Idiosyncratic, retinotopic bias in face identification modulated by familiarity

Abbreviated title: Retinotopic bias in face identification

Matteo Visconti di Oleggio Castello[1,*], Morgan Taylor[1], Patrick Cavanagh[1,2], M. Ida Gobbini[1,3,*]

[1]Department of Psychological & Brain Sciences,
Dartmouth College,
Hanover NH, 03755, USA
[2]Department of Psychology
Glendon College
Toronto ON, M4N 3M6 Canada
[3]Dipartimento di Medicina Specialistica, Diagnostica, e Sperimentale,
University of Bologna,
40100 Bologna, Italy

[*]corresponding authors:
M. Ida Gobbini, mariaida.gobbini@unibo.it
Matteo Visconti di Oleggio Castello, mvdoc.gr@dartmouth.edu
6207 Moore Hall
Dartmouth College
Hanover, NH 03755, USA

Number of pages:	47
Number of figures:	5
Number of tables:	3
Words (Abstract):	155
Words (Introduction):	831
Words (Discussion):	2,133

Conflict of Interest

The authors declare no competing financial interests.

Acknowledgments

We would like to thank the Martens Family Fund and Dartmouth College for their support. We would like to thank Carlo Cipolli for helpful discussion.

Abstract

The perception of gender and age of unfamiliar faces is reported to vary idiosyncratically across retinal locations such that, for example, the same androgynous face may appear to be male at one location but female at another. Here we test spatial heterogeneity for the recognition of the *identity* of personally familiar faces in human participants. We found idiosyncratic biases that were stable within participants and that varied more across locations for low as compared to high familiar faces. These data suggest that like face gender and age, face identity is processed, in part, by independent populations of neurons monitoring restricted spatial regions and that the recognition responses vary for the same face across these different locations. Moreover, repeated and varied social interactions appear to lead to adjustments of these independent face recognition neurons so that the same familiar face is eventually more likely to elicit the same recognition response across widely separated visual field locations.

Significance statement

In this work we tested spatial heterogeneity for the recognition of personally familiar faces. We found retinotopic biases that varied more across locations for low as compared to highly familiar faces. The retinotopic biases were idiosyncratic and stable within participants. Our data suggest that, like face gender and age, face identity is processed by independent populations of neurons monitoring restricted

spatial regions and that recognition may vary for the same face at these different locations. Unlike previous findings, our data show how the effect of learning modifies the representation of face identity in cortical areas with spatially restricted receptive fields. This new perspective has broader implications for understanding how learning optimizes visual processes for socially salient stimuli.

Introduction

We spend most of our days interacting with acquaintances, family and close friends. Because of these repeated and protracted interactions, the representation of personally familiar faces is rich and complex, as reflected by stronger and more widespread neural activation in the distributed face processing network, as compared to responses to unfamiliar faces (Gobbini and Haxby, 2007; Taylor et al., 2009; Gobbini, 2010; Natu and O'Toole, 2011; Bobes et al., 2013; Sugiura, 2014; Ramon and Gobbini, 2017; Visconti di Oleggio Castello et al., 2017a). Differences in representations are also reflected in faster detection and more robust recognition of familiar faces (Burton et al., 1999; Gobbini et al., 2013; Ramon et al., 2015; Visconti di Oleggio Castello and Gobbini, 2015; Visconti di Oleggio Castello et al., 2017b). Thus, despite the subjective feeling of expertise with faces in general (Diamond and Carey, 1986), our visual system seems to be optimized for the processing of familiar faces. The mechanisms underlying the prioritized processing of familiar faces are still a matter of investigation (Guntupalli and Gobbini, 2017; Ramon and Gobbini, 2017; Young and Burton, 2017).

The advantage for familiar faces could originate at different stages of the face processing system. The classic psychological model by Bruce and Young (1986) posits that recognition of familiar faces occurs when the structural encoding of a perceived face matches stored representations (Bruce and Young, 1986). In this model the stored representations of familiar faces consist of "an interlinked set of expression-independent structural codes for distinct head angles, with some codes reflecting the global configuration at each angle and others representing particular distinctive features" (Bruce and Young, 1986, p. 309). Behavioral evidence supports the hypothesis that local features are processed differentially for personally familiar faces. For example, in a study of perception of gaze direction and head angle, changes in eye gaze were detected around 100ms faster in familiar than in unfamiliar faces (Visconti di Oleggio Castello and Gobbini, 2015). In another study, the advantage for personally familiar faces was maintained after face inversion, a manipulation that is generally thought to reduce holistic processing in favor of local processing (Visconti di Oleggio Castello et al., 2017b).

Could local features be sufficient to initially drive a differential response to personally familiar faces? In a study measuring saccadic reaction time, correct and reliable saccades to familiar faces were recorded as fast as 180 ms when unfamiliar faces were distractors (Visconti di Oleggio Castello and Gobbini, 2015). In an EEG study using multivariate analyses, significant decoding of familiarity could be detected at around 140 ms from stimulus onset (Barragan-Jason et al., 2015). At such short latencies it is unlikely that a viewpoint-invariant representation of an individual face's identity

98 drives these differential responses. To account for facilitated, rapid detection of
99 familiarity, we have previously hypothesized that personally familiar faces may be
100 recognized quickly based on diagnostic, idiosyncratic features, which become highly
101 learned through extensive personal interactions (Visconti di Oleggio Castello and
102 Gobbini, 2015; Visconti di Oleggio Castello et al., 2017b). Detection of these features
103 may occur early in the face-processing system, allowing an initial, fast differential
104 processing for personally familiar faces.

105 Processes occurring at early stages of the visual system can show idiosyncratic
106 retinotopic biases (Greenwood et al., 2017). Afraz et al. (2010) reported retinotopic
107 biases for perceiving face gender and age that varied depending on stimulus location
108 in the visual field and were specific to each subject. These results suggest that
109 diagnostic facial features for gender and age are encoded in visual areas with limited
110 position invariance. Neuroimaging studies have shown that face-processing areas
111 such as OFA, pFus, and mFus have spatially restricted population receptive fields that
112 could result in retinotopic differences (Kay et al., 2015; Silson et al., 2016; Grill-Spector
113 et al., 2017b). Here we hypothesized that detectors of diagnostic visual features that
114 play a role in identification of familiar faces may also show idiosyncratic retinotopic
115 biases and that these biases may be tuned by repeated interactions with personally
116 familiar faces. Such biases may affect recognition of the identities presented in
117 different parts of the visual field and may be modulated by the familiarity of those
118 identities.

We tested this hypothesis by presenting participants with morphed stimuli of personally familiar individuals that were briefly shown at different retinal locations. In two separate experiments we found that participants showed idiosyncratic biases for specific identities in different visual field locations, and these biases were stable on retesting after weeks. Importantly, the range of the retinal biases was inversely correlated with the reported familiarity of each target identity, suggesting that prolonged personal interactions with the target individuals reduced retinal biases. These findings provide additional support for the hypothesis that asymmetries in the processing of personally familiar faces can arise at stages of the face-processing system where there is reduced position invariance and where local features are being processed, such as in OFA or perhaps even earlier. Our results show that prolonged, personal interactions can modify the neural representation of faces at this early level of processing.

Materials and Methods

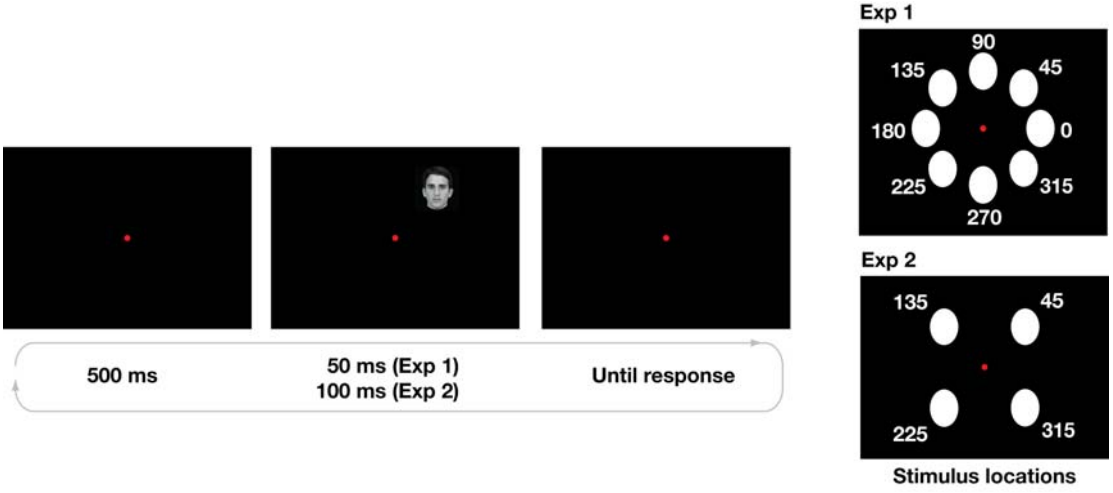

Figure 1. Experimental paradigm. The left panel shows an example of the

experimental paradigm, while the right panel shows the locations used in Experiment 1 (eight locations, top panel) and in Experiment 2 (four locations, bottom panel).

Stimuli

Pictures of the faces of individuals who were personally familiar to the participants (graduate students in the same department) were taken in a photo studio room with the same lighting condition and the same camera. Images of two individuals were used for Experiment 1, and images of three individuals were used for Experiment 2. All individuals portrayed in the stimuli signed written informed consent for the use of their pictures for research and in publications.

