and to gender/race changes (grey bars) persists—for both male and female faces—even though identity and gender/race morphs share almost the same levels of image-based similarity (as indicated by the overlapping black and dotted lines). For image similarity, a 2 (identity vs. gender/race) by 5 (morph level) ANOVAs revealed neither significant main effect of face manipulation nor significant interaction (all $F_s \leq 2.36$, all $p_s \geq .07$). For behavioral responses, we found significant interactions between face manipulation and morph level across all four sets of data shown in Figure 12 (all $F_s \geq 3.62$, all $p_s \leq .007$).

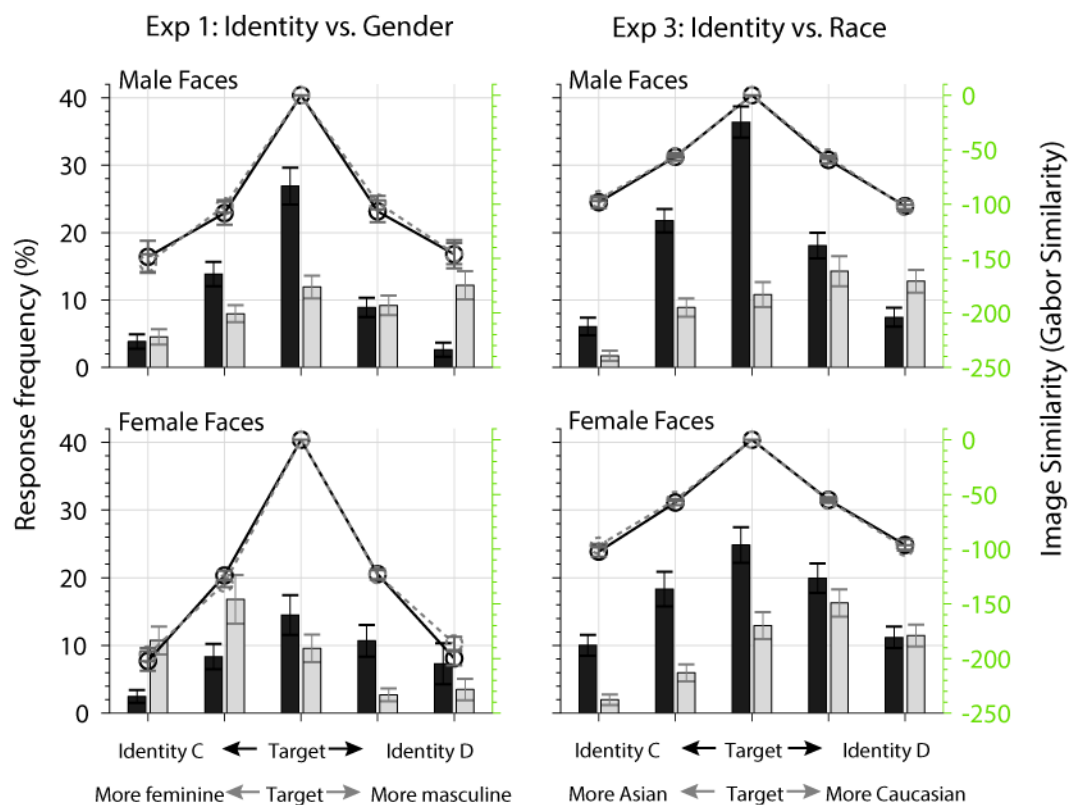


Figure 12. Response frequency of identity and gender/race manipulation when Gabor similarity between the face morphs and the target face is controlled. Black and grey bars (left Y-axis) represent participants' choice frequency for identity and gender/race sets respectively. Solid and dashed lines (right Y-axis) represent the image similarity for the selected identity and gender/race morphs respectively. Upper and lower rows shows data for male and female target faces respectively. For Experiment 1 (left column), the chosen identity morph levels were 20% and 30% toward each side for male faces and 20% and 40% for female faces. The corresponding gender morph levels were 30% and 50% toward each side for both male and female faces. For Experiment 3 (right column), the chosen identity morphs were 10% and 20% and the chosen race morphs were 20% and 40% toward each side for both male and female faces.

