NeuroImage 124 (2016) 526–535

Contents lists available at ScienceDirect

NeuroImage

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

Trial-level information for individual faces in the fusiform face area depends on subsequent memory

Adam E. Hasinski, Per B. Sederberg *

Department of Psychology, The Ohio State University, USA

ARTICLE INFO

Article history:
Received 16 April 2015
Accepted 28 August 2015
Available online 4 September 2015

Keywords:
Face perception
Face memory
Fusiform face area
Anterior inferior temporal lobes
fMRI

ABSTRACT

Previous research has shown that face-sensitive brain regions, such as the fusiform face area (FFA) and anterior inferior temporal lobe (aIT), not only respond selectively to face stimuli, but also respond uniquely to individual faces. A common factor in the existing literature is that face stimuli in these experiments are highly familiar to participants, usually by design. We set out to investigate to what extent familiarity correlates with the emergence of face-specific information in face-sensitive regions by testing novel faces with only a single repetition. Our results, consistent with a familiarity hypothesis, demonstrate that the FFA and aIT show face-specific information only when participants demonstrate subsequent memory for those faces. Functionally-defined regions that are not believed to process faces holistically showed no face-specific information, regardless of subsequent memory. To our knowledge, this is the first demonstration of face-specific information in face-sensitive regions for stimuli that were not highly familiar. These results contribute to our understanding of how individuating information comes to be represented in face-sensitive regions and suggest that this process can take place even after a single repetition of a particular face.

Introduction

How the human brain represents face information is a core question in cognitive neuroscience, with implications ranging from machine vision (Tistarelli et al., 2009) to social biases (Van Bavel et al., 2008). Converging evidence from human neuroimaging (Collins and Olson, 2014; Haxby et al., 2000; Kanwisher and Yovel, 2006; Kanwisher et al., 1997) and cellular recordings of nonhuman primates (Freiwald et al., 2009; Tsao et al., 2006) supports the existence of specialized regions in the primate brain that preferentially process faces. It has been proposed that these regions form a face-processing network, with different regions assumed to play unique but potentially overlapping roles in the processing of face stimuli (Collins and Olson, 2014; Haxby et al., 2000; Nestor et al., 2011). However, what exactly each region represents—and how they come to represent it—is not yet understood.

Several key face-processing regions lie in the ventral visual processing pathway. The most well-studied of these, the fusiform face area (FFA) is a portion of the fusiform gyrus that responds preferentially to face stimuli relative to many other categories, such as scenes or objects (Kanwisher and Yovel, 2006; Kanwisher et al., 1997). FFA responds to constituent face features, but it is also sensitive to complete face configurations, suggesting that it may contain holistic representations of faces (Harris and Aguirre, 2010; Liu et al., 2009). In contrast, the more posterior occipital face area (OFA) does not appear sensitive to intact face configurations, suggesting that it represents only the components of faces (Haxby et al., 2000; Liu et al., 2009; Pitcher et al., 2007, 2011). Additionally, the anterior inferior temporal region (aIT)[1] is strongly implicated in both the perception *and* memory of faces (Collins and Olson, 2014). Indeed, lesions to aIT impair the ability to bind episodic and semantic information to individual faces (Gainotti and Marra, 2011; Olson et al., 2013; Ross and Olson, 2010) and produce deficits in face discrimination (Busigny et al., 2014; Olson et al., 2014). In contrast, the evidence for FFA's involvement in face-memory is mixed (Collins and Olson, 2014), with some studies showing differential FFA activity for highly familiar vs. novel faces (Lehmann et al., 2004; Verosky et al., 2013) and others showing no such differentiation (Eger et al., 2005; Gorno-Tempini and Price, 2001). Finally, the OFA is believed to be uninvolved in memory for faces (Collins and Olson, 2014).

The precise function of these regions is still an open question. For instance, the FFA has been viewed as performing face detection (Kanwisher et al., 1998), face identification (Gauthier et al., 2000), or both (Grill-Spector et al., 2004). Taking advantage of advancements in multivariate analysis of fMRI, recent research has shown that face-sensitive regions contain information that differentiates individuals (Anzellotti and Caramazza, 2014). An early study presented one male

* Corresponding author.
E-mail address: sederberg.1@osu.edu (P.B. Sederberg).

[1] We refrain from describing aIT as a "face area" per se because we reserve that phrase for regions that are typically functionally defined as more responsive to faces than other stimulus categories. Due to known issues in imaging the anterior temporal lobe with fMRI (Devlin et al., 2000; Visser et al., 2009), functional localization of a face-sensitive ROI in anterior inferior temporal lobe is less common.

http://dx.doi.org/10.1016/j.neuroimage.2015.08.065
1053-8119/

and one female to participants many times during scanning. Multivariate analyses revealed that information in the aIT, but not the FFA, was distinct for each face (Kriegeskorte et al., 2007). Although this result could reflect the presence of gender information, more recent work also found information in face-selective voxels in the FFA and anterior temporal lobe that individuated different faces of the same gender (Nestor et al., 2011). More recently, the OFA, FFA, and aIT were shown to contain information identifying individuals even when some aspect of individuals' faces changed across presentations. For instance, a pattern classifier trained on faces from several viewpoints was used to effectively discriminate between those faces when tested on novel viewpoints (Anzellotti et al., 2013). Other researchers extracted individuating information from faces that persisted across changes in facial expression (Nestor et al., 2011). Indeed, converging neuroimaging and lesion data suggest that the integration of information across these face areas may be essential for face identification in humans (Anzellotti and Caramazza, 2014; Collins and Olson, 2014; Haxby et al., 2001; Natu et al., 2009).

One commonality across these studies is that they all present the same faces multiple times during scanner sessions, usually with additional presentations during a pre-scanning training phase. Consequently, participants inevitably become familiar with the stimuli used. Often, this training or familiarization process is an intentional, and perhaps critical, element of the experiment. Because familiarization is a constant feature of these studies, it leads to a question about when and how this differentiating information arises in face regions. One possibility is that these regions contain individuating information regardless of familiarity.[2] Face regions may give rise to information for individual faces (i.e. face-specific information) through the bottom-up processing of face stimuli and the representation of each face's unique constellation of features. Under this hypothesis, patterns in face-sensitive regions should individuate faces without any prior exposure. Consequently, face-specific information should be present in both OFA and FFA, which are believed to represent face components (Harris and Aguirre, 2010; Liu et al., 2009). OFA, which appears to more specifically represent lower-level features (Haxby et al., 2000; Liu et al., 2009; Pitcher et al., 2007, 2011), may show greater face-specific information than FFA. Finally, face-specific information may or may not be present in aIT, which is typically associated with more abstracted or view-invariant representations of identities (Anzellotti et al., 2013; Collins and Olson, 2014; Freiwald and Tsao, 2010).

An alternate hypothesis is that familiarization of a face is what leads to this individuating information. According to this latter hypothesis, face regions should only individuate faces that are familiar to participants. Additionally, the face regions more strongly associated with face or person memory, FFA (Grill-Spector et al., 2004) and especially aIT (Collins and Olson, 2014; Gainotti and Marra, 2011; Olson et al., 2013), should contain this face-specific information, whereas OFA may not contain information for individual faces.

The current study directly compares these hypotheses. To do this, we forego extensive training and repeated presentations of faces during scanning. Instead, we use a single repetition (i.e. two presentations) of target faces during scanning, and then test memory for those faces after scanning. We operationalize familiarity as subsequent memory for these target faces. That is, remembered faces will retrospectively indicate whether faces were familiar to a participant at the initial repetition and forgotten faces will have remained unfamiliar during study. According to the familiarity hypothesis, only faces that are subsequently remembered should show significant pattern similarity between presentations. Furthermore, we would expect this similarity to be present in the more downstream regions: FFA and especially aIT. However, under the alternate hypothesis faces should show significant pattern similarity between presentations regardless of subsequent memory performance, and this similarity should be stronger in more upstream regions: FFA and especially OFA.