The images were converted to grayscale, resized and centered so that the eyes were aligned in the same position for the three identities, and the background was manually removed. These operations were performed using ImageMagick and Adobe Photoshop CS4. The resulting images were matched in luminance (average pixel intensity) using the SHINE toolbox (function *lumMatch*) (Willenbockel et al., 2010) after applying an oval mask, so that only pixels belonging to the face were modified. The luminance-matched images were then used to create morph continua (between two identities in Experiment 1, see Figure 2; and among three identities in Experiment 2, see Figure 3) using Abrosoft Fantamorph (v. 5.4.7) with seven percentages of morphing: 0, 17, 33, 50, 67, 83, 100 (see Figures 2, 3).

Experiment 1

Paradigm

The experimental paradigm was similar to that by Afraz et al., (2010). In every trial participants would see a briefly flashed image in one of eight locations at the periphery of their visual field (see Figure 1). Each image was shown for 50 ms at a distance of 7˚ of visual angle from the fixation point, and subtended approximately 4˚ x 4˚ of visual angle. The images could appear in one of eight locations evenly spaced by 45 angular degrees around fixation. For Experiment 1, only the morph *ab* was used (see Figure 1). Participants were required to maintain fixation on a central red dot subtending approximately 1˚ of visual angle.

After the image disappeared, participants reported which identity they saw using the left (identity *a*) and right (identity *b*) arrow keys. There was no time limit for responding, and participants were asked to be as accurate as possible. After responding, participants had to press the spacebar key to continue to the next trial.

Participants performed five blocks containing 112 trials each, for a total of 560 trials. In each block all the images appeared twice for every angular location (8 angular locations x 7 morph percentages x 2 = 112). This provided ten data points for each percentage morphing at each location, for a total of 70 trials at each angular location.

Before the experimental session participants were shown the identities used in the experiment (corresponding to 0% and 100% morphing, see Figure 2), and practiced

the task with 20 trials. These data were discarded from the analyses. Participants performed two identical experimental sessions at least four weeks apart.

Participants sat at a distance of approximately 50 cm from the screen, with their chin positioned on a chin-rest. The experiment was run using Psychtoolbox (Kleiner et al., 2007) (version 3.0.12) in MATLAB (R2014b). The screen operated at a resolution of 1920x1200 and a 60Hz refresh rate.

Subjects

We recruited six subjects for this experiment (three males, including one of the authors, MVdOC). The sample size for Experiment 1 was not determined by formal estimates of power, and was limited by the availability of participants familiar with the stimulus identities. After the first experimental session, two participants (one male, one female) were at chance level in the task, thus only data from four subjects (two males, mean age 27.50 ± 2.08 SD) were used for the final analyses.

All subjects had normal or corrected-to-normal vision, and provided written informed consent to participate in the experiment. The study was approved by the Dartmouth College Committee for the Protection of Human Subjects.

Experiment 2

Paradigm

Experiment 2 differed from Experiment 1 in the following parameters (see Figures 1, 3): 1. three morph continua (*ab*, *ac*, *bc*) instead of one; 2. images appeared in four locations (45˚, 135˚, 225˚, 315˚) instead of eight; 3. images were shown for 100 ms instead of 50 ms to make the task easier.

All other parameters were the same as in Experiment 1. Participants had to indicate which of the three identities they saw by pressing the left (identity *a*), right (identity *b*), or down (identity *c*) arrow keys.

Participants performed ten blocks containing 84 trials each, for a total of 840 trials. In each block all the images appeared once for every angular location (4 angular locations x 7 morph percentages x 3 morphs = 84). We used 70 data points at every angular location to fit the model for each pair of identities. Thus, we used the responses to different unmorphed images for each pair of identities, ensuring independence of the models.

Before the experimental session participants were shown the identities used in the experiment (corresponding to 0% and 100% morphing, see Figure 3), and practiced the task with 20 trials. These data were discarded from the analyses. Participants performed two experimental sessions at least four weeks apart.

Subjects

Ten participants (five males, mean age 27.30 ± 1.34 SD) participated in Experiment 2, five of which were recruited for Experiment 1 as well. No authors participated in Experiment 2. The sample size (n = 10) was determined using G*Power3 (Faul et al., 2007, 2009) to obtain 80% power at □ = 0.05 based on the correlation of the PSE estimates across sessions in Experiment 1, using a bivariate normal model (one-tailed).

All subjects had normal or corrected-to-normal vision, and provided written informed consent to participate in the experiment. The study was approved by the Dartmouth College Committee for the Protection of Human Subjects.

Familiarity and contact scales

After the two experimental sessions, participants completed a questionnaire designed to assess how familiar each participant was with the identities shown in the experiment. Participants saw each target identity, and were asked to complete various scales for that identity. The questionnaire comprised the "Inclusion of the Other in the Self" scale (IOS) (Aron et al., 1992; Gächter et al., 2015), the "Subjective Closeness Inventory" (SCI) (Berscheid et al., 1989), and the "We-scale" (Cialdini et al., 1997). The IOS scale showed two circles increasingly overlapping labeled "You" and "X", and participants were given the following instructions: *Using the figure below select which pair of circles best describes your relationship with this person. In the figure "X" serves as a placeholder for the person shown in the image at the beginning of this*

section, and you should think of "X" being that person. By selecting the appropriate number please indicate to what extent you and this person are connected (Aron et al., 1992; Gächter et al., 2015). The SCI scale comprised the two following questions: *Relative to all your other relationships (both same and opposite sex) how would you characterize your relationship with the person shown at the beginning of this section?*, and *Relative to what you know about other people's close relationships, how would you characterize your relationship with the person shown at the beginning of this section?* Participants responded with a number between one *(Not close at all)* and seven *(Very close)* (Berscheid et al., 1989). The We-scale comprised the following question: *Please select the appropriate number below to indicate to what extent you would use the term "WE" to characterize you and the person shown at the beginning of this section.* Participants responded with a number between one *(Not at all)* and seven *(Very much so).* For each participant and each identity we created a composite "familiarity score" by averaging the scores in the three scales.

We also introduced a scale aimed at estimating the amount of interaction or contact between the participant and the target identity. The scale was based on the work by Idson and Mischel (2001), and participants were asked to respond Yes/No to the following six questions: *Have you ever seen him during a departmental event?*, *Have you ever seen him during a party?*, *Have you ever had a group lunch/dinner/drinks with him?*, *Have you ever had a one-on-one lunch/dinner/drinks with him?*, *Have you ever texted him personally (not a group message)?*, and *Have you ever emailed him personally (not a*

group email)? The responses were converted to 0/1 and for each participant and for each identity we created a "contact score" by summing all the responses.

For each subject separately, to obtain a measure of familiarity and contact related to each morph, we averaged the familiarity and contact scores of each pair of identities (e.g., the familiarity score of morph *ab* was the average of the scores for identity *a* and identity *b*).

Psychometric fit

For both experiments we fitted a group-level psychometric curve using Logit Mixed-Effect models (Moscatelli et al., 2012) as implemented in *lme4* (Bates et al., 2015). For each experiment and each session, we fitted a model of the form

$$y^k = \text{logit}\left(\beta_0 x + \sum_{i=1}^{n}\left(\beta_i + z_i^k\right) I_i\right)$$

where k indicates the subject, n is the number of angular locations ($n = 8$ for the first experiment, and $n = 4$ for the second experiment), I_i is an indicator variable for the angular location, β_i are the model fixed-effects, and z_i are the subject-level random-effects (random intercept) . From this model, we defined for each subject the Point of Subjective Equality (PSE) as the point x such that logit(x) = 0.5, that is for each angular location

$$PSE_i^k = -\frac{\beta_i}{\beta_0} - \frac{z_i^k}{\beta_0} = PSE_i^p + \Delta PSE_i^k$$

Thus, the PSE for subject *k* at angular location *i* can be decomposed in a population-level PSE and a subject-specific deviation from the population level, indicated with PSE^{p} and ΔPSE^{k} respectively.

In Experiment 2 we fitted three separate models for each of the morph continua. In addition, prior to fitting we removed all trials in which subjects mistakenly reported a third identity. For example, if an image belonging to morph *ab* was presented, and subjects responded with *c*, the trial was removed.

To quantify the bias across locations, we computed a variance score by squaring the ΔPSE_i, and summing them across locations, that is $bias = \sum_{i=1}^{4}(\Delta PSE_i)^2$. Because this quantity is proportional to the variance against 0, throughout the manuscript we refer to it as ΔPSE variance.

Computational modeling

To account for the retinotopic biases we simulated a population of neural units using the Compressive Spatial Summation model (Kay et al., 2013, 2015). This model was originally developed as an encoding model (Naselaris et al., 2011) to predict BOLD responses and estimate population receptive fields in visual areas and face-responsive areas such as OFA, pFus, and mFus (Kay et al., 2015). We refer to activations of neural units that can be thought as being voxels, small populations of neurons, or individual neurons.

The CSS model posits that the response of a neural unit is equal to

$$r = g \cdot a^n$$

with $a = \int G(x, y\,|x_0, y_0, \sigma) S(x, y) dx\, dy$, and $G(x, y|x_0, y_0, \sigma)$ being a 2D gaussian centered at x_0, y_0, with covariance $\Sigma = \sigma I$, and $S(x, y)$ being the stimulus converted into contrast map. The term g represents the gain of the response, while the power exponent n accounts for subadditive responses (Kay et al., 2013). In our simulations we set n = 0.2 as in Kay et al. (2015) when the parameter was not explicitly optimized. Note that the median estimates of n in the three ROIs of interest (IOG, pFus, and mFus) reported by Kay et al. (2015) were 0.20, 0.16, and 0.23 respectively.