General discussion

In the present study, we investigated and directly contrasted how well people represent idiosyncratic and categorical facial information of personally familiar persons in memory. Our study revealed four major findings. First, participants remembered very precisely the idiosyncratic facial features of the people they know well (i.e., colleagues). They could, based exclusively on their memory, identify a veridical familiar faces among very similar distracters which were generated by morphing the original face with the face of another person. Second, participants judged the anti-caricatures of familiar faces (i.e., toward the average face and away from the original face) more often as the true target face than the targets or their caricatures. This finding concurs with previous studies using caricature and anti-caricature morphs created with photographs, although such a bias may partially due to the fact that our caricatures had a more blotchy appearance than the target faces and the anti-caricatures. Third, in contrast to relative precise memory of idiosyncratic information, participants showed less precise memory of categorical facial information such as gender and race. While participants agreed on which face in a caricature or identity set represented the familiar target person the best, as indicated by a robust peaked choice distribution, response distributions to gender and race manipulations were flatter and thus without a

Memory of personally familiar faces

clearly preferred representation. Finally, the categorical information of familiar faces, like their gender and race, was enhanced in memory (e.g. participants preferred morphs further away from the opposite category than the targets in the gender and race sets). When gender or race information of a familiar face was altered in our test sets, participants tended to choose more feminine/masculine morphs than the female/male targets, and similarly they chose more Caucasian-/Asian-looking morphs than the original Caucasian/ Asian targets. These results not only paint a picture of how very familiar faces are stored in memory, but also uncover critical differences between the representation of idiosyncratic and categorical facial information.

When the idiosyncratic information of familiar faces was parametrically manipulated, participants showed high sensitivity to such change. Compared to the original faces, participants' response frequency often starts to drop when the morphs contain only 10% idiosyncratic information from a different face (i.e., in identity sets), and further drops to chance-level when the morphs consist of 20% idiosyncratic information from a novel person (Figures 4 and 8). Participants also showed similar sensitivity to the change of identity strength when faces varied along the caricature and anti-caricature axis (i.e., in caricature sets). However, the memorized identity strength for familiar faces appears to be reduced, as the peak response frequency occurs at the morphs of 20% or 30% toward the average face. This contradicts the early hypothesis that diagnostic features of familiar faces are exaggerated in memory (e.g., Benson & Perrett, 1994; Gillian Rhodes et al., 1987, 1997)

Neurophysiological studies have suggested that the face cells in the primate brain are tuned to identity strength – the distance between a specific identity and the norm or average of faces in a multidimensional face space (L. Chang & Tsao, 2017; Giese & Leopold, 2005; Leopold et al., 2006; Loffler, Yourganov, Wilkinson, & Wilson, 2005). For instance, Leopold

Memory of personally familiar faces

and colleagues (Giese & Leopold, 2005; Leopold et al., 2006) have recorded the firing of face cells in the anterior inferotemporal cortex when macaque monkeys viewed face morphs that were very similar to those used in the current study (e.g., similar morphing manipulations of identity and caricature with the same face database and morphable model). They found that face cells respond monotonically to gradual changes along the caricature axis. The further away a face morph is from the average face, the higher is the elicited firing rate of the responding face cells (see Loffler et al., 2005, for analogous findings in a human fMRI study). Recently, Chang and Tsao (Chang & Tsao, 2017) showed that face cells respond monotonically along specific face axes (e.g., inter-eye distance) across the average face, demonstrating a ramp-like tuning (see also Loffler et al., 2005). Despite the difference in the shape of face cells' tuning function, both studies agree, at least for (anti-)caricaturing faces, that face cells show the lowest level of neural responses to the average face (i.e., zero identity strength) and the highest neural activity to the strongest caricatures (i.e., highest identity strength).

However, our participants here consistently considered the strongest face caricatures (i.e., 50% exaggeration of face identity strength) as displaying one of the lowest likeness to the original face identity. Similar behavioral responses have also been observed in previous studies (e.g., Allen et al., 2009; Benson & Perrett, 1991, 1994; Frowd et al., 2007). These results underscore the gap between the neurophysiological coding of faces and the behaviour of face processing. How a monotonic neural coding of face identity strength may underlie the peaked choice at anti-caricature levels remains to be elucidated. We speculate that at least two factors may underlie such a gap between behavioural responses to identity manipulation and the neurophysiological coding system for face identity. First, the discrepancy might result from the different tasks used. The task for the monkeys during

Memory of personally familiar faces

neurophysiological recording was not to find a face representing the learned identity the best; they performed either an identification task or no task at all. Second, high familiarity with faces might change the tuning of face cells to the identity variation of those faces. The tuning function of face cells was based on monkeys' response to either unfamiliar faces (e.g., Chang & Tsao, 2017) or visually familiar faces via training (e.g., Leopold et al., 2006), whereas the peaked choices in human behaviour is observed with personally familiar faces.