Potential support for the familiarity hypothesis comes from recent research that found greater pattern similarity between repetitions of *words* that were subsequently remembered relative to the similarity observed between repetitions of forgotten words (Xue et al., 2010).[3] The researchers further argued that their finding provided evidence against the encoding variability hypothesis, which states that repeatedly experienced items are better remembered if the repetitions occur in varied contexts (i.e., with different preceding items), as opposed to all repetitions occurring in the same context (Bray et al., 1976; Martin, 1968; Melton, 1970). Although they did not explicitly manipulate the context in which spaced items were presented, Xue and colleagues equated greater pattern similarity across repetitions with decreased encoding variability. As such, when they observed that items with lower pattern similarity exhibited worse subsequent memory performance, they concluded that this contradicted the encoding variability hypothesis.

Given that the context across spaced repetitions was never the same in the Xue et al. study, the relationship between neural pattern similarity, memory, and context has yet to be fully explored. To that end, we added a secondary question to the current experiment: Does degree of encoding variability influence subsequent memory via altering the degree of similarity between presentations? We manipulated encoding variability by controlling the temporal contexts in which faces appeared across repetitions (Howard and Kahana, 2002; Sederberg et al., 2008; Turk-Browne et al., 2012). This manipulation allowed us to test directly whether encoding variability in the traditional sense is related to subsequent memory. If we found a significant relationship, we would then test whether the effect is mediated by the degree of pattern similarity in face-sensitive regions. However, if we found context to be unrelated to memory behaviorally, then we would then drop the mediation analysis and focus instead on our primary question: What is the relationship between subsequent memory and face-specific information, irrespective of context?

Materials and methods

Subjects

Twenty-seven Ohio State community members (11 female, all right-handed, mean age 22.3 years) participated in the main experiment and a functional localizer task. All subjects had normal or corrected-to-normal visual acuity, provided informed consent, and received monetary compensation. The study protocol was approved by the Institutional Review Board for Human Subjects at the Ohio State University.

Stimuli

Face stimuli were drawn from a pool of 496 color photographs of nonfamous male (243) and female (253) faces. We created this pool by combining images from two sources: The Center for Vital Longevity Face Database (Minear and Park, 2004), and the FEI face database (Thomaz and Giraldi, 2010). We cropped and adjusted individual images so all stimuli would have equal size and brightness. All faces were forward-facing, from the shoulders up, with neutral expression and various hair styles (see Fig. 1). All faces appeared before a white background. A separate set of faces–along with sets of scenes, objects, and scrambled objects–was also used in the functional localizer task. Using separate pools of faces for each task prevented any additional

[2] We use the term "familiarity" in the colloquial sense to refer to memory strength in general, rather than a specific memory process that is distinct from recollection. The current research makes no claim regarding single- vs. dual-process accounts of recognition memory.

[3] In fact, Xue et al. (2010) did show greater pattern similarity for remembered faces in a separate experiment. However, the experiment using face stimuli did not support trial-level analyses, so they could not assess whether or not representations for *individual* faces were more similar across repetitions when subsequently remembered.

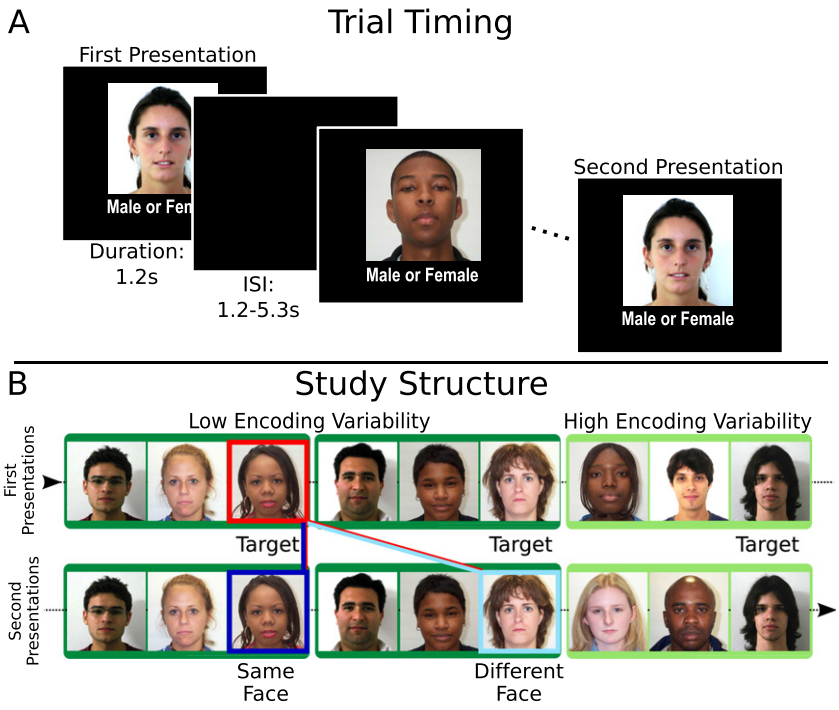

Fig. 1. Overview of experimental design. (A) Faces were presented one at a time, for 1.2 s, with a randomly varied interstimulus interval. Each target face occurred twice. (B) Each presentation of a given target face occurred in either the same context (low encoding variability) or in different contexts (high encoding variability). Same-face similarity was computed between a first-presentation target (red square) and its identical second presentation (dark-blue square). Different-face similarities were computed between the first-presentations and non-identical second presentations with matching genders and encoding variability (light-blue square).

exposures to our target stimuli before the surprise memory test (see below). A Christie digital projector displayed images at 60 Hz with a resolution of 1280 × 1024 onto a screen behind the scanner bore. Participants viewed images with a mirror attached to the head coil.

Procedure

Functional localizer

Participants completed a functional localizer task based on previous work (Epstein and Kanwisher, 1998). This task, which consisted of three 4 min and 24 s runs, enabled us to functionally define ROIs, independently of our main experiment. Consistent with similar tasks used by others (Epstein and Kanwisher, 1998), each localizer run alternated between blocks of scenes, faces, objects and scrambled objects. Each block contained 16 stimuli, all from the same category. Each stimulus was presented for 500 ms, followed by an ISI of 500 ms. Blocks were arranged in a pseudorandom fashion. A block of each category would appear once, in random order, followed by 12 s of fixation. This occurred 3 times per run, with the category blocks arranged in a different order each time. To keep participants engaged during the localizer, they were instructed to detect immediate repetitions of stimuli. In each block, two stimuli were randomly chosen to be repeated immediately. Participants indicated when they detected an immediate repetition via button press. All responses in the scanner were made using a Current Designs fiber-optic response pad.

Main task

The main scanner task consisted of four runs, each lasting 8 min and 48 s. Each run consisted of 96 faces, presented one at a time for 1.2 s, with a randomly jittered inter-stimulus interval (ISI) between 2.3 s and 5.3 s (see Fig. 1A). To keep participants engaged in the task, they were instructed to indicate whether each face was male or female. Participants indicated their response via the response pad.

The main task featured a structure identical to work published previously, yet with different stimuli (Smith et al., 2013). Each target face appeared twice, with between 5 and 20 other faces occurring in between the first and second presentations (see Fig. 1B). Consistent with previous research (Smith et al., 2013; Turk-Browne et al., 2012) and theoretical work on temporal context (Howard and Kahana, 2002; Sederberg et al., 2008), we denoted the two faces that preceded each target as the "context" in which each target presentation occurred. Faces making up the high encoding variability (high-EV) condition were preceded by different context faces on their first and second presentations. That is, each face occurred in a unique temporal context on both presentations. In contrast, faces making up the low encoding variability (low-EV) condition were preceded by the same context faces on their first and second presentations. However, different low-EV face pairs appeared in different contexts.