We simulated a population of $N = N_a + N_b$ neural units, where N_a indicates the number of units selective to identity a, and N_b indicates the number of units selective to identity b. For simplicity we set N_b = 1 and varied N_a, effectively changing the ratio of units selective to one of the two identities. The stimuli consisted of contrast circles of diameter 4° centered at 7° from the center, and placed at an angle of 45°, 135°, 225°, and 315°, simulating Experiment 2. We simulated the activation of the units assuming i.i.d. random noise normally distributed with mean of 0 and standard deviation of 0.1.

Each experiment consisted of a learning phase in which we simulated the (noisy) response to the full identities a and b in each of the four locations, with 10 trials for each identity and location. We used these responses to train a Support Vector Machine (Cortes and Vapnik, 1995) with linear kernel to differentiate between the two identities based on the pattern of population responses. Then, we simulated the actual experiment by generating responses to morphed faces. For simplicity, we

assumed a linear response between the amount of morphing and the population response. That is, we assumed that if a morph with *m* percentage morphing towards *b* was presented, the population response was a combination of the responses to *a* and *b*, weighted by *(1-m, m)*. The amounts of morphing paralleled those used in the two experiments (0, 17, 33, 50, 67, 83, 100). We simulated 10 trials for each angular location and each amount of morphing, and recorded the responses of the trained decoder. These responses were used to fit a logit model similar to the model used in the main analyses (without random effects), and to estimate the Point of Subjective Equality for each angular location. The sum of these squared estimates around 50% was computed and stored.

We varied systematically the ratio N_a/N_b of units responsive to identity *a*, ranging from 1 to 9, and repeated 500 experiments for each ratio. For each experiment, eccentricity and receptive field size of the units were randomly sampled from a normal bivariate distribution with mean (μ_{ecc}, μ_{size}) and covariance Σ. Once a given sample of eccentricity was drawn, it was converted to a random location lying on the circle of given eccentricity. We also simulated attentional modulations by modifying the gain for the units responsive to identity *a* between 1 and 4 in 0.5 steps, and fixing the gain for identity *b* to 1. We simulated receptive fields in this way from three face-responsive ROIs (Inferior Occipital Gyrus, IOG—also termed OFA—mFus, and pFus).

We obtained the parameter estimates for (μ_{ecc}, μ_{size}) from published results reported in Kay et al. (2015). In particular, we used the median estimates of eccentricity and

receptive field size from Kay et al. (2015) to estimate (μ_{ecc}, μ_{size}) (obtained from Figure S2 for the "face" task). For the three simulated ROIs (IOG, pFus, mFus) these values were (2.05˚, 2.75˚), (2.45˚, 3.68˚), and (1.86˚, 3.41˚) respectively. To estimate Σ, we assumed a standard deviation of 0.5˚ for both eccentricity and receptive field size, and we used the regression fit between eccentricity and receptive field size to estimate their covariance (see Figure S2 in Kay et al., 2015). Prior to estimation, the receptive field sizes were scaled back to pixel values by multiplying them by *n*, with *n* = 0.2 in all simulations. The covariances thus estimated were 0.0546, 0.0541, and 0.0768 for IOG, mFus, and pFus respectively. Data values were extracted from Figure S2 in Kay et al. (2015) using WebPlotDigitizer (https://automeris.io/WebPlotDigitizer).

Code and data availability

Code for the analyses, raw data for both experiments, single subject results, and simulations are available at [REDACTED] as well as Extended Data.

Results

Experiment 1

In this experiment, participants performed a two-alternative forced-choice (AFC) task on identity discrimination. In each trial they saw a face presented for 50 ms, and were asked to indicate which of the two identities they just saw. Each face could appear in one of eight stimulus locations. Participants performed the same experiment with the

same task a second time, at least 33 days after the first session (average 35 days ± 4 days standard deviation).

Participants showed stable and idiosyncratic retinal heterogeneity for identification. The PSE estimates for the two sessions were significantly correlated (see Table 1 and Figure 2B), showing stable estimates, and the within-subject correlations of ΔPSEs (see Methods) was significantly higher than the between-subject correlation (correlation difference: 0.87 [0.64, 1.10], 95% BCa confidence intervals (Efron, 1987); see Table 2), showing that the biases were idiosyncratic (see Figure 2A for example fits for two different subjects).

Table 1.
Correlation of parameter estimates across sessions for the two experiments.

Parameter	r	t	df	p
Experiment 1				
PSE	0.89 [-0.23, 1]	4.86**	6	0.002831
ΔPSE	0.71 [0.47, 0.84]	5.47***	30	6.106e-06
Experiment 2				
PSE	0.98 [0.93, 0.99]	15.22***	10	3.042e-08
ΔPSE	0.64 [0.5, 0.75]	9.02***	118	3.997e-15

Note: *All confidence intervals are 95% BCa with 10,000 repetitions.*
** p < .05. ** p < .01. *** p < .001*

Table 2.
Comparison of within-subjects correlations of parameter estimates across sessions with between-subjects correlations.

Morph	Within-subjects r	Between-subjects r	Difference

Experiment 1			
ab	0.65† [0.57, 0.8]	-0.22 [-0.41, -0.01]	0.87† [0.63, 1.1]
Experiment 2			
ab	0.32 [-0.10, 0.62]	-0.02 [-0.15, 0.11]	0.34 [-0.07, 0.69]
ac	0.62† [0.35, 0.79]	-0.07 [-0.21, 0.08]	0.68† [0.41, 0.92]
bc	0.85† [0.61, 0.95]	-0.08 [-0.27, 0.12]	0.92† [0.68, 1.15]
Note: *All confidence intervals are 95% BCa with 10,000 repetitions.* † *indicates that the CIs do not contain 0.*			

353

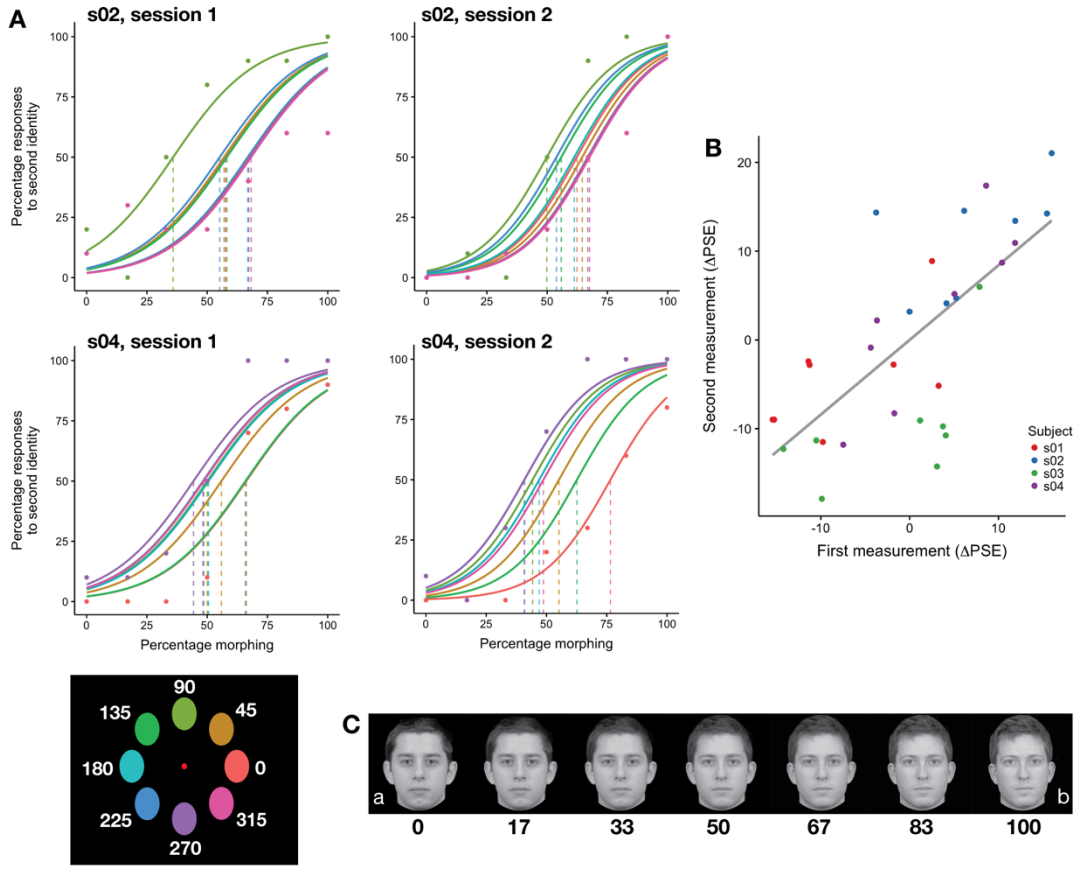

Figure 2. Stable and idiosyncratic biases in identification in Experiment 1. A) Psychometric fit for two subjects from both sessions. Colors indicate location (see colors in bottom left corner); actual data (points) are shown only for the extreme locations to avoid visual clutter. B) The parameter estimates across sessions (at least 33 days apart) were stable (r = 0.71 [0.47, 0.84], see Table 1). Dots represent

individual parameter estimates for each location, color coded according to each subject. Correlations were performed on the data shown in this panel. C) Example morphs used in the experiment. Note that the morphs depicted here are shown for illustration only, and participants saw morphs of identities that were personally familiar to them.