Our results with gender and race variations raise a further question about the aforementioned neural coding of faces: how is gender, race, and other categorical information of faces coded in the brain? Freeman, Rule, Adams, & Ambady (2010) have found that neural populations in the lateral fusiform gyrus show linearly increased neural activity to faces varying between androgynous to more gendered (either more male or more female looking), which is similar to Leopold et al.'s (2006) and Chang & Tsao's (2017) neurophysiological findings for identity manipulations. However, our results are not readily predicted by such monotonic tuning to the gender (or race) of faces. Our participants chose neither the original gender/race nor the most gendered face morphs most frequently. Instead, the coding of gender and race of faces seems to be slightly enhanced in memory. A male (female) familiar face is represented as more masculine (feminine) and a Caucasian (Asian) face is represented as more Caucasian (Asian)-looking. Distortions in subjective perception of facial gender have been reported for unfamiliar faces (e.g., Campanella, Chrysochoos, & Bruyer, 2001; Freeman et al., 2010). Here we show surprisingly that even for very familiar faces, people do not know exactly how "male/female" or how "Caucasian/Asian" these faces are. Further studies are required to address how a linear sensitivity to objective gender information may produce such a biased and enhanced gender/race representation for familiar faces (cf. Freeman et al., 2010).

Memory of personally familiar faces

The different patterns of response to identity/caricatures sets and to gender/race sets indicate that different characteristics of familiar faces are represented differently in memory. The results for the caricatures and the identity morphs both show a clearly peaked distribution, indicating a strong consensus among participants about the crucial idiosyncratic face features denoting a unique familiar identity in memory. The results obtained with the gender and race manipulations are often flatter without robust peak responses, suggesting that representation of categorical facial information is less precise than idiosyncratic information. The consistent bias of choosing faces of exaggerated categorical information suggests that categorical representation of familiar faces is not veridical. It is also worth noting that we used the term “gender” in this study, as we created morphed faces that were mixes of both sexes, whereas the targets were derived from male and female faces belonging to biological categories. In contrast, races are considered to be a cultural categories (Cosmides, Tooby, & Krurzban, 2003) and mixes of races in any amount occur. The similar results for gender (biological categories) and race categories (cultural categories) suggest that the bias of gender/race representation is not due to some top-down influence of culture or expertise.

Do different patterns of responses to idiosyncratic and categorical changes of faces result from different levels of image similarity in the stimuli sets? If the identity distracters differ more from the original faces in terms of image similarity than the gender or race distracters, it would be easier to identify the original face among the identity distracters than among the gender/race distracters, resulting in a peaked response for the former and less so for the latter. Our image similarity analyses indicate that this is not the case. Upright and inverted image sets have the same physical similarity but produce completely different patterns of responses (Experiment 2, see also **Figure 10**). Similarly, equivalent image

Memory of personally familiar faces

similarity between the race sets yielded opposite bias to Caucasian and Asian target faces (Experiment 3, **Figure 11** right column). If participants formed accurate representation of gender or race of faces and their responses were based only on the image differences between target and morphs, we would expect no bias toward one race or gender subcategory over the other. Thus, image similarity does not decisively drive participants' responses. Moreover, when image similarity between identity and gender/race morphs is set to be identical, the patterns of responses to identity and gender/race sets remain clearly different (**Figure 12**). Therefore, although gender/race morphs have an overall smaller variation in terms of image similarity than the identity morphs, such a difference cannot account for the different patterns of responses observed here.

Why do people respond differently to the idiosyncratic and categorical changes applied to very familiar faces? One possibility is that identity is a crucial facial characteristic that differs from race, gender, and other categorical information in our social life. We encounter a huge number of identities that belong to few different races and only two sexes. While discriminating between different identities has significant social importance, people usually categorize faces as either male or female, or Asian or Caucasian and rarely assess the intensity of those characteristics. Consequently, we have developed expertise for fine-grained sensitivity to identity changes but coarsely defined boundaries for categorical facial changes (e.g., Dobs et al., 2019). Remembering detailed levels of categorical information (e.g., specific level of masculinity associated with a familiar face) appears to be unnecessary and comes at a cost of memory storage. Discarding categorical details will make us less sensitive to the categorical than idiosyncratic changes. However, it also minimizes the processing and storage of visual information about familiar faces, resulting in an earlier and faster

Memory of personally familiar faces

categorization processing, which may facilitate more important identification process (e.g., Baudouin & Tiberghien, 2002; Dobs et al., 2019; Zhao & Hayward, 2013).

Another possibility is that the processes underlying participants' decision are more tuned to identity variations than variations along other dimensions of faces (e.g., gender or race). Faces that are equally distant from a target face in image similarity are not always perceived as equally distant from the target face in terms of identity similarity (e.g., Jacques & Rossion, 2006; Rotshtein, Henson, Treves, Driver, & Dolan, 2005). Similarly, Chang and Tsao (2017) showed that face cells in the primate brain are tuned to face variations along single axes of face space, but are blind to variations orthogonal to the tuned axis. It is therefore possible that participants were focusing on the perceived identity, not the gender or race appearance when choosing the face that best resembles a familiar person, even though we varied the target face along different axes in face space than identity (e.g., gender or race). We speculate that our participants might have been processing the gender and race manipulations in the same way as they process variations of facial expressions which do not induce a change of the perceived identity. Under this scenario, the deformations of faces in gender and race morphs, within a certain grade, are perceived as natural variations of the target face, making it hard to decide which of them is the original face.