Across all runs, each participant viewed 16 targets in the high-EV condition and 16 targets in the low-EV condition, with each target being presented twice. Importantly, these stimuli were completely novel to the participants. They had no training or previous exposure to them prior to the study and only minimal familiarization–two presentations–during the scanner task.

Participants were not informed of the organization of the stimulus lists, nor were they informed that some faces were "target" stimuli. They were simply told that they would see a sequence of faces, that some of them may be repeated, and that they should indicate the gender of each face as quickly and accurately as possible. Additional filler faces were interspersed throughout each run, consistent with similar paradigms used previously (Smith et al., 2013). These faces were either presented once or twice to obscure the organizational structure of the target stimuli, but they had no bearing on the hypotheses in question here.

Recognition memory task

Following the last structural scan, participants exited the scanner and completed a recognition memory task on a laptop in an adjacent testing room. They were shown all of the target and comparison faces for all four conditions (96) and 20% of the context faces (38), along

with an equal number (134) of lures and asked whether or not each face appeared in the study task. Again, faces were presented one by one for 1.2 s and subjects had the same amount of time to make a decision. Subjects were given four response options: "sure old," "old," "new," or "sure new", with each corresponding to a different keyboard key. Participants indicated their decisions with a single keypress.

MRI data acquisition

MRI data were acquired with a 3 T Siemens Trio scanner and a 12-channel matrix head coil at the Center for Cognitive and Behavioral Brain Imaging (CCBBI) at the Ohio State University. For the main encoding task, functional images were obtained with a T2-weighted EPI sequence: TR = 2200 ms, TE = 26 ms, field of view = 250 × 220 mm [read × phase], flip angle = 75°, thickness = 2.7 mm (2.5 × 2.5 × 2.7 mm voxels). 41 oblique axial slices were collected in interleaved order. 237 volumes were acquired.

For the functional localizer scans, T2-weighted images were obtained with the following parameters: TR = 3000 ms, TE = 28 ms, field of view = 250 × 220 mm [read × phase]; flip angle = 80°; thickness = 2.5 mm (2.5 × 2.5 × 2.5 mm voxels). 47 oblique axial slices were collected in interleaved order. 85 volumes were acquired.

For each subject, we used the high-resolution T1-weighted anatomical scan to align all functional scans. 160 sagittal slices were collected, with a thickness of 1 mm (1 × 1 × 1 mm voxels).

Data processing

Preprocessing

Processing of both functional and anatomical MRI images was carried out with a combination of AFNI (Cox, 1996) and FSL (Jenkinson et al., 2012; Smith et al., 2004). FSL's brain extraction tool (Jenkinson et al., 2005) was applied to the T1-weighted structural images to isolate voxels containing brain tissue.

For functional scans, we began by discarding the first two volumes of each functional run. Functional runs were preprocessed and analyzed using AFNI (Cox, 1996). Processing then followed standard AFNI protocols (Cox, 2012) and included despiking the time series, aligning the runs together and correcting for head motion. The functional localizer data were then smoothed using a 4 mm FWHM Gaussian kernel. Because our primary analyses involved patterns of activation, we did not smooth the main task data. Finally, all functional data were scaled to have a mean of 100 as part of AFNI's standard data analysis (Cox, 2012).

Modeling the BOLD response

For both functional tasks, we modeled the neural response to each stimulus by fitting canonical hemodynamic response functions (HRFs), convolved with a square function equal to the stimulus duration, in a general linear model (3dDeconvolve in AFNI) (Ward, 2002) for each participant. Each HRF was estimated by fitting its amplitude (i.e. one regressor per HRF) to the data. The result of normalizing our data to have a mean of 100 is that the regression coefficients are equivalent to estimates of percent signal change (Ward, 2002).

For the functional localizer, separate canonical HRF regressors were entered for every stimulus category (face, scene, object and scrambled object), convolved to a .5 s square function, representing stimulus duration. The resulting statistical maps were similar if we used a single 16 s regressor for each category block.

For the main task, all hemodynamic responses were convolved to a 1.2 s square function, representing stimulus duration. We modeled the hemodynamic response to each target face presentation with its own regressor, using the least-squares all method (Rissman et al., 2004). This produced 64 single-trial regressors (2 presentations for each of 16 targets in both the high-EV and low-EV conditions). Extracting estimated BOLD response for these individual trials allowed us to find patterns of activation associated with each individual presentation.

Conventional condition regressors were used to model all non-target stimuli. We estimated a single hemodynamic response for each type of context and filler face. This resulted in 20 additional regressors (16 for context faces and 4 for filler faces). This was done to improve the stability of the estimates of the 64 trials of interest. We also included 6 motion parameters, which were estimated during preprocessing. Finally, we included 4 polynomial regressors to remove slow changes in the BOLD signal (i.e. due to drift in the magnetic field). This procedure resulted in 94 total regressors for the main encoding task.

Note that even though our task is a rapid-presentation, event-related design, the spacing between our target stimuli is considerable (mean: 76.6 s, SD: 21.7 s). Thus, we effectively have a slow-presentation, event-related design where our regressors of interest are concerned, with BOLD responses due to non-target stimuli accounted for with stable condition regressors.

Participants' resulting coefficient maps were transformed to their own high-resolution T1-weighted anatomicals. Registration between functional and anatomical maps was performed using FSL's linear volumetric registration tool (Jenkinson and Smith, 2001; Jenkinson et al., 2002). Coefficient maps in anatomical space were used for all ROI analyses.

Regions of interest

We used the functional localizer to define regions of interest (ROIs) for our main analysis. Consistent with previous research (Anzellotti et al., 2013), we used a faces > scenes contrast to obtain a statistical map of face selectivity. Left and right fusiform face areas (FFAs) were identified separately for each subject. Contrasts were clusterized, using 20 voxels as the minimum cluster size. This ensured that all ROIs had at least 20 voxels in functional space, providing ROIs of sufficient size to extract reliable signal from our task data.

The process for selecting ROIs involved three steps: First, for each hemisphere, we focused on a broad anatomical region that coincides with established FFA location (e.g. the fusiform gyrus). Then we set an initial threshold equivalent to $p = .0001$. If no clusters survived this initial threshold we did not define an ROI for that subject and region. Finally, if a cluster was found, we decreased the p-value until the cluster shrank to 20 voxels, or as close to 20 as possible without breaking up. All ROIs were between 20 and 40 2.5 × 2.5 × 2.7 mm voxels in size.

ROIs were then transformed into each subject's high-resolution anatomical space. We then created a mask out of the 200 most significant 1 × 1 × 1 mm voxels in anatomical space. This allowed us to reliably identify anatomical landmarks and ensured that ROIs would align with our task data, which were also transformed to anatomical space. Finally, we combined–when possible–left and right hemispheres into a single mask for each ROI. This was done because we did not have specific hypotheses regarding hemisphere and we wanted to keep the number of statistical comparisons to a minimum.

This process enabled us to create OFA and FFA cluster in 26 of 27 participants. We were also able to create a parahippocampal place area (PPA) (Epstein and Kanwisher, 1998) ROI in 26 of 27 participants by following the same process but reversing the statistical contrast to be scenes > faces. We use PPA as a control region in the following analyses to ensure that we do not get face-specific activation in regions not known to be sensitive to faces.