354

355

Experiment 2

In Experiment 1 participants exhibited stable, retinotopic biases for face identification that were specific to each participant. Experiment 1, however, used only two target identities, thus it could not address the question of whether the biases were specific to target identities or to general variations in face recognition that would be the same for all target faces. For this reason we conducted a second experiment in which we increased the number of target identities. In Experiment 2, participants performed a similar task as in Experiment 1 with the following differences. First, each face was presented for 100 ms instead of 50 ms in order to make the task easier, since some participants could not perform the task in Experiment 1; second, each face could belong to one of three morphs, and participants were required to indicate which of three identities the face belonged to; third, each face could appear in four retinal locations instead of eight (see Figure 1) to maintain an appropriate duration of the experiment. Each participant performed another experimental session at least 28 days after the first session (average 33 days ± 8 days SD).

We found that participants exhibited stable biases across sessions for the three morphs (see Table 1 and Figure 3). Interestingly, within-subjects correlations were higher than between-subjects correlations for the two morphs that included the identity *c* (morphs *ac* and *bc*), but not for morph *ab* (see Table 2), suggesting stronger differences in spatial heterogeneity caused by identity *c*. To test this further, we performed a two-way ANOVA on the PSE estimates across sessions with participants and angular locations as factors. The ANOVA was run for each pair of morphs

containing the same identity (e.g., for identity *a* the ANOVA was run on data from morphs *ab* and *ac*), and the PSE estimates were transformed to be with respect to the same identity (e.g., for identity *b* we considered PSE_{bc} and 100 - PSE_{ab}). We found significant interactions between participants and angular locations for identity *b* ($F(27, 120) = 1.77$, $p = 0.01947$) and identity *c* ($F(27, 120) = 3.34$, $p = 3.229e\text{-}06$), but not identity *a* ($F(27, 120) = 1.17$, $p = 0.2807$), confirming that participants showed increased spatial heterogeneity for identities *b* and *c*. The increased spatial heterogeneity for identities b and c, but not a, can be appreciated by inspecting the ΔPSE estimates for each participant. Figure 4A shows lower bias across retinal locations for morph *ab* than the other two morphs, suggesting more similar performance across locations for morph *ab*. To investigate factors explaining the difference in performance across spatial locations between the three identities, we compared the ΔPSE estimates with the reported familiarity of the identities.

The variance of the average ΔPSE estimates across sessions for each subject was significantly correlated with the reported familiarity of the identities ($r = -0.56$ [-0.71, -0.30], $t(28) = -3.59$, $p = 0.001248$), showing that the strength of the retinal bias for identities was inversely modulated by personal familiarity (see Figure 4B). We estimated personal familiarity by averaging participants' ratings of the identities on three scales (Inclusion of the Other in the Self, the We-Scale, and the Subjective Closeness Inventory, see Methods for details). The three scales were highly correlated (min correlation $r = 0.89$, max correlation $r = 0.96$).

399 Because the amount of personal familiarity was correlated with the amount of contact
400 with a target identity (r = 0.45 [0.17, 0.68], t(28) = 2.65,
401 p = 0.01304), we tested whether a linear model predicting ΔPSE with both contact and
402 familiarity as predictors could fit the data better. Both models were significant, but
403 the model with two predictors provided a significantly better fit ($X^2(1) = 6.30$, p =
404 0.0121, log-likelihood ratio test), and explained more variance as indicated by higher
405 R^2: $R^2 = 0.45$, adjusted $R^2 = 0.40$ for the model with both Familiarity and Contact
406 scores ($F(2, 27) = 10.82$, p = 0.0003539), and $R^2 = 0.32$, adjusted $R^2 = 0.29$ for the model
407 with the Familiarity score only ($F(1, 28) = 12.88$, p = 0.001248). Importantly, both
408 predictors were significant (see Table 3), indicating that familiarity modulated the
409 variance of the ΔPSE estimates in addition to modulation based on the amount of
410 contact with a person. After adjusting for the contact score, the variance of the ΔPSE
411 estimates and the familiarity score were still significantly correlated (r_p = -0.42 [-0.61, -
412 0.16], t(28) = -2.42, p = 0.02235).

Table 3. Models predicting variance of the ΔPSE estimates across locations in Experiment 2.

Model	R^2	Score	β	η_p^2	t	p
1	0.32	Familiarity	-0.0574	0.32	-3.59	0.0013
2	0.45	Familiarity	-0.0390	0.17	-2.38	0.0249
		Contact	-0.0452	0.19	-2.512	0.0183

413

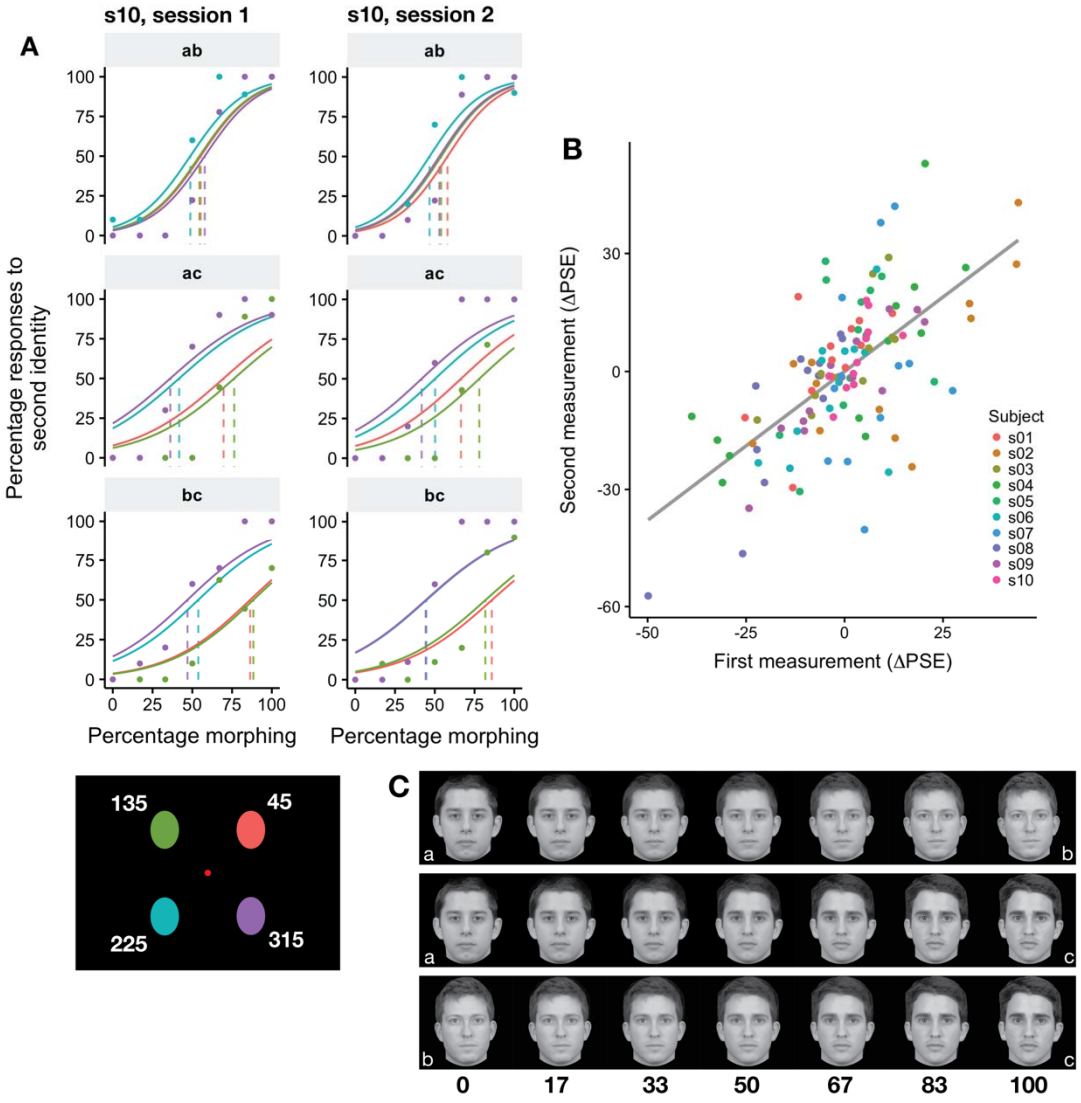

Figure 3. Stable and idiosyncratic biases in identification in Experiment 2. A) Psychometric fit for one subject from both sessions for each of the morphs. Colors indicate location (see colors in bottom left corner); actual data (points) are shown only for the extreme locations to avoid visual clutter. B) The parameter estimates across sessions (at least 28 days apart) were stable (r = 0.64 [0.5, 0.75], see Table 1). Dots represent individual parameter estimates for each location, color coded according to each participant. Correlations were performed on the data shown in this panel. C) Example morphs used in the experiment. Note that the morphs depicted here are shown only for illustration (participants saw morphs of identities who were personally familiar).