Our results may help addressing one fundamental question faced by all models of face recognition-- how idiosyncratic and categorical face information is stored in memory and whether this information contributes to face recognition and identification equally. This is of interests for both biological and artificial face recognition systems (Kramer, Young, Day, & Burton, 2017; O'Toole, Castillo, Parde, Hill, & Chellappa, 2018). Earlier models of face recognition often assume that categorizing faces is an easier task and/or involves separate processes from those of face identification (Bruce & Young, 1986; Haxby, Hoffman, &

Memory of personally familiar faces

Gobbini, 2000; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Ullman, 1996). Recent studies promote a view that these invariant aspects of face information are encoded and processed in an integrative way (Armann & Bühlhoff, 2012; Bühlhoff & Newell, 2004; Dahl, Rasch, Bühlhoff, & Chen, 2016; Dobs et al., 2019; Kramer et al., 2017; Zhao & Hayward, 2013). Such integrative processing of race, gender, and identity is also consistent with the notion of “holistic” face representation, that is, we may form an appearance (norm) based face representation that contains both types of information as a whole (Rossion, 2013; see also (Zhao, Bühlhoff, & Bühlhoff, 2016; Zhao, Bühlhoff, & Bühlhoff, 2016). These various aspects of identity-related and categorical information may be processed together (Calder & Young, 2005) or may represent different levels of generalization (Riesenhuber & Poggio, 2000). Here we show that these invariant aspects of facial information are not weighted equally in the recognition process; even for very familiar faces, we are more sensitive to the idiosyncratic than categorical change of faces. Our ability to categorize faces rapidly comes at a cost: losing fine-detailed representation of category and being prone to representational bias. Our results also suggest that different types of facial information are not equally accessible. Otherwise, we should not observe different pattern of responses to identity and gender/race variations.

Conclusion

The present study demonstrates that the representation of idiosyncratic information (i.e., identity-specific features) of familiar faces is more precise than the representation of categorical face information (i.e., gender/race). While people often remember precisely what facial features make a familiar face different from other faces, they have no clear memory about the non-identity-specific, categorical face information, even for very familiar colleagues. Such a difference might be driven by both the social significance of person

Memory of personally familiar faces

individualization and an adapted neural coding system that is tuned to idiosyncratic face features. What we remember about a familiar face is what is needed in our real life, not more. Our face processing system minimizes the storage of unnecessary categorical information and is less tuned to identity-irrelevant facial changes (e.g., expressions, illumination, viewpoints, and natural or artificial face deformations) to achieve rapid and precise identification of familiar individuals—persons that are of the most social significance to us.

Acknowledgments

We wish to thank Karin Bierig for her help with all experiments and all our colleagues for agreeing to the use of their faces and for their participation in the experiments. This research was supported by the Max Planck Society.

References

- Allen, H., Brady, N., & Tredoux, C. (2009). Perception of “best likeness” to highly familiar faces of self and friend. *Perception, 38*(12), 1821–1830. <https://doi.org/10.1068/p6424>
- Armann, R., & Bühlhoff, I. (2009). Gaze behavior in face comparison: The roles of sex, task, and symmetry. *Attention, Perception, & Psychophysics, 71*(5), 1107–1126. <https://doi.org/10.3758/APP>
- Armann, R., & Bühlhoff, I. (2012). Male and female faces are only perceived categorically when linked to familiar identities—and when in doubt, he is a male. *Vision Research, 63*, 69–80. <https://doi.org/10.1016/j.visres.2012.05.005>
- Armann, R., Jeffery, L., Calder, A. J., & Rhodes, G. (2011). Race-specific norms for coding face identity and a functional role for norms. *Journal of Vision, 11*(13), 1–14.
- Armann, R., Jenkins, R., & Burton, A. M. (2016). A familiarity disadvantage for remembering

Memory of personally familiar faces

specific images of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 42(4), 571–580. <https://doi.org/10.1037/xhp0000174>

Bahrick, H. P., Bahrick, P. O., & Wittlinger, R. P. (1975). Fifty Years of Memory for Names and Faces: A Cross-Sectional Approach. *Journal of Experimental Psychology: General*, 104(1), 54–75. <https://doi.org/10.1037/0096-3445.104.1.54>

Baudouin, J.-Y., & Tiberghien, G. (2002). Gender is a dimension of face recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(2), 362–365. <https://doi.org/10.1037//0278-7393.28.2.362>

Benson, P. J., & Perrett, D. I. (1991). Perception and recognition of photographic quality facial caricatures: Implications for the recognition of natural images. *European Journal of Cognitive Psychology*, 3(1), 103–105.