We were unable to functionally localize an anterior temporal face area in more than a small fraction of our participants. This was likely due to signal loss in the ventral anterior temporal lobe, a known issue (Devlin et al., 2000; Visser et al., 2009). As an alternative, we created an anatomically defined 400-voxel (1 mm^3 voxels) anterior inferior temporal lobe ROI (aIT) masks for each participant, using the Harvard–Oxford atlas distributed by FSL (Jenkinson et al., 2012; Smith et al., 2004) as a reference. We treat this analogue as a substitute for a functionally defined aIT.

The use of participant-specific ROIs obviated the need for transforming any data to a standard, across-participant space. Therefore,

we perform this second transformation on our ROI masks only to aggregate ROIs across participants for graphical purposes (see Fig. 2). This anatomical-to-standard space transformation was accomplished with FSL nonlinear volumetric registration tool (Andersson et al., 2007). ROIs are aligned to the Montreal Neurological Institute (MNI-152) atlas provided by FSL (Mark Jenkinson et al., 2012; Smith et al., 2004).

Calculating face-specific information

To investigate the presence of face-specific information, we used each participant's masks to extract a pattern of neural activity at each ROI for each target presentation. We then used a trial-level variant of correlation analysis (Haxby et al., 2001; Xue et al., 2010). Specifically, for each subject we correlated activity for each first-presentation target with activity for every second-presentation target. This produced two types of correlations: "same-face" correlations between each first presentation and the identical second presentation, and "different-face" correlations between each first presentation and all non-identical second presentations. We then applied a Fisher z-transformation to all correlations, converting them to similarity values on the interval $(-\infty, \infty)$.

The same-face similarities reflect the dependency or overlapping information between the neural representations of first and second presentations of identical faces. The different-face similarities are important because collectively they quantify the amount of overlapping information between neural representations of first-presentation faces to *other* non-identical faces on average. In other words, different-face similarities let us estimate the amount of overlapping information that is due to the stimulus category of faces. We use these different-face similarities to estimate the amount of information in the same-face similarity values that can be attributed to representing the same exact face rather than the same stimulus category. Specifically, for each first presentation, we created a vector of face-specific similarity values by subtracting each different-face similarity from the same-face similarity:

$$\mathbf{f_i} = s_i - \mathbf{d_i} = [s_i - d_{i1}, s_i - d_{i2}, \ldots, s_i - d_{in}], \quad (1)$$

where s_i is the scalar same-face similarity between ith first presentation and its identical second presentation, $\mathbf{d_i}$ is the vector of different-face similarity values between the ith first presentation and all non-identical second presentations, $\mathbf{d_{in}}$ is the scalar different-face similarity between the ith first presentation and nth different-face second presentation, and $\mathbf{f_i}$ is the resulting vector of face-specific similarity values.

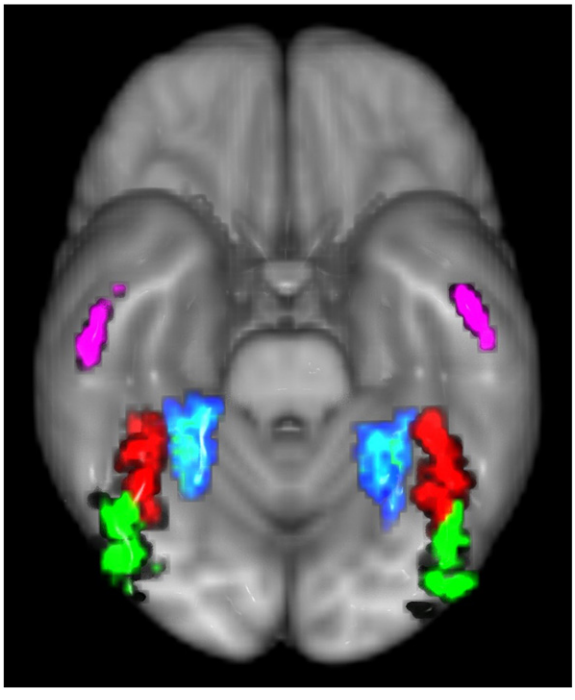

Fig. 2. Functional and anatomical ROIs. Superposition of every subject's FFA (red), OFA (green) and PPA (blue), along with aIT (violet) in MNI standard space.

The different-face similarity values are meant to control for any pattern similarity that is not due to specific faces. Therefore, we avoid using different-face similarities that may be low due to systematic differences between first and second presentations. Past research has shown that full multivariate techniques can be used to distinguish representations of male and female faces (Kriegeskorte et al., 2007). In searching for face-specific information in face-sensitive regions, we did not want our results to be due to gender differences between faces. Similarly, we only use different-face similarity values when the same-face and different-face targets have the same gender. Additionally, we only use different-face similarity values when the different-face targets occur in the same encoding variability condition as the same-face target. Finally, we only use different-face similarity values when the subsequent memory performances of the same-face targets and different-face targets match.

The eligible different-face similarities were used to calculate face-specific similarity values according to Eq. (1). We use these resulting values as the dependent variable in our neural analyses.

Mixed-effects regression

To test the competing hypotheses that face-specific information is either dependent on subsequent memory or not, we fit our data with a mixed-effects regression model. We performed a separate analysis for each of our four ROIs. Our predictor variable of interest is an indicator variable for whether or not each target face was recognized during the memory test. If this recognition term is significant in a given model, it would indicate that face-specific information is contingent on familiarity in that ROI. In contrast, a significant intercept term would indicate that face-specific information occurs in a given ROI regardless of familiarity. Additionally, we wanted to rule out the possibility that any effect, or lack of effect, could actually be due to the spacing between faces or overall activation within ROIs. We included two nuisance regressors to control for these issues.

First, we needed to control for the variation in lag between the first and second presentations. Autocorrelation is a known feature of fMRI time-series data. In general, when two presentations are closer to one another in time, the neural patterns they evoke maybe more similar regardless of the actual evoked neural activity. Thus, the lag between the first and second presentations is a potential source of variance we must account for. More specifically, because our face-specific similarity values are the difference between two values–same-face similarity and different-face similarity–we needed to account for the relationship between same-face lags and different-face lags.

To accomplish this we created a lag–log-ratio term, which we will explain further: We begin by calculating the same-face lags, the spacing between first presentations and their identical second presentations, and the different-face lags, the spacing between first presentations and non-identical second presentations. We then log-transform both types of lags. Finally we subtract the different-face log–lags from the same-face log–lags:

$$\mathbf{l_i^r} = \log(l_i^s) - \log\left(\mathbf{l_i^d}\right) = \left[\log\left(\frac{l_i^s}{l_{i1}^d}\right), \log\left(\frac{l_i^s}{l_{i2}^d}\right), \ldots, \log\left(\frac{l_i^s}{l_{in}^d}\right)\right], \quad (2)$$

where l_i^s is the lag between the first and second presentations of the ith face, $\mathbf{l_i^d}$ is the vector of lags between the ith first presentation and all different-face second presentations, $\mathbf{l_{in}^d}$ is the scalar lag between the ith first presentation and the nth different-face second presentation,

and $\mathbf{l_i^r}$ is the resulting vector of lag–log-ratios. Note that these values are the log of the ratio between same-face lags over different-face lags.

The resulting lag–log-ratios have three useful properties. Firstly, they combine both same-face and different-face lags into a single value, keeping the complexity of our inferential models to a minimum. Secondly the ratio preserves the *relative* difference in spacing between same-face pairs and different-face pairs. Finally, the log transformation provides an intuitive and interpretable zero point. That is, when the same-face lag equals the different-face lag, the resulting lag–log-ratio value is zero. Thus, this lag-ratio variable quantifies uneven lag between same- and different-face similarity values. As an additional safeguard against autocorrelation, we discarded pairings where the lag between the first and second presentation was less than 30 s.