414

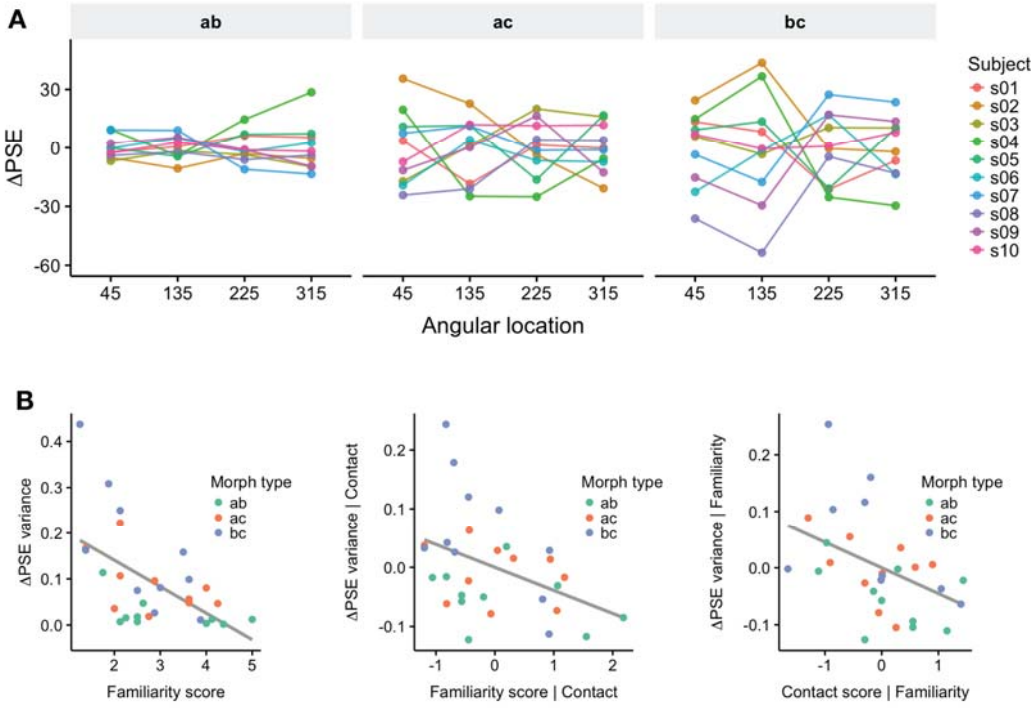

Figure 4. The strength of idiosyncratic biases was modulated by personal familiarity. A) Individual subjects' ΔPSE for each morph, averaged across sessions. Note the difference in variance across locations for the three different morphs (left to right)). B) The variance across locations of ΔPSE estimates was inversely correlated with the reported familiarity of the identities (left panel; r = -0.56 [-0.71, -0.30]), even when adjusting for the Contact score (middle panel; r_p = -0.42 [-0.61, -0.16]). The right panel shows the scatterplot between the Contact score and the ΔPSE variance, adjusted for the Familiarity score, which were significantly correlated as well (r_p = -0.44 [-0.62, -0.17]). See Methods for definition of the Familiarity score and the Contact score. Dots represent individual participant's data, color coded according to morph type. Correlations were performed on the data shown in these panels.

Model simulation

In two behavioral experiments we found a stable, idiosyncratic bias towards specific identities that varied according to the location in which the morphed face stimuli appeared. The bias was reduced with more familiar identities, showing effects of learning. To account for this effect, we hypothesized that small populations of neurons selective to specific identities sample a limited portion of the visual field (Afraz et al., 2010). We also hypothesized that with extended interactions with a person, more neural units become selective to the facial appearance of the identity. In turn, this increases the spatial extent of the field covered by the population and thus reduces the retinotopic bias.

To quantitatively test this hypothesis, we simulated a population of neural units in IOG (OFA), pFus, and mFus using the Compressive Spatial Summation model (Kay et al., 2013, 2015). The parameters of this model were estimated from published results reported in Kay et al. (2015). We simulated learning effects by progressively increasing the number of units selective to one of the two identities, and measuring the response of a linear decoder trained to distinguish between the two identities. As can be seen in Figure 5A, increasing the number of units reduced the overall bias (expressed as variance against 0.5 of the PSE estimates, see *Methods* for details) by increasing the spatial coverage (see Figure 5B). Interestingly, the larger bias was found within the simulated IOG, because the stimuli shown at 7° of eccentricity were at the border of the receptive field coverage (Figure 5B).

438 As an alternative explanation, we tested whether differences in gain could reduce the
439 bias to a similar extent as increasing the number of units. Figure 5C shows that
440 modulating the gain failed to reduce the retinotopic bias in all simulated ROIs.

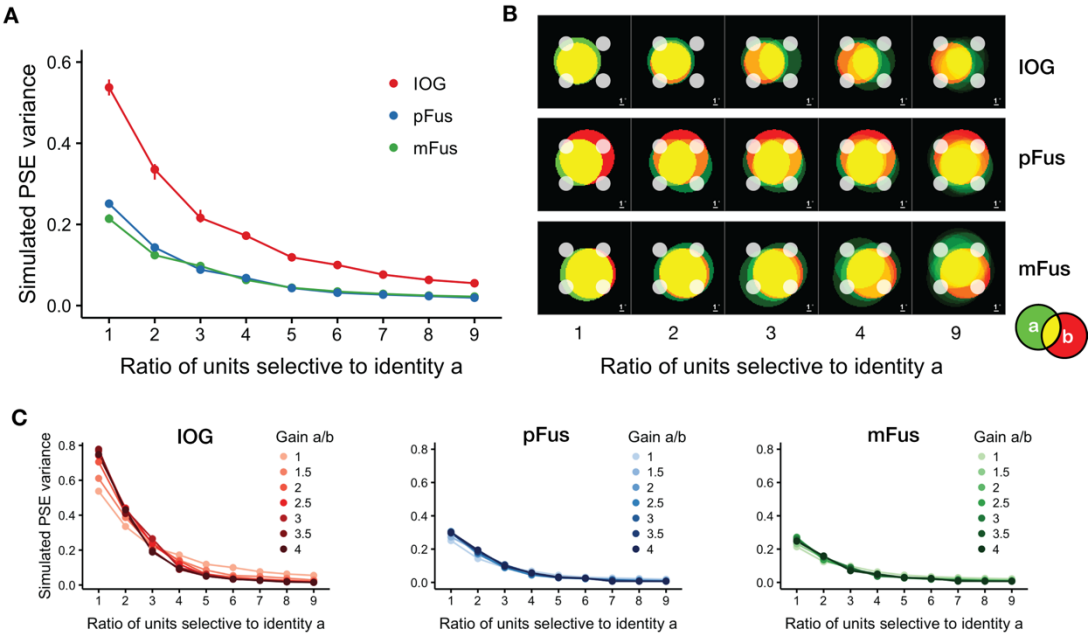

Figure 5. Simulating retinotopic biases and learning effects in face-responsive ROIs. We hypothesized that neural units (voxels, small populations of neurons, or individual neurons) cover a limited portion of the visual field, and that learning increases the number of neural units selective to a particular identity. A) Increasing the number of units selective to one identity reduces the retinotopic bias. Results of simulating 500 experiments by varying the ratio of neural units selective to one of two identities and fixing the gain to 1 for both identities. Dots represent median values with 68% bootstrapped CIs (1,000 replicates; note that for some points the CIs are too small to be seen). In all simulated ROIs the variance of the PSE around 50% decreases with increasing number of units selective to a, but remains larger in IOG because of its receptive field size. B) Example of increasing the number of units selective to one identity. Each colored circle represents the receptive field of a neural unit, color coded according to its preferential identity (green: identity *a*, red: identity *b*, yellow: overlap). Gray circles show location of the stimuli. Each column is normalized to the maximum number of units covering a portion of receptive field. Receptive field are shown as circles with radius 2σ, following the convention in Kay et al., (2015). In IOG, stimuli are at the border of the field covered by the simulated units, resulting in a larger bias across locations compared to pFus and mFus. C) Increasing the gain of the response to one identity fails to reduce the retinotopic bias. We repeated 500 simulated experiments as in A) and modulated the gain of the response of the units selective to identity *a*. Each dot represents median values

of PSE variance for 500 simulated experiments. CIs are not shown to reduce visual clutter.

Discussion

Afraz et al. (2010) reported spatial heterogeneity for recognition of facial attributes such as gender and age, suggesting that relatively independent neural populations tuned to facial features might sample different regions of the visual field. Prolonged social interactions with personally familiar faces lead to facilitated, prioritized processing of those faces. Here we wanted to investigate if this learning of face identity through repeated social interactions also affects these local visual processes, by measuring spatial heterogeneity of identity recognition. We measured whether face identification performance for personally familiar faces differed according to the location in the visual field where face images were presented. We found that participants exhibited idiosyncratic, retinotopic biases for different face identities that were stable across experimental sessions. Importantly, the variability of the retinotopic bias was reduced with increased familiarity with the target identities. These data support the hypothesis that familiarity modulates processes in visual areas with limited position invariance (Visconti di Oleggio Castello et al., 2017a).

These results extend the reports of spatial heterogeneity in visual processing to face identification. Similar biases exist for high-level judgments such as face gender and age (Afraz et al., 2010), as well as shape discrimination (Afraz et al., 2010), crowding, and saccadic precision (Greenwood et al., 2017). Afraz et al. (2010) suggested that

neurons in IT exhibit biases that are dependent on retinal location because their receptive field sizes are not large enough to provide complete translational invariance, and stimuli in different locations will activate a limited group of neurons. In this work, we show that these perceptual biases for face processing not only exist for gender and age judgments (Afraz et al., 2010), but also for face identification and that these biases are affected by learning.

Location-dependent coding in face-responsive areas

Neurons in temporal cortex involved in object recognition are widely thought to be invariant to object translation, that is their response to an object will not be modulated by the location of the object in the visual field (Riesenhuber and Poggio, 1999; Hung et al., 2005). However, evidence suggests that location information is preserved in activity of neurons throughout temporal cortex (Kravitz et al., 2008; Hong et al., 2016). Location information can be encoded as a retinotopic map, such as in early visual cortex, where neighboring neurons are selective to locations that are neighboring in the visual field. In the absence of a clear cortical retinotopic map, location information can still be preserved at the level of population responses (Schwarzlose et al., 2008; Rajimehr et al., 2014; Henriksson et al., 2015; Kay et al., 2015).