Benson, P. J., & Perrett, D. I. (1994). Visual processing of facial distinctiveness. *Perception*, 23(1), 75–93. <https://doi.org/10.1068/p230075>

Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99* (pp. 187–194). New York, New York, USA: ACM Press. <https://doi.org/10.1145/311535.311556>

Bonner, L., Burton, A. M., & Bruce, V. (2003). Getting to know you: How we learn new faces. *Visual Cognition*, 10(5), 527–536.

Bruce, V. (1982). Changing faces: Visual and non-visual coding processes in face recognition. *British Journal of Psychology*, 73(1), 105–116.

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), 339–360. <https://doi.org/10.1037//1076->

Memory of personally familiar faces

898X.5.4.339

- Bruce, V., & Young, A. W. (1986). Understanding face recognition. *British Journal of Psychology*, 77(3), 305–327. <https://doi.org/10.1111/j.2044-8295.1986.tb02199.x>
- Bülthoff, I., & Newell, F. N. (2004). Categorical perception of sex occurs in familiar but not unfamiliar faces. *Visual Cognition*, 11(7), 823–855.
<https://doi.org/10.1080/13506280444000012>
- Burton, A. M. (2013). Why has research in face recognition progressed so slowly? The importance of variability. *The Quarterly Journal of Experimental Psychology*, 66(8), 1467–1485. <https://doi.org/10.1080/17470218.2013.800125>
- Burton, A. M., Schweinberger, S. R., Jenkins, R., & Kaufmann, J. M. (2015). Arguments Against a Configural Processing Account of Familiar Face Recognition. *Perspectives on Psychological Science*, 10(4), 482–496. <https://doi.org/10.1177/1745691615583129>
- Burton, A. M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face Recognition in Poor-Quality Video: Evidence From Security Surveillance. *Psychological Science*, 10(3), 243–248.
<https://doi.org/10.1111/1467-9280.00144>
- Calder, A. J., & Young, A. W. (2005). Understanding the recognition of facial identity and facial expression. *Nature Reviews Neuroscience*, 6(8), 641–651.
<https://doi.org/10.1038/nrn1724>
- Campanella, S., Chrysochoos, A., & Bruyer, R. (2001). Categorical perception of facial gender information: Behavioural evidence and the face-space metaphor. *Visual Cognition*, 8(2), 237–262. <https://doi.org/10.1080/13506280042000072>
- Chang, L., & Tsao, D. Y. (2017). The Code for Facial Identity in the Primate Brain. *Cell*, 169(6), 1013–1028.e14. <https://doi.org/10.1016/j.cell.2017.05.011>
- Chang, P. P. W., Levine, S. C., & Benson, P. J. (2002). Children's recognition of caricatures.

Memory of personally familiar faces

Developmental Psychology, 38(6), 1038–1051. <https://doi.org/10.1037//0012-1649.38.6.1038>

Cosmides, L., Tooby, J., & Krurzban, R. (2003). Perceptions of Race. *TRENDS in Cognitive Sciences*, 7(4), 173–179. [https://doi.org/10.1016/S1364-6613\(03\)00057-3](https://doi.org/10.1016/S1364-6613(03)00057-3)

Dahl, C. D., Rasch, M. J., Bülthoff, I., & Chen, C.-C. (2016). Integration or separation in the processing of facial properties - a computational view. *Scientific Reports*, 6(February), 20247. <https://doi.org/10.1038/srep20247>

Dehon, H., Bredart, S., & Brédart, S. (2001). An “other-race” effect in age estimation from faces. *Perception*, 30(9), 1107–1113. <https://doi.org/10.1068/p3122>

Dobs, K., Bülthoff, I., Breidt, M., Vuong, Q. C., Curio, C., & Schultz, J. (2014). Quantifying human sensitivity to spatio-temporal information in dynamic faces. *Vision Research*, 100, 78–87. <https://doi.org/10.1016/j.visres.2014.04.009>

Dobs, K., Bülthoff, I., & Schultz, J. (2016). Identity information in facial motion varies with the type of facial movement. *Scientific Reports*. <https://doi.org/10.1038/srep34301>

Dobs, K., Isik, L., Pantazis, D., & Kanwisher, N. (2019). How face perception unfolds over time. *Nature Communications*, 10(1), 1–10. <https://doi.org/10.1038/s41467-019-09239-1>

Dunn, J. D., Ritchie, K. L., Kemp, R. I., & White, D. (2019). Familiarity does not inhibit image-specific encoding of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 45(7), 841–854. <https://doi.org/10.1037/xhp0000625>

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. <https://doi.org/10.1068/p080431>

Freeman, J. B., Rule, N. O., Adams, R. B., & Ambady, N. (2010). The neural basis of categorical face perception: graded representations of face gender in fusiform and orbitofrontal cortices. *Cerebral Cortex (New York, N.Y. : 1991)*, 20(6), 1314–1322.