We also wanted to ensure that any face-specific pattern similarity was not due to, or obscured by, changes in overall level of activation from first to second presentations. This repetition modulation (RM) has already been shown to covary with familiarity (Eger et al., 2005), so we created a nuisance regressor to partial out any covariation between subsequent memory and face-specific similarity that is shared with RM. We follow a similar approach to our calculation of lag–log-ratios above. First we calculate the mean difference in activation between first and second presentations (RM) for same-face and different-face pairings. We then log-transform these RM values. Finally, we take the difference between all same-face log-RMs and all the corresponding different-face RMs:

$$\mathbf{\overline{m}_i^r} = \log(\overline{m}_i^s) - \log\left(\mathbf{\overline{m}_i^d}\right) = \left[\log\left(\frac{\overline{m}_i^s}{\overline{m}_{i1}^d}\right), \log\left(\frac{\overline{m}_i^s}{\overline{m}_{i2}^d}\right), \ldots, \log\left(\frac{\overline{m}_i^s}{\overline{m}_{in}^d}\right)\right], \tag{3}$$

where $\overline{m}_i^s$ is the mean RM between the first and second presentations of the ith face, $\mathbf{\overline{m}_i^d}$ is the vector of mean RMs between the ith first presentation and all different-face second presentations, $\mathbf{\overline{m}_{in}^d}$ is the scalar mean RM between the ith first presentation and the nth different-face second presentation, and $\mathbf{\overline{m}_i^r}$ is the resulting vector of RM–log-ratios. These values quantify the relative difference in RM between same- and different-face pairings, such that the values would be zero when the relative difference was zero.

We include both lag–log-ratio and RM–log-ratio as nuisance regressors in our statistical model; however, we have no hypotheses concerning these variables. Therefore, we avoid making inferences about them, and do not report statistics or significance in this paper.

Finally, our model included three random effects terms: Firstly, individual brains may differ in the degree to which they exhibit face-specific similarity. Therefore, we treated participant as a random effect. This allowed us to account for any variation in face-specific similarity that may be due to the particular participants in our sample. Additionally, some faces may tend to evoke more distinct neural responses than others. To control for any differences in face-specific similarity due to stimulus distinctiveness, we treated each same-face target and each different-face target as random effects.

Determining degrees of freedom for mixed effects regression models is not straightforward. Instead, parametric statistical inference requires an estimation of the degrees of freedom (Kenward and Roger, 1997). To circumvent this estimation, we performed non-parametric permutation tests to ascertain the p-values for each of our comparisons of interest (Ernst, 2004). For each ROI analysis, we performed 10,000 permutations of the dependent data by randomly flipping the sign of the similarity difference. We then calculated two-tailed p-values based on the null distribution of t-values for each regression term of interest. Finally, to correct for multiple comparisons we used a Benjamini–Hochberg correction (Genovese et al., 2002) to keep the false discovery rate at $\alpha = 0.05$ across all tests. For significant results, we provide the raw p-values, followed by corrected p-values in parentheses.

Results

Behavioral results

Participants responded correctly on 93.1% of all trials of the gender categorization task during encoding. We defined incorrect responses as those where the participant provided the wrong gender or failed to make a response within the 1.2 s the face was on the screen. Only face pairs where the participant responded correctly during both encoding presentations were included in subsequent neural similarity analyses. This ensured that differences in similarity were not due to a lack of attention to the faces.

The mean reaction time for completed trials was 647.1 ms (SD = 158.1 ms). By means of a mixed effects regression analysis, we looked for evidence that the first presentation of a face would facilitate the gender judgment response for the second presentation of the same face. We did not find evidence of overall repetition priming (mean = 4.03 ms, SD = 41.48 ms), $t = 0.37$, $p = 0.701$. We also found no difference in repetition priming between target faces that were subsequently remembered vs. forgotten, $t = 0.82$, $p = 0.411$, or between faces learned under high vs. low encoding variability, $t = -0.15$, $p = 0.883$. Finally, the interaction of subsequent memory and encoding variability also did not produce repetition priming, $t = -0.58$, $p = 0.558$.

Turning to the recognition test, we calculated hit rates and false-alarm rates for each subject. The mean hit rate for all target items was 51.6% (SD = 16.5%), with a corresponding mean false-alarm rate of 35.9% (SD = 14.0%). As expected, given the use of unfamiliar faces, performance on the memory task was poor, yet, on average, participants performed above chance. For a comparison, the mean hit rate for non-target, once-presented faces was 41.6% (SD = 12.18%). As would be expected, a paired samples t-test revealed a significant boost in memory performance for twice-presented items, $t(26) = 4.17$, $p < 0.001$. Additionally, hit rates for once-presented items were significantly higher than false-alarm rates for lures, $t(26) = 2.45$, $p = .021$.

Fig. 3 shows the memory performance for target trials split by condition. We find a mean hit rate of 50.9% (SD = 18.9%) for faces presented

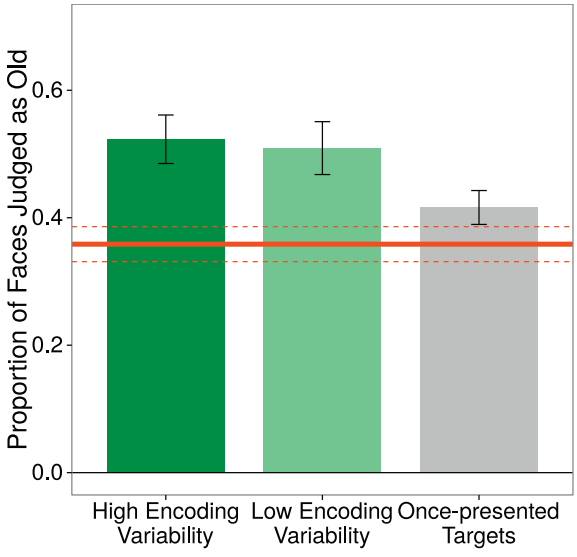

Fig. 3. Subsequent memory performance as a function of encoding condition. Bars show mean proportion of faces recognized during the memory test for faces that occurred under high encoding variability (dark green) and low encoding variability (light green). For a reference, the mean proportion of once-presented items that were subsequently remembered is shown in gray. Error bars reflect standard errors. The solid red line indicates the mean proportion of lures that were incorrectly judged as having been seen previously, with dashed red lines indicating standard error.

twice in the same context (low-EV), and a mean hit rate of 52.3% (SD = 17.4%) for those presented in different contexts (high-EV). A paired-samples t-test did not detect a significant difference between conditions, $t(26) = -0.47$, $p = 0.642$. Thus, at least in this experiment, variability in encoding context does not appear to be related to memory performance. Because encoding variability is unrelated to memory performance, we forego investigating whether this (null) relationship is mediated by neural similarity. We collapse across encoding variability conditions for the remainder of our analyses, focusing instead on a potential relationship between similarity and subsequent memory.[4]

At test, mean reaction time for target probes was 860.9 ms (SD = 160.4 ms). A mixed effects regression analysis revealed that participants were significantly faster to incorrectly judge target probes as 'new' (mean = 834.1, SD = 174.9) than to correctly judge them as old (mean = 881.8, SD = 144.8), $t = 2.68$, $p = 0.006$. No difference in reaction times was found between target probes that occurred under high or low encoding variability, $t = 0.65$, $p = 0.534$. Additionally, reaction time was not affected by the interaction of test judgment and encoding variability, $t = -0.28$, $p = 0.772$.