Areas of occipital and temporal cortices show responses to objects that are modulated by position (Kravitz et al., 2008, 2010; Sayres and Grill-Spector, 2008). In particular, also face-responsive areas of the ventral core system (Haxby et al., 2000;

Visconti di Oleggio Castello et al., 2017a) such as OFA, pFus, and mFus show responses that are modulated by the position in which a face appears. Responses to a face are stronger in these areas when faces are presented foveally rather than peripherally (Levy et al., 2001; Hasson et al., 2002; Malach et al., 2002). In addition, early face processing areas such as PL in monkeys or OFA in humans code specific features of faces in typical locations. Neurons in PL are tuned to eyes in the contralateral hemifield, with receptive fields covering the typical location of the eyes at fixation (Issa and DiCarlo, 2012). Similarly, OFA responses to face parts are stronger when they are presented in typical locations (de Haas et al., 2016), and OFA activity codes the position and relationship between face parts (Henriksson et al., 2015).

The modulation of responses by object location in these areas seems to be driven by differences in receptive field sizes. In humans, population receptive fields (pRF) can be estimated with fMRI by modeling voxel-wise BOLD responses (Dumoulin and Wandell, 2008; Wandell and Winawer, 2011, 2015; Kay et al., 2013). These studies have shown that pRF centers are mostly located in the contralateral hemifield (Kay et al., 2015; Grill-Spector et al., 2017b), corresponding to the reported preference of these areas for faces presented contralaterally (Hemond et al., 2007). In addition, pRF sizes increase the higher in the face processing hierarchy, favoring perifoveal regions (Kay et al., 2015; Silson et al., 2016). The location-dependent coding of faces in these face-processing areas might be based on population activity, since these areas do not overlap with retinotopic maps in humans (for example, OFA does not seem to overlap with estimated retinotopic maps, Silson et al., 2016, but see Janssens et al., 2014;

Rajimehr et al., 2014; Arcaro and Livingstone, 2017; Arcaro et al., 2017 for work in monkeys showing partial overlap between retinotopic maps and face patches).

Cortical origin of idiosyncratic biases and effects of familiarity

Populations of neurons in visual areas and in temporal cortex cover limited portions of the visual field, with progressively larger receptive fields centered around perifoveal regions (Grill-Spector et al., 2017b). This property suggests that biases in high-level judgments of gender, age, and identity may be due to the variability of feature detectors that cover limited portions of the visual field (Afraz et al., 2010). While the results from our behavioral study cannot point to a precise location of the cortical origin of these biases, our computational simulation suggests that a larger bias could arise from responses in the OFA, given the estimates of receptive field size and eccentricity in this area (Kay et al., 2015; Grill-Spector et al., 2017b). We cannot exclude that this bias might originate in earlier areas of the visual processing stream.

In this work, we showed that the extent of variation in biases across retinal locations was inversely correlated with the reported familiarity with individuals, suggesting that a history of repeated interaction with a person may tune the responses of neurons to that individual in different retinal locations, generating more homogeneous responses. Repeated exposure to the faces of familiar individuals during real-life social interactions results in a detailed representation of the visual appearance of a personally familiar face. Our computational simulation suggests a simple process for augmenting and strengthening the representation of a face. Learning through social

525 interactions might cause a greater number of neural units to become responsive to a
526 specific identity, thus covering a larger area of the visual field and reducing the
527 retinotopic biases. Our results showed that both ratings of familiarity and ratings of
528 amount of contact were strong predictors for reduced retinotopic bias; however,
529 familiarity still predicted the reduced bias when accounting for amount of contact.
530 While additional experiments are needed to test whether pure perceptual learning is
531 sufficient to reduce the retinotopic biases to the same extent as personal familiarity,
532 these results suggest that repeated personal interactions can strengthen neural
533 representations to a larger extent than mere increased frequency of exposure to a
534 face. This idea is consistent with neuroimaging studies showing a stronger and more
535 widespread activation for personally familiar faces compared to unfamiliar or
536 experimentally learned faces (Gobbini and Haxby, 2006; Cloutier et al., 2011; Natu and
537 O'Toole, 2011; Leibenluft et al., 2004; Gobbini and Haxby, 2007; Bobes et al., 2013;
538 Ramon and Gobbini, 2017; Visconti di Oleggio Castello et al., 2017a) .

539 *Effects of attention*

540 Could differences in attention explain the modulation of retinotopic biases reported
541 here? Faces, and personally familiar faces in particular, are important social stimuli
542 whose correct detection and processing affects social behavior (Brothers, 2002;
543 Gobbini and Haxby, 2007). Behavioral experiments from our lab have shown that
544 personally familiar faces break through faster in a continuous flash suppression
545 paradigm (Gobbini et al., 2013), and hold attention more strongly than unfamiliar
546 faces do in a Posner cueing paradigm (Chauhan et al., 2017). These results show that

547 familiar faces differ not only at the level of representations, but also in allocation of
548 attention. At the neural level, changes in attention might be implemented as
549 increased gain for salient stimuli or increased receptive field size (Kay et al., 2015). In
550 an fMRI experiment Kay et al. (2015) reported that population receptive field (pRF)
551 estimates were modulated by the type of task. Gain, eccentricity, and size of the pRFs
552 increased during a 1-back repetition detection task on facial identity as compared to a
553 1-back task on digits presented foveally.

554 To address differences in gain in our computational simulation, we modified the
555 relative gain of units responsive to one of the two identities and found that it did not
556 influence the PSE bias across locations. This bias was more strongly modulated by the
557 number of units responsive to one of the identities. On the other hand, increases in
558 receptive field size and eccentricity could reduce the bias, as shown by differences
559 between the simulated ROIs. However, while the bias was reduced in pFus compared
560 to IOG, for example, the difference in receptive field size (3.68˚ vs 2.75˚) was not
561 sufficient to eliminate the bias with a limited number of neural units. While this result
562 cannot rule out attentional effects completely, it suggests that the retinotopic biases
563 for identification are more strongly modulated by the uneven coverage of the visual
564 field by a limited number of neural units. Additional experiments are needed to
565 further characterize the differences in attention and representations that contribute
566 to the facilitated processing of personally familiar faces.

567 *Implications for computational models of vision*

Many computational models of biological vision posit translational invariance: neurons in IT are assumed to respond to the same extent, regardless of the object position (Riesenhuber and Poggio, 1999; Serre et al., 2007; Kravitz et al., 2008). Even the models that currently provide better fits to neural activity in IT such as hierarchical, convolutional neural networks (Yamins et al., 2014; Kriegeskorte, 2015; Yamins and DiCarlo, 2016) use weight sharing in convolutional layers to achieve position invariance (LeCun et al., 2015; Schmidhuber, 2015; Goodfellow et al., 2016). While this reduces complexity by limiting the number of parameters to be fitted, neuroimaging and behavioral experiments have shown that translational invariance in IT is preserved only for small displacements (DiCarlo and Maunsell, 2003; Kay et al., 2015; Silson et al., 2016; for a review see Kravitz et al., 2008), with varying receptive field sizes and eccentricities (Grill-Spector et al., 2017a). Our results highlight the limited position invariance for high-level judgments such as identity, and add to the known spatial heterogeneity for gender and age judgments (Afraz et al., 2010). Our results also show that a higher degree of invariance can be achieved through learning, as shown by the reduced bias for highly familiar faces. This finding highlights that to increase biological plausibility of models of vision, differences in eccentricity and receptive field size should be taken into account (Poggio et al., 2014), as well as more dynamic effects such as changes induced by learning and attention (Grill-Spector et al., 2017a).

Conclusions

589 Taken together, the results reported here support our hypothesis that facilitated
590 processing for personally familiar faces might be mediated by the development or
591 tuning of detectors for personally familiar faces in the visual pathway in areas that still
592 have localized analyses (Gobbini et al., 2013; Visconti di Oleggio Castello et al., 2014,
593 2017b; Visconti di Oleggio Castello and Ida Gobbini, 2015). The OFA might be a
594 candidate for the cortical origin of these biases as well as for the development of
595 detectors for diagnostic fragments. Patterns of responses in OFA (and neurons in the
596 monkey putative homologue PL, Issa and DiCarlo, 2012) are tuned to typical locations
597 of face fragments (Henriksson et al., 2015; de Haas et al., 2016). Population receptive
598 fields of voxels in this region cover an area of the visual field that is large enough to
599 integrate features of intermediate complexity at an average conversational distance
600 (Kay et al., 2015; Grill-Spector et al., 2017b), such as combinations of eyes and
601 eyebrows, which have been shown to be theoretically optimal and highly informative
602 for object classification (Ullman et al., 2001, 2002; Ullman, 2007).

603 Future research is needed to further disambiguate differences in representations or
604 attention that generate these biases and how learning reduces them. Nonetheless,
605 our results suggest that prioritized processing for personally familiar faces may exist
606 at relatively early stages of the face processing hierarchy, as shown by the local biases
607 reported here. Learning associated with repeated personal interactions modifies the
608 representation of these faces, suggesting that personal familiarity affects face-
609 processing areas well after developmental critical periods (Arcaro et al., 2017;
610 Livingstone et al., 2017). We hypothesize that these differences may be one of the

611 mechanisms that underlies the known behavioral advantages for perception of
612 personally familiar faces (Burton et al., 1999; Gobbini and Haxby, 2007; Gobbini, 2010;
613 Gobbini et al., 2013; Visconti di Oleggio Castello et al., 2014, 2017b; Ramon et al.,
614 2015; Visconti di Oleggio Castello and Gobbini, 2015; Chauhan et al., 2017; Ramon and
615 Gobbini, 2017).