Memory of personally familiar faces

<https://doi.org/10.1093/cercor/bhp195>

Freire, A., Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception, 29*(2), 159–170.

<https://doi.org/10.1068/p3012>

Frowd, C., Bruce, V., Ross, D. A., McIntyre, A., & Hancock, P. J. B. (2007). An application of caricature: How to improve the recognition of facial composites. *Visual Cognition, 15*(8), 954–984. <https://doi.org/10.1080/13506280601058951>

Giese, M. A., & Leopold, D. A. (2005). Physiologically inspired neural model for the encoding of face spaces. *Neurocomputing, 65–66*, 93–101.

<https://doi.org/10.1016/j.neucom.2004.10.060>

Gobbini, M. I., & Haxby, J. V. (2006). Neural response to the visual familiarity of faces. *Brain Research Bulletin, 71*(1–3), 76–82. <https://doi.org/10.1016/j.brainresbull.2006.08.003>

Gobbini, M. I., & Haxby, J. V. (2007). Neural systems for recognition of familiar faces.

Neuropsychologia, 45(1), 32–41. <https://doi.org/10.1016/j.neuropsychologia.2006.04.015>

Hancock, P. J. B., Bruce, V., & Burton, A. M. (2010). Recognition of unfamiliar faces. *Trends in Cognitive Sciences, 4*(9), 330–337. [https://doi.org/10.1016/S1364-6613\(00\)01519-9](https://doi.org/10.1016/S1364-6613(00)01519-9)

Harmon, L. D. (1973). The recognition of faces. *Scientific American, 229*(5), 70–83.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*(6), 223–233.

[https://doi.org/10.1016/S1364-6613\(00\)01482-0](https://doi.org/10.1016/S1364-6613(00)01482-0)

Hochberg, J., & Galper, R. R. (1967). Recognition of faces: An explanatory Study.

Psychological Science, 9, 619–620.

Jacques, C., & Rossion, B. (2006). The speed of individual face categorization. *Psychological Science, 17*(6), 485–492. <https://doi.org/10.1111/j.1467-9280.2006.01733.x>

Memory of personally familiar faces

Jaquet, E., Rhodes, G., & Hayward, W. G. (2008). Race-contingent aftereffects suggest distinct perceptual norms for different race faces. *Visual Cognition*, 16(6), 734–753.

<https://doi.org/10.1080/13506280701350647>

Jiang, F., Blanz, V., & O’Toole, A. J. (2007). The role of familiarity in three-dimensional view-transferability of face identity adaptation. *Vision Research*, 47(4), 525–531.

<https://doi.org/10.1016/J.VISRES.2006.10.012>

Johnston, R. A., & Edmonds, A. J. (2009). Familiar and unfamiliar face recognition: A review.

Memory, 17(5), 577–596. <https://doi.org/10.1080/09658210902976969>

Kaufmann, J. M., & Schweinberger, S. R. (2008). Distortions in the brain? ERP effects of caricaturing familiar and unfamiliar faces. *Brain Research*, 1228, 177–188.

<https://doi.org/DOI:10.1016/j.brainres.2008.06.092>

Kramer, R. S. S., Young, A. W., Day, M. G., & Burton, A. M. (2017). Robust social

categorization emerges from learning the identities of very few faces. *Psychological*

Review, 124(2), 115–129. <https://doi.org/10.1037/rev0000048>

Lee, K. J., & Perrett, D. I. (2000). Manipulation of colour and shape information and its consequence upon recognition and best-likeness judgments. *Perception*, 29(11), 1291–

1312. <https://doi.org/10.1068/p2792>

Leopold, D. A., Bondar, I. V., & Giese, M. a. (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature*, 442(7102), 572–575.

<https://doi.org/10.1038/nature04951>

Leopold, D. A., O’Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4(1), 89–94.

<https://doi.org/10.1038/82947>

Leveroni, C. L., Seidenberg, M., Mayer, A. R., Mead, L. A., Binder, J. R., & Rao, S. M. (2000).