Neuroimaging results

For each of our four ROIs, we tested for a baseline effect of face-specific similarity. We also tested for a difference in face-specific similarity between forgotten and remembered faces. Finally, we tested whether face-specific similarity was present for subsequently remembered faces. This resulted in a total of 12 tests across ROIs. Faces that participants judged as "old" or "sure old" were defined as remembered, while forgotten faces were those that participants judged as "new" or "sure new."[5] We conducted these tests using a linear mixed-effects regression procedure, which let us control for nuisance regressors (lag between presentations and repetition modulation), as well as participant-specific and stimulus-specific effects. We controlled for multiple comparisons using a Benjamini–Hochberg correction (Genovese et al., 2002) to keep the false discovery rate at $\alpha = 0.05$. Corrected p-values are shown in parentheses for significant results.

In FFA, no baseline effect of face-specific similarity was found, $t = 0.11$, $p > 0.91$. However, face-specific similarity was greater for subsequently remembered faces relative to forgotten faces, $t = 3.29$, $p = 0.001$ (.007). Furthermore, significant face-specific similarity was found for remembered faces alone, $t = 2.57$, $p = 0.012$ (0.046) (see Fig. 4A).

In aIT, no baseline effect of face-specific similarity was found, $t = 0.85$, $p > .35$. Once again, face-specific similarity was greater for remembered than forgotten items; however, this finding did not survive correction for multiple comparisons $t = 1.99$, $p = 0.045$ (0.135). Nevertheless, we did find significant face-specific similarity for remembered faces, $t = 2.39$, $p = 0.007$ (0.043) (see Fig. 4B).

In OFA, an upstream face-sensitive region that is generally believed to process face components, we found no baseline effect of face-specific similarity, $t = 0.58$, $p > 0.57$. Furthermore, we found no effect of subsequent memory on face-specific similarity, $t = -1.16$, $p > 0.23$. Finally, we found no evidence of face-specific similarity for remembered faces, $t = -0.33$, $p > 0.74$.

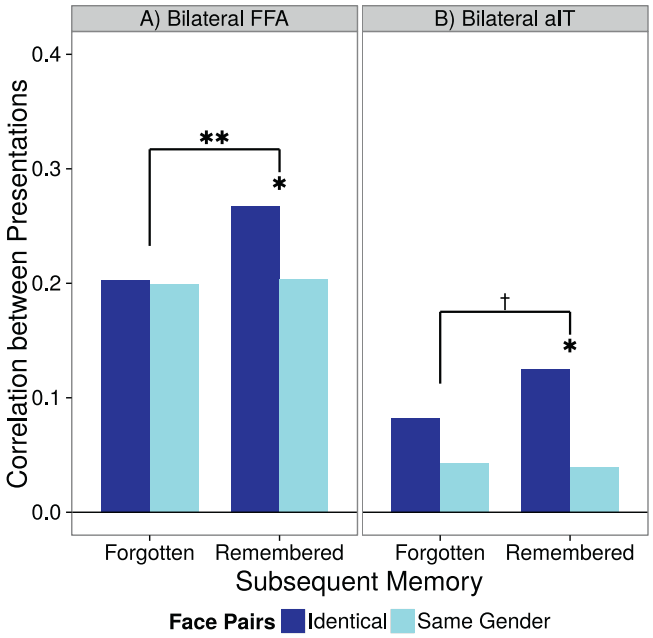

Fig. 4. Average similarity between face repetitions. Bars show similarity between first presentation targets and either (dark blue) same or (light blue) different second presentations, split by memory performance. (A) FFA shows significant face-specific information–greater similarity between identical faces than between different faces of the same gender–only for faces that are subsequently remembered. The difference in face-specific similarity between remembered and forgotten faces is also significant. (B) aIT shows the same pattern, although the difference in face-specific information is no longer significant after correcting for multiple comparisons. † $p < 0.05$ uncorrected, * $p < 0.05$ corrected, ** $p < 0.01$ corrected.

As expected, our control region—the scene-selective PPA—similarly showed no baseline effect of face-specific similarity, $t = -0.37$, $p > 0.69$, and no effect of subsequent memory on face-specific similarity, $t = 0.89$, $p > .36$. PPA also showed no face-specific similarity for remembered faces, $t = 0.33$, $p > 0.74$.

In support of the familiarization hypothesis, we found that memory performance was a significant predictor of face-specific similarity in downstream face-sensitive regions. We found no evidence of face-specific information in OFA, which is consistent with the idea that subsequent memory for faces—not their individual features—is an important moderator.

Discussion

The present results demonstrate that downstream face-processing regions, such as the FFA and aIT, contain unique face-specific information after minimal familiarization. We find no evidence of such face-specific information for faces that participants are unable to subsequently recognize. This was unlikely to be due to participants not attending to those stimuli because we only included faces in our analyses that the participants successfully categorized as male or female within 1.2 s. We also find no evidence of face-specific information in an early face-selective region (OFA) or in a control region that is not specifically sensitive to faces (PPA).

Our finding is consistent with recent research (Cowen et al., 2014) that reconstructed face stimuli from patterns of activation in the fusiform gyrus. These reconstructed faces were subjectively (i.e. according to participant ratings), but not objectively (i.e. according to Euclidean distance), similar to the face stimuli that elicited those patterns of activation. However this research focused on comparing face reconstructions rather than the actual distributed neural representations elicited by each face presentation. We, therefore, feel that the present research

[4] Although our encoding variability manipulation did not produce a memory effect, we performed additional analyses to see if encoding variability influenced the amount of face-specific information present in any ROI. We found no significant relationship between encoding variability and the amount of face-specific similarity. Thus, degree of encoding variability did not predict either variable of interest.

[5] For each ROI, we also performed identical regression analyses using participants' confidence ratings. The outcome of these analyses qualitatively paralleled our results using hits and misses: Higher confidence was significantly and marginally related to greater face-specific similarity in FFA and aIT, respectively. No relationship was found in OFA or PPA.

complements these previous findings. Furthermore, the previous research (Cowen et al., 2014) did not discuss the role of familiarity. Our results suggest that future work into the reconstruction of faces based on neural activity may benefit from the inclusion of a subsequent memory test.

We used subsequent memory as an index of the strength of encoding into memory across both presentations of a face; however, the mechanism that drives this encoding–similarity relationship has not yet been explained. One possibility is a reactivated-trace account proposed previously (Xue et al., 2010). According to this account, memory is most improved when the same trace of a stimulus is reactivated on subsequent presentations. Thus, greater face-specific similarity actually reflects this stable reactivation, which gives rise to stronger memory encoding during each presentation of a given face. Importantly, the current findings suggest that this reactivation refers to the representation of the face itself, rather than the face and its context as it has been previously discussed (Xue et al., 2010). This account is also consistent with recent evidence demonstrating that the fusiform gyrus is involved in memory encoding for pictorial stimuli (Garoff et al., 2005; Kim, 2011).

A second possibility is a feedback account where patterns in face areas reflect online perceptual processing that is associated with enduring memory representations in other regions, such as the hippocampus (Buzsáki and Moser, 2013; Eichenbaum, 2000, 2004) or regions that have been implicated in person memory, such as the anterior paracingulate cortex or amygdala (Gobbini and Haxby, 2007). These memory representations may feed back to and alter the perceptual representations in upstream areas such as the FFA during the second presentation of a face. Thus, the stronger a memory representation is after a first presentation, the more that memory representation may influence a subsequent perceptual representation to be like the first. This feedback account is also consistent with previous research that found greater similarity between the encoding and retrieval representations of scenes that were subsequently remembered vs. forgotten. This increased similarity was further shown to be mediated by hippocampal activity (Ritchey et al., 2012).