616

References

Afraz A, Pashkam MV, Cavanagh P (2010) Spatial heterogeneity in the perception of face and form attributes. Curr Biol 20:2112–2116.

Arcaro MJ, Livingstone MS (2017) A hierarchical, retinotopic proto-organization of the primate visual system at birth. Elife 6

Arcaro MJ, Schade PF, Vincent JL, Ponce CR, Livingstone MS (2017) Seeing faces is necessary for face-domain formation. Nat Neurosci

Aron A, Aron EN, Smollan D (1992) Inclusion of Other in the Self Scale and the structure of interpersonal closeness. J Pers Soc Psychol 63:596.

Barragan-Jason G, Cauchoix M, Barbeau EJ (2015) The neural speed of familiar face recognition. Neuropsychologia 75:390–401.

Bates D, Mächler M, Bolker B, Walker S (2015) Fitting Linear Mixed-Effects Models Using lme4. J Stat Softw 67:1–48.

Berscheid E, Snyder M, Omoto AM (1989) The Relationship Closeness Inventory: Assessing the closeness of interpersonal relationships. J Pers Soc Psychol 57:792.

Bobes MA, Lage Castellanos A, Quiñones I, García L, Valdes-Sosa M (2013) Timing and tuning for familiarity of cortical responses to faces. PLoS One 8:e76100.

Brothers L (2002) The social brain: a project for integrating primate behavior and neurophysiology in a new domain. Foundations in social neuroscience:367–385.

Bruce V, Young A (1986) Understanding face recognition. Br J Psychol 77 (Pt 3):305–327.

Burton AM, Wilson S, Cowan M, Bruce V (1999) Face Recognition in Poor-Quality Video: Evidence From Security Surveillance. Psychol Sci 10:243–248.

Chauhan V, Visconti di Oleggio Castello M, Soltani A, Gobbini MI (2017) Social Saliency of the Cue Slows Attention Shifts. Front Psychol 8:738.

Cialdini RB, Brown SL, Lewis BP, Luce C, Neuberg SL (1997) Reinterpreting the empathy--altruism relationship: When one into one equals oneness. J Pers Soc Psychol 73:481.

Cloutier J, Kelley WM, Heatherton TF (2011) The influence of perceptual and knowledge-based familiarity on the neural substrates of face perception. Soc Neurosci 6:63–75.

Cortes C, Vapnik V (1995) Support-vector networks. Mach Learn 20:273–297.

de Haas B, Schwarzkopf DS, Alvarez I, Lawson RP, Henriksson L, Kriegeskorte N, Rees G (2016) Perception and Processing of Faces in the Human Brain Is Tuned to Typical Feature Locations. J Neurosci 36:9289–9302.

Diamond R, Carey S (1986) Why faces are and are not special: An effect of expertise. J Exp Psychol Gen 115:107.

DiCarlo JJ, Maunsell JHR (2003) Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. J Neurophysiol 89:3264–3278.

Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. Neuroimage 39:647–660.

Efron B (1987) Better Bootstrap Confidence Intervals. J Am Stat Assoc 82:171–185.

Faul F, Erdfelder E, Buchner A, Lang A-G (2009) Statistical power analyses using G*Power 3.1: tests for correlation and regression analyses. Behav Res Methods 41:1149–1160.

Faul F, Erdfelder E, Lang A-G, Buchner A (2007) G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behav Res Methods 39:175–191.

Gächter S, Starmer C, Tufano F (2015) Measuring the Closeness of Relationships: A Comprehensive Evaluation of the "Inclusion of the Other in the Self" Scale. PLoS One 10:e0129478.

Gobbini MI (2010) Distributed process for retrieval of person knowledge. Social neuroscience: Toward understanding the underpinnings of the social mind:40–53.

Gobbini MI, Gors JD, Halchenko YO, Rogers C, Guntupalli JS, Hughes H, Cipolli C (2013) Prioritized Detection of Personally Familiar Faces. PLoS One 8:e66620.

Gobbini MI, Haxby JV (2006) Neural response to the visual familiarity of faces. Brain Res Bull 71:76–82.

Gobbini MI, Haxby JV (2007) Neural systems for recognition of familiar faces. Neuropsychologia 45:32–41.

Goodfellow I, Bengio Y, Courville A, Bengio Y (2016) Deep learning. MIT press Cambridge.

Greenwood JA, Szinte M, Sayim B, Cavanagh P (2017) Variations in crowding, saccadic precision, and spatial localization reveal the shared topology of spatial vision. Proc Natl Acad Sci U S A 114:E3573–E3582.

Grill-Spector K, Kay K, Weiner KS (2017a) The Functional Neuroanatomy of Face Processing: Insights from Neuroimaging and Implications for Deep Learning. In: Deep Learning for Biometrics (Bhanu B, Kumar A, eds), pp 3–31 Advances in Computer Vision and Pattern Recognition. Cham: Springer International Publishing.

Grill-Spector K, Weiner KS, Kay K, Gomez J (2017b) The Functional Neuroanatomy of Human Face Perception. Annu Rev Vis Sci 3:167–196.

Guntupalli JS, Gobbini MI (2017) Reading Faces: From Features to Recognition. Trends Cogn Sci 21:915–916.

Hasson U, Levy I, Behrmann M, Hendler T, Malach R (2002) Eccentricity bias as an organizing principle for human high-order object areas. Neuron 34:479–490.

Haxby JV, Hoffman EA, Gobbini MI (2000) The distributed human neural system for face perception. Trends Cogn Sci 4:223–233.

Hemond CC, Kanwisher NG, Op de Beeck HP (2007) A Preference for Contralateral Stimuli in Human Object- and Face-Selective Cortex. PLoS One 2:e574.

Henriksson L, Mur M, Kriegeskorte N (2015) Faciotopy-A face-feature map with face-like topology in the human occipital face area. Cortex 72:156–167.

Hong H, Yamins DLK, Majaj NJ, DiCarlo JJ (2016) Explicit information for category-orthogonal object properties increases along the ventral stream. Nat Neurosci 19:613–622.

Hung CP, Kreiman G, Poggio T, DiCarlo JJ (2005) Fast readout of object identity from macaque inferior temporal cortex. Science 310:863–866.

Idson LC, Mischel W (2001) The personality of familiar and significant people: the lay perceiver as a social-cognitive theorist. J Pers Soc Psychol 80:585–596.

Issa EB, DiCarlo JJ (2012) Precedence of the Eye Region in Neural Processing of Faces. Journal of Neuroscience 32:16666–16682.

Janssens T, Zhu Q, Popivanov ID, Vanduffel W (2014) Probabilistic and single-subject retinotopic maps reveal the topographic organization of face patches in the macaque cortex. J Neurosci 34:10156–10167.

Kay KN, Weiner KS, Grill-Spector K (2015) Attention reduces spatial uncertainty in human ventral temporal cortex. Curr Biol 25:595–600.

Kay KN, Winawer J, Mezer A, Wandell BA (2013) Compressive spatial summation in human visual cortex. J Neurophysiol 110:481–494.

Kleiner M, Brainard D, Pelli D, Ingling A, Murray R (2007) What's new in Psychtoolbox-3. Perception

Kravitz DJ, Kriegeskorte N, Baker CI (2010) High-level visual object representations are constrained by position. Cereb Cortex 20:2916–2925.

Kravitz DJ, Vinson LD, Baker CI (2008) How position dependent is visual object recognition? Trends Cogn Sci 12:114–122.

Kriegeskorte N (2015) Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing. Annual Review of Vision Science 1:417–446.

LeCun Y, Bengio Y, Hinton G (2015) Deep learning. Nature 521:436–444.

Leibenluft E, Gobbini MI, Harrison T, Haxby JV (2004) Mothers' neural activation in response to pictures of their children and other children. Biol Psychiatry 56:225–232.

Levy I, Hasson U, Avidan G, Hendler T, Malach R (2001) Center–periphery organization of human object areas. Nat Neurosci 4:533–539.

Livingstone MS, Vincent JL, Arcaro MJ, Srihasam K, Schade PF, Savage T (2017) Development of the macaque face-patch system. Nat Commun 8:14897.

Malach R, Levy I, Hasson U (2002) The topography of high-order human object areas. Trends Cogn Sci 6:176–184.

Moscatelli A, Mezzetti M, Lacquaniti F (2012) Modeling psychophysical data at the population-level: the generalized linear mixed model. J Vis 12

738 Naselaris T, Kay KN, Nishimoto S, Gallant JL (2011) Encoding and decoding in fMRI.
739 Neuroimage 56:400–410.

740 Natu V, O'Toole AJ (2011) The neural processing of familiar and unfamiliar faces: A
741 review and synopsis. Br J Psychol 102:726–747.

742 Poggio T, Mutch J, Isik L (2014) Computational role of eccentricity dependent cortical
743 magnification. arXiv [csLG]

744 Rajimehr R, Bilenko NY, Vanduffel W, Tootell RBH (2014) Retinotopy versus Face
745 Selectivity in Macaque Visual Cortex. J Cogn Neurosci 22:1–10.

746 Ramon M, Gobbini MI (2017) Familiarity matters: A review on prioritized processing of
747 personally familiar faces. Vis cogn:1–17.

748 Ramon M, Vizioli L, Liu-Shuang J, Rossion B (2015) Neural microgenesis of personally
749 familiar face recognition. Proc Natl Acad Sci U S A 112:E4835–E4844.