Memory of personally familiar faces

- Neural Systems Underlying the Recognition of Familiar and Newly Learned Faces. *The Journal of Neuroscience*, 20(2), 878–886. Retrieved from <http://www.jneurosci.org/cgi/content/abstract/20/2/878>
- Lewis, M. B., & Johnston, R. A. (1998). Understanding Caricatures of Faces. *The Quarterly Journal of Experimental Psychology*, 50A(2), 321–346. <https://doi.org/10.1080/713755758>
- Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience*, 8(10), 1386–1390. <https://doi.org/10.1038/nn1538>
- Megreya, A. M., & Burton, A. M. (2006). Unfamiliar faces are not faces: Evidence from a matching task. *Memory & Cognition*, 34(4), 865–876. <https://doi.org/10.3758/BF03193433>
- Natu, V. S., & O'Toole, A. J. (2011). The neural processing of familiar and unfamiliar faces: a review and synopsis. *British Journal of Psychology (London, England : 1953)*, 102(4), 726–747. <https://doi.org/10.1111/j.2044-8295.2011.02053.x>
- Natu, V. S., & O'Toole, A. J. (2015). Spatiotemporal changes in neural response patterns to faces varying in visual familiarity. *NeuroImage*, 108, 151–159. <https://doi.org/10.1016/j.neuroimage.2014.12.027>
- O'Toole, A. J., Castillo, C. D., Parde, C. J., Hill, M. Q., & Chellappa, R. (2018). Face Space Representations in Deep Convolutional Neural Networks. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2018.06.006>
- O'Toole, A. J., Vetter, T., Volz, H., & Salter, E. M. (1997). Three-dimensional caricatures of human heads: Distinctiveness and the perception of facial age. *Perception*, 26(6), 719–732. <https://doi.org/10.1068/p260719>
- Osborne, C. D., & Stevenage, S. V. (2008). Internal feature saliency as a marker of familiarity and configural processing. *Visual Cognition*, 16(1), 23–43.

Memory of personally familiar faces

<https://doi.org/10.1080/13506280701238073>

Phillips, P. J., Scruggs, W. T., O'Toole, A. J., Flynn, P. J., Bowyer, K. W., Schott, C. L., & Sharpe, M. (2007). *FRVT 2006 and ICE 2006 Large-Scale Results. NISTIR 7408 Report, March 2007.*

Platek, S. M., & Kemp, S. M. (2009). Is family special to the brain? An event-related fMRI study of familiar, familial, and self-face recognition. *Neuropsychologia*.

<https://doi.org/10.1016/j.neuropsychologia.2008.12.027>

Ramon, M., & Van Belle, G. (2016). Real-life experience with personally familiar faces enhances discrimination based on global information. *PeerJ*, 4, e1465.

<https://doi.org/10.7717/peerj.1465>

Ramon, M., Vizioli, L., Liu-Shuang, J., & Rossion, B. (2015). Neural microgenesis of personally familiar face recognition. *Proceedings of the National Academy of Sciences*, (AUGUST), 201414929. <https://doi.org/10.1073/pnas.1414929112>

Rhodes, G. (1996). *Superportraits: Caricatures and recognition*. Hove, U.K.: Psychology Press.

Rhodes, G., Jaquet, E., Jeffery, L., Evangelista, E., Keane, J., & Calder, A. J. (2011). Sex-specific norms code face identity. *Journal of Vision*, 11(1), 1–11.

<https://doi.org/10.1167/11.1.1>

Rhodes, Gillian, Brennan, S. E., & Carey, S. (1987). Identification and ratings of caricatures: implications for mental representations of faces. *Cognitive Psychology*, 19, 473–497.

[https://doi.org/10.1016/0010-0285\(87\)90016-8](https://doi.org/10.1016/0010-0285(87)90016-8)

Rhodes, Gillian, Byatt, G., Tremewan, T., & Kennedy, A. (1997). Face distinctiveness and the power of caricatures. *Perception*, 26(2), 207–223.

Rhodes, Gillian, Hickford, C., & Jeffery, L. (2000). Sex-typicality and attractiveness - Are supermale and superfemale faces super-attractive? *British Journal of Psychology*, 91, 125–

140.

Rhodes, Gillian, & Tremewan, T. (1994). Understanding face recognition: Caricature effects, inversion, and the homogeneity problem. *Perception, 1*, 275–311.

Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience, 3*, 1199–1204.

Robbins, R., & McKone, E. (2003). Can holistic processing be learned for inverted faces? *Cognition*. [https://doi.org/10.1016/S0010-0277\(03\)00020-9](https://doi.org/10.1016/S0010-0277(03)00020-9)

Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cogn. Psychol., 8*, 382–439.