The critical difference between these accounts is the directionality of the flow of information. According to the reactivated-trace account, stimulus information moves from perceptual regions to memory regions. The more similar the perceptual information is across two presentations of a face, the more stable the memory representation for that face will be. In contrast, according to the feedback account, information moves bidirectionally between perceptual and memory regions. This creates a dynamical system whereby better encoding during the first presentation facilitates reactivation of the original pattern during the second presentation. Consequently, the former account posits that subsequent memory is a function of perceiving face-specific information. The latter account posits that the perception of and memory for face-specific information influence each other.

The current findings do not clearly support either account over the other. However, the lack of repetition priming for subsequently remembered faces is perhaps more consistent with the reactivated-trace account. The feedback account claims that information stored in memory facilitates perceptual representations. Hence, the hypothesis could be made that such feedback would also facilitate gender discrimination judgments, as evidenced by shorter response latencies on second presentations. The lack of any such priming could be viewed as minor evidence against the feedback account. Nevertheless, future work will be required to fully explain the neural mechanisms underlying the current findings.

An alternate explanation for the current results is that attention, not memory encoding per se, is the mechanism that produces face-specific similarity. Specifically, faces observed under greater attention may be encoded with greater fidelity. This may also produce more distinctive representations of those faces in face-sensitive regions. Greater attention during encoding also facilitates subsequent memory (Chun and Turk-Browne, 2007; Turk-Browne et al., 2013). While we are agnostic with regard to potential attention-based mechanisms at this time, future research could use eye-tracking and pupillometry to provide quantitative indices of attentional focus and cortical arousal to help shed light on this hypothesis.

A more extreme characterization of this mechanism is that the lack of face-specific similarity for forgotten items is due to a complete lack of attention during forgotten trials. Specifically, face areas will not represent stimuli if a participant's eyes were closed or looking away. This would also hamper memory performance for faces, producing a correlation. However, our neural analyses only used face pairs that were correctly categorized by gender during *both* presentations. Performing this task accurately requires at least some degree of attention. Thus, the current results are likely not driven by an *absence* of attention during subsequently forgotten trials.

Furthermore, the current results suggest that attention alone may not be sufficient to produce face-specific information. If attention were the sole mechanism at work, we would expect face-specific similarity in OFA precisely because we used the same image for both presentations of our target faces. The OFA principally represents lower-level features of faces (Haxby et al., 2000; Liu et al., 2009; Pitcher et al., 2007, 2011), so greater attention should also yield greater representational fidelity of those features in FFA. That is, the low-level features are the same, so attention to those features could have produced "features-specific" information in OFA. We found no evidence of face-specific similarity for familiar faces in OFA. This is consistent with earlier research that found no effect of familiarity in OFA when using univariate methods (Davies-Thompson et al., 2009). Our null result in OFA suggests that something more than attention is necessary to produce face-specific information in FFA and aIT.

Another potential explanation is that both face-specific similarity and memory are driven by the features of our stimuli. That is, some faces may be more visually distinctive than others. Greater distinctiveness could produce both greater face-specific similarity and better memory performance. By including face labels as a random factor in our statistical analyses, we were able to control for variation in similarity that was attributable to the varying distinctiveness of different faces. However, this procedure does not control for within-participant variation in face distinctiveness. That is, a particular face may not be seen as distinctive by the entire sample, but could be particularly distinctive for a given participant. This distinctiveness could still affect that participant's memory and neural response pattern.

The present results are limited by the difficulty in imaging the anterior temporal lobes (Devlin et al., 2000; Visser et al., 2009). If we had been able to reliably extract functionally-localized anterior temporal face areas from our participants, we may have found stronger or different effects in that region. Nonetheless, the fact that our atlas-based aIT results followed the same pattern as the FFA results is evidence that familiarity may serve the same function in both regions. fMRI protocols optimized to image the anterior temporal lobes may provide better functional localization for future work.

Although the current work relied on repetitions of identical images of the target faces, future research may benefit from including alternate images of the same face (e.g. from different viewpoints) for first and second presentations. Previous work has demonstrated that the FFA and aIT contain patterns that discriminate highly familiar identities, even if alternate images are used for the same face (Anzellotti and Caramazza, 2014). This image-invariant representation of identity has been shown with respect to variations in viewing angle (Anzellotti et al., 2013; Natu et al., 2009) and emotional expression (Nestor et al., 2011). However, each of these studies relied on identities that were highly familiar to participants due to training or previous experience. Future work should investigate the extent to which image variation (e.g. head rotation, morphing along a dimension, or different expressions) affect representations for untrained identities.

Conclusion

These results begin to elucidate how the face-processing network comes to represent individual faces. Face-specific information is detectable in FFA and aIT, but not OFA, on a trial level. Extensive training of the stimuli before or during scanning is not required to elicit this face-specific information. However, we only detect face-specific information for faces that were subsequently recognized. Thus, these results implicate the development of familiarity as being directly related to the emergence of face-specific information in ventral-stream face regions.

Acknowledgments

We would like to thank Julie Golomb for many helpful discussions and suggestions and Emily Weichart for helpful comments on this manuscript.

References

Andersson, J., Jenkinson, M., Smith, S., 2007. Non-linear Registration, Aka Spatial Normalisation FMRIB Technical Report TR07JA2 (Technical Report No. TR07JA2). FMRIB Analysis Group of the University of Oxford.

Anzellotti, S., Caramazza, A., 2014. The neural mechanisms for the recognition of face identity in humans. Front. Psychol. 5. http://dx.doi.org/10.3389/fpsyg.2014.00672.

Anzellotti, S., Fairhall, S.L., Caramazza, A., 2013. Decoding representations of face identity that are tolerant to rotation. Cereb. Cortex bht046. http://dx.doi.org/10.1093/cercor/bht046.

Bray, J.F., Robbins, D., Witcher, W.B., 1976. Encoding variability theory and the spacing effect in associate learning. Mem. Cogn. 4 (5), 548–552. http://dx.doi.org/10.3758/BF03213217.

Busigny, T., Van Belle, G., Jemel, B., Hosein, A., Joubert, S., Rossion, B., 2014. Face-specific impairment in holistic perception following focal lesion of the right anterior temporal lobe. Neuropsychologia 56, 312–333. http://dx.doi.org/10.1016/j.neuropsychologia.2014.01.018.

Buzsáki, G., Moser, E.I., 2013. Memory, navigation and theta rhythm in the hippocampal–entorhinal system. Nat. Neurosci. 16 (2), 130–138. http://dx.doi.org/10.1038/nn.3304.

Chun, M.M., Turk-Browne, N.B., 2007. Interactions between attention and memory. Curr. Opin. Neurobiol. 17 (2), 177–184. http://dx.doi.org/10.1016/j.conb.2007.03.005.

Collins, J.A., Olson, I.R., 2014. Beyond the FFA: the role of the ventral anterior temporal lobes in face processing. Neuropsychologia 61, 65–79. http://dx.doi.org/10.1016/j.neuropsychologia.2014.06.005.

Cowen, A.S., Chun, M.M., Kuhl, B.A., 2014. Neural portraits of perception: reconstructing face images from evoked brain activity. NeuroImage 94, 12–22. http://dx.doi.org/10.1016/j.neuroimage.2014.03.018.

Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res. 29 (3), 162–173. http://dx.doi.org/10.1006/cbmr.1996.0014.

Cox, R.W., 2012. AFNI: what a long strange trip it's been. NeuroImage 62 (2), 743–747. http://dx.doi.org/10.1016/j.neuroimage.2011.08.056.

Davies-Thompson, J., Gouws, A., Andrews, T.J., 2009. An image-dependent representation of familiar and unfamiliar faces in the human ventral stream. Neuropsychologia 47 (6), 1627–1635. http://dx.doi.org/10.1016/j.neuropsychologia.2009.01.017.