750 Riesenhuber M, Poggio T (1999) Hierarchical models of object recognition in cortex.
751 Nat Neurosci 2:1019–1025.

752 Sayres R, Grill-Spector K (2008) Relating retinotopic and object-selective responses in
753 human lateral occipital cortex. J Neurophysiol 100:249–267.

754 Schmidhuber J (2015) Deep learning in neural networks: An overview. Neural Netw
755 61:85–117.

756 Schwarzlose RF, Swisher JD, Dang S, Kanwisher N (2008) The distribution of category
757 and location information across object-selective regions in human visual cortex.
758 Proc Natl Acad Sci U S A 105:4447–4452.

759 Serre T, Oliva A, Poggio T (2007) A Feedforward Architecture Accounts for Rapid
760 Categorization. Proc Natl Acad Sci U S A 104:6424–6429.

761 Silson EH, Groen IIA, Kravitz DJ, Baker CI (2016) Evaluating the correspondence
762 between face-, scene-, and object-selectivity and retinotopic organization
763 within lateral occipitotemporal cortex. J Vis 16:14.

764 Sugiura M (2014) Neuroimaging studies on recognition of personally familiar people.
765 Front Biosci 19:672–686.

766 Taylor MJ, Arsalidou M, Bayless SJ, Morris D, Evans JW, Barbeau EJ (2009) Neural
767 correlates of personally familiar faces: parents, partner and own faces. Hum
768 Brain Mapp 30:2008–2020.

769 Ullman S (2007) Object recognition and segmentation by a fragment-based hierarchy.
770 Trends Cogn Sci 11:58–64.

771 Ullman S, Sali E, Vidal-Naquet M (2001) A Fragment-Based Approach to Object
772 Representation and Classification. In: Visual Form 2001, pp 85–100. Springer,
773 Berlin, Heidelberg.

774 Ullman S, Vidal-Naquet M, Sali E (2002) Visual features of intermediate complexity
775 and their use in classification. Nat Neurosci

776 Visconti di Oleggio Castello M, Guntupalli JS, Yang H, Gobbini MI (2014) Facilitated
777 detection of social cues conveyed by familiar faces. Front Hum Neurosci 8:678.

778 Visconti di Oleggio Castello M, Halchenko YO, Guntupalli JS, Gors JD, Gobbini MI
779 (2017a) The neural representation of personally familiar and unfamiliar faces in
780 the distributed system for face perception. Sci Rep 7:12237.

781 Visconti di Oleggio Castello M, Ida Gobbini M (2015) Familiar Face Detection in 180ms.
782 PLoS One 10:e0136548.

783 Visconti di Oleggio Castello M, Wheeler KG, Cipolli C, Gobbini MI (2017b) Familiarity
784 facilitates feature-based face processing. PLoS One 12:e0178895.

785 Wandell BA, Winawer J (2011) Imaging retinotopic maps in the human brain. Vision
786 Res 51:718–737.

787 Wandell BA, Winawer J (2015) Computational neuroimaging and population receptive
788 fields. Trends Cogn Sci 19:349–357.

789 Willenbockel V, Sadr J, Fiset D, Horne GO, Gosselin F, Tanaka JW (2010) Controlling
790 low-level image properties: The SHINE toolbox. Behav Res Methods 42:671–
791 684.

792 Yamins DLK, DiCarlo JJ (2016) Using goal-driven deep learning models to understand
793 sensory cortex. Nat Neurosci 19:356–365.

794 Yamins DLK, Hong H, Cadieu CF, Solomon EA, Seibert D, DiCarlo JJ (2014)
795 Performance-optimized hierarchical models predict neural responses in higher
796 visual cortex. Proc Natl Acad Sci U S A 111:8619–8624.

797 Young AW, Burton AM (2017) Are We Face Experts? Trends Cogn Sci

798

Legends

Figure 1. Experimental paradigm. The left panel shows the experimental paradigm, while the right panel shows the locations used in Experiment 1 (eight locations, top panel) and in Experiment 2 (four locations, bottom panel).

Figure 2. Stable and idiosyncratic biases in identification in Experiment 1. A) Psychometric fit for two subjects from both sessions. Colors indicate location (see colors in bottom left corner); actual data (points) are shown only for the extreme locations to avoid visual clutter. B) The parameter estimates across sessions (at least 33 days apart) were stable (r = 0.71 [0.47, 0.84], see Table 1). Dots represent individual parameter estimates for each location, color coded according to each subject. Correlations were performed on the data shown in this panel. C) Example morphs used in the experiment. Note that the morphs depicted here are shown for illustration only, and participants saw morphs of identities that were personally familiar to them.

Figure 3. Stable and idiosyncratic biases in identification in Experiment 2. A) Psychometric fit for one subject from both sessions for each of the morphs. Colors indicate location (see colors in bottom left corner); actual data (points) are shown only for the extreme locations to avoid visual clutter. B) The parameter estimates across sessions (at least 28 days apart) were stable (r = 0.64 [0.5, 0.75], see Table 1). Dots represent individual parameter estimates for each location, color coded according to each participant. Correlations were performed on the data shown in this panel. C) Example morphs used in the experiment. Note that the morphs depicted here are shown only for illustration (participants saw morphs of identities who were personally familiar).

Figure 4. The strength of idiosyncratic biases was modulated by personal familiarity. A) Individual subjects' ΔPSE for each morph, averaged across sessions. Note the difference in variance across locations for the three different morphs (left to right)). B) The variance across locations of ΔPSE estimates was inversely correlated with the reported familiarity of the identities (left panel; r = -0.56 [-0.71, -0.30]), even when adjusting for the Contact score (middle panel; r_p = -0.42 [-0.61, -0.16]). The right panel shows the scatterplot between the Contact score and the ΔPSE variance, adjusted for the Familiarity score, which were significantly correlated as well (r_p = -0.44 [-0.62, -0.17]). See Methods for definition of the Familiarity score and the Contact score. Dots represent individual participant's data, color coded according to morph type. Correlations were performed on the data shown in these panels.

Figure 5. Simulating retinotopic biases and learning effects in face-responsive ROIs. We hypothesized that neural units (voxels, small populations of neurons, or individual neurons) cover a limited portion of the visual field, and that learning increases the number of neural units selective to a particular identity. A) Increasing the number of units selective to one identity reduces the retinotopic bias. Results of simulating 500 experiments by varying the ratio of neural units selective to one of two

identities and fixing the gain to 1 for both identities. Dots represent median values with 68% bootstrapped CIs (1,000 replicates; note that for some points the CIs are too small to be seen). In all simulated ROIs the variance of the PSE around 50% decreases with increasing number of units selective to a, but remains larger in IOG because of its receptive field size. B) Example of increasing the number of units selective to one identity. Each colored circle represents the receptive field of a neural unit, color coded according to its preferential identity (green: identity *a*, red: identity *b*, yellow: overlap). Gray circles show location of the stimuli. Each column is normalized to the maximum number of units covering a portion of receptive field. Receptive field are shown as circles with radius 2σ, following the convention in Kay et al., (2015). In IOG, stimuli are at the border of the field covered by the simulated units, resulting in a larger bias across locations compared to pFus and mFus.C) Increasing the gain of the response to one identity fails to reduce the retinotopic bias. We repeated 500 simulated experiments as in A) and modulated the gain of the response of the units selective to identity *a*. Each dot represents median values of PSE variance for 500 simulated experiments. CIs are not shown to reduce visual clutter.

Table 1. Correlation of parameter estimates across sessions for the two experiments.

Table 2. Comparison of within-subjects correlations of parameter estimates across sessions with between-subjects correlations.

Table 3. Models predicting variance of the ΔPSE estimates across angular locations in Experiment 2.

Extended Data. The archive contains data from both experiments, as well as the analysis scripts.

Tables

Table 1. Correlation of parameter estimates across sessions for the two experiments.			
Parameter	r	t	df
Experiment 1			
PSE	0.89 [-0.23, 1]	4.86**	6
ΔPSE	0.71 [0.47, 0.84]	5.47***	30
Experiment 2			
PSE	0.98 [0.93, 0.99]	15.22***	10
ΔPSE	0.64 [0.5, 0.75]	9.02***	118
Note: *All confidence intervals are 95% BCa with 10,000 repetitions.* ** p < .05. ** p < .01. *** p < .001*			

Table 2. Comparison of within-subjects correlations of parameter estimates across sessions with between-subjects correlations.			
Morph	Within-subjects r	Between-subjects r	Difference
Experiment 1			
ab	0.65† [0.57, 0.8]	-0.22 [-0.41, -0.01]	0.87† [0.63, 1.1]
Experiment 2			
ab	0.32 [-0.10, 0.62]	-0.02 [-0.15, 0.11]	0.34 [-0.07, 0.69]
ac	0.62† [0.35, 0.79]	-0.07 [-0.21, 0.08]	0.68† [0.41, 0.92]
bc	0.85† [0.61, 0.95]	-0.08 [-0.27, 0.12]	0.92† [0.68, 1.15]
Note: *All confidence intervals are 95% BCa with 10,000 repetitions.* *† indicates that the CIs do not contain 0.*			

Table 3. Models predicting variance of the ΔPSE estimates across angular locations in Experiment 2.

Model	R^2	Score	β	η_p^2	t	p
1	0.32	Familiarity	-0.0574	0.32	-3.59	0.0013
2	0.45	Familiarity	-0.0390	0.17	-2.38	0.0249
		Contact	-0.0452	0.19	-2.512	0.0183

866