Rossion, B. (2002). Is sex categorization from faces really parallel to face recognition? *Visual Cognition, 9*(8), 1003–1020. <https://doi.org/10.1080/13506280143000485>

Rossion, B. (2013). The composite face illusion: A whole window into our understanding of holistic face perception. *Visual Cognition, 21*(2), 139–253.
<https://doi.org/10.1080/13506285.2013.772929>

Rossion, B., & Gauthier, I. (2002). How does the brain process upright and inverted faces? *Behavioral and Cognitive Neuroscience Reviews, 1*(1), 63–75.
<https://doi.org/10.1177/1534582302001001004>

Rotshtein, P., Henson, R. N. a, Treves, A., Driver, J., & Dolan, R. J. (2005). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nature Neuroscience, 8*(1), 107–113. <https://doi.org/10.1038/nn1370>

Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychon Bull Rev, 16*, 252–257.

Sandford, A., & Burton, A. M. (2014). Tolerance for distorted faces: Challenges to a configural processing account of familiar face recognition. *Cognition, 132*, 262–268.

Memory of personally familiar faces

- Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (doi: 10.1109/CVPR.2015.7298682), 815–823.
- Sinha, P., Balas, B., Ostrovsky, Y., & Russell, R. (2006). Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11), 1948–1961. <https://doi.org/10.1109/JPROC.2006.884093>
- Srismith, D., Zhao, M., & Bühlhoff, I. (2016). Precise Representation of Personally, but not Visually, Familiar Faces. *Perception*, 45 (ECPV A), 313.
- Taigman, Y., Yang, M., Ranzato, M., & Wolf, L. (2014). DeepFace: Closing the Gap to Human-Level Performance in Face Verification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (doi: 10.1109/CVPR.2014.220), 1701–1708.
- Taylor, M. J., Arsalidou, M., Bayless, S. J., Morris, D., Evans, J. W., & Barbeau, E. J. (2009). Neural correlates of personally familiar faces: parents, partner and own faces. *Human Brain Mapping*, 30(7), 2008–2020. <https://doi.org/10.1002/hbm.20646>
- Troje, N. F., & Bühlhoff, H. H. (1996). Face recognition under varying poses: The role of texture and shape. *Vision Research*, 36(12), 1761–1771. [https://doi.org/10.1016/0042-6989\(95\)00230-8](https://doi.org/10.1016/0042-6989(95)00230-8)
- Ullman, S. (1996). *High-Level Vision: Object Recognition and Visual Cognition*. Cambridge, MA, USA: MIT Press.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, 43(2), 161–204. <https://doi.org/10.1080/14640749108400966>
- Visconti Di Oleggio Castello, M., Halchenko, Y. O., Guntupalli, J. S., Gors, J. D., & Gobbini, M. I. (2017). The neural representation of personally familiar and unfamiliar faces in the

Memory of personally familiar faces

distributed system for face perception. *Scientific Reports*, 7(1), 1–14.

<https://doi.org/10.1038/s41598-017-12559-1>

White, D., Kemp, R., Jenkins, R., Matheson, M., & Burton, A. (2014). Passport Officers' Errors in Face Matching. *PLoS ONE*, 9(8), e103510.

Xu, X., & Biederman, I. (2010). Loci of the release from fMRI adaptation for changes in facial expression, identity, and viewpoint. *Journal of Vision*, 10(14), 36–36.

<https://doi.org/10.1167/10.14.36>

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81(1), 141–145. <https://doi.org/10.1037/h0027474>

Yip, A. W., & Sinha, P. (2002). Contribution of color to face recognition. *Perception*, 31(8), 995–1003. <https://doi.org/10.1068/p3376>

Young, A. W., & Burton, A. M. (2018). Are We Face Experts? *Trends in Cognitive Sciences*, 22(2), 100–110.

Yue, X., Biederman, I., Mangini, M. C., Malsburg, C. Von Der, & Amir, O. (2012). Predicting the psychophysical similarity of faces and non-face complex shapes by image-based measures. *Vision Research*, 55, 41–46. <https://doi.org/10.1016/j.visres.2011.12.012>

Zhao, M., Bülthoff, H. H., & Bülthoff, I. (2016). A shape-based account for holistic face processing. *Journal of Experimental Psychology: Learning Memory and Cognition*, 42(4). <https://doi.org/10.1037/xlm0000185>

Zhao, M., & Hayward, W. G. (2013). Integrative processing of invariant aspects of faces: Effect of gender and race processing on identity analysis. *Journal of Vision*, 13(1), 15–15. <https://doi.org/10.1167/13.1.15>

Zhao, M., Bülthoff, H. H., & Bülthoff, I. (2016). Beyond Faces and Expertise: Facelike Holistic Processing of Nonface Objects in the Absence of Expertise. *Psychological Science*, 27(2),

Memory of personally familiar faces

213–222. <https://doi.org/10.1177/0956797615617779>