Devlin, J.T., Russell, R.P., Davis, M.H., Price, C.J., Wilson, J., Moss, H.E., Tyler, L.K., 2000. Susceptibility-induced loss of signal: comparing PET and fMRI on a semantic task. NeuroImage 11 (6), 589–600. http://dx.doi.org/10.1006/nimg.2000.0595.

Eger, E., Schweinberger, S.R., Dolan, R.J., Henson, R.N., 2005. Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. NeuroImage 26 (4), 1128–1139. http://dx.doi.org/10.1016/j.neuroimage.2005.03.010.

Eichenbaum, H., 2000. A corticalhippocampal system for declarative memory. Nat. Rev. Neurosci. 1 (1), 41–50. http://dx.doi.org/10.1038/35036213.

Eichenbaum, H., 2004. Hippocampus: cognitive processes and neural representations that underlie declarative memory. Neuron 44 (1), 109–120. http://dx.doi.org/10.1016/j.neuron.2004.08.028.

Epstein, R., Kanwisher, N., 1998. A cortical representation of the local visual environment. Nature 392 (6676), 598–601. http://dx.doi.org/10.1038/33402.

Ernst, M.D., 2004. Permutation methods: a basis for exact inference. Stat. Sci. 19 (4), 676–685. http://dx.doi.org/10.1214/088342304000000396.

Freiwald, W.A., Tsao, D.Y., 2010. Functional compartmentalization and viewpoint generalization within the macaque face-processing system. Science (New York, N.Y.) 330 (6005), 845–851. http://dx.doi.org/10.1126/science.1194908.

Freiwald, W.A., Tsao, D.Y., Livingstone, M.S., 2009. A face feature space in the macaque temporal lobe. Nat. Neurosci. 12 (9), 1187–1196. http://dx.doi.org/10.1038/nn.2363.

Gainotti, G., Marra, C., 2011. Differential contribution of right and left temporo-occipital and anterior temporal lesions to face recognition disorders. Front. Hum. Neurosci. 5, 55. http://dx.doi.org/10.3389/fnhum.2011.00055.

Garoff, R.J., Slotnick, S.D., Schacter, D.L., 2005. The neural origins of specific and general memory: the role of the fusiform cortex. Neuropsychologia 43 (6), 847–859. http://dx.doi.org/10.1016/j.neuropsychologia.2004.09.014.

Gauthier, I., Tarr, M.J., Moylan, J., Skudlarski, P., Gore, J.C., Anderson, A.W., 2000. The fusiform "face area" is part of a network that processes faces at the individual level. J. Cogn. Neurosci. 12 (3), 495–504. http://dx.doi.org/10.1162/089892900562165.

Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. NeuroImage 15 (4), 870–878. http://dx.doi.org/10.1006/nimg.2001.1037.

Gobbini, M.I., Haxby, J.V., 2007. Neural systems for recognition of familiar faces. Neuropsychologia 45 (1), 32–41. http://dx.doi.org/10.1016/j.neuropsychologia.2006.04.015.

Gorno-Tempini, M.L., Price, C.J., 2001. Identification of famous faces and buildings. Brain 124 (10), 2087–2097. http://dx.doi.org/10.1093/brain/124.10.2087.

Grill-Spector, K., Knouf, N., Kanwisher, N., 2004. The fusiform face area subserves face perception, not generic within-category identification. Nat. Neurosci. 7 (5), 555–562. http://dx.doi.org/10.1038/nn122410.1038/nn1224.

Harris, A., Aguirre, G.K., 2010. Neural tuning for face wholes and parts in human fusiform gyrus revealed by fMRI adaptation. J. Neurophysiol. 104 (1), 336–345. http://dx.doi.org/10.1152/jn.00626.2009.

Haxby, J.V., Hoffman, E.A., Gobbini, M.I., 2000. The distributed human neural system for face perception. Trends Cogn. Sci. 4 (6), 223–233. http://dx.doi.org/10.1016/S1364-6613(00)01482-0.

Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 293 (5539), 2425–2430. http://dx.doi.org/10.1126/science.1063736.

Howard, M.W., Kahana, M.J., 2002. A distributed representation of temporal context. J. Math. Psychol. 46 (3), 269–299. http://dx.doi.org/10.1006/jmps.2001.1388.

Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. Med. Image Anal. 5 (2), 143–156.

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. NeuroImage 17 (2), 825–841.

Jenkinson, M., Pechaud, M., Smith, S., 2005. BET2: MR-based Estimation of Brain, Skull and Scalp Surfaces (June).

Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. NeuroImage 62 (2), 782–790. http://dx.doi.org/10.1016/j.neuroimage.2011.09.015.

Kanwisher, N., Yovel, G., 2006. The fusiform face area: a cortical region specialized for the perception of faces. Phil. Trans. R. Soc. B Biol. Sci. 361 (1476), 2109–2128. http://dx.doi.org/10.1098/rstb.2006.1934.

Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J. Neurosci. 17 (11), 4302–4311 (Retrieved from http://www.jneurosci.org/content/17/11/4302).

Kanwisher, N., Tong, F., Nakayama, K., 1998. The effect of face inversion on the human fusiform face area. Cognition 68 (1), B1–B11. http://dx.doi.org/10.1016/S0010-0277(98)00035-3.

Kenward, M.G., Roger, J.H., 1997. Small sample inference for fixed effects from restricted maximum likelihood. Biometrics 53 (3), 983–997. http://dx.doi.org/10.2307/2533558.

Kim, H., 2011. Neural activity that predicts subsequent memory and forgetting: a meta-analysis of 74 fMRI studies. NeuroImage 54 (3), 2446–2461. http://dx.doi.org/10.1016/j.neuroimage.2010.09.045.

Kriegeskorte, N., Formisano, E., Sorger, B., Goebel, R., 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. Proc. Natl. Acad. Sci. 104 (51), 20600–20605. http://dx.doi.org/10.1073/pnas.0705654104.

Lehmann, C., Mueller, T., Federspiel, A., Hubl, D., Schroth, G., Huber, O., Dierks, T., 2004. Dissociation between overt and unconscious face processing in fusiform face area. NeuroImage 21 (1), 75–83. http://dx.doi.org/10.1016/j.neuroimage.2003.08.038.

Liu, J., Harris, A., Kanwisher, N., 2009. Perception of face parts and face configurations: an fMRI study. J. Cogn. Neurosci. 22 (1), 203–211. http://dx.doi.org/10.1162/jocn.2009.21203.

Martin, E., 1968. Stimulus meaningfulness and paired-associate transfer: an encoding variability hypothesis. Psychol. Rev. 75 (5), 421–441. http://dx.doi.org/10.1037/h0026301.

Melton, A.W., 1970. The situation with respect to the spacing of repetitions and memory. J. Verbal Learn. Verbal Behav. 9 (5), 596–606. http://dx.doi.org/10.1016/S0022-5371(70)80107-4.

Minear, M., Park, D.C., 2004. A lifespan database of adult facial stimuli. Behav. Res. Methods Instrum. Comput. 36 (4), 630–633. http://dx.doi.org/10.3758/BF03206543.

Natu, V.S., Jiang, F., Narvekar, A., Keshvari, S., Blanz, V., O'Toole, A.J., 2009. Dissociable neural patterns of facial identity across changes in viewpoint. J. Cogn. Neurosci. 22 (7), 1570–1582. http://dx.doi.org/10.1162/jocn.2009.21312.

Nestor, A., Plaut, D.C., Behrmann, M., 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. Proc. Natl. Acad. Sci. 108 (24), 9998–10003. http://dx.doi.org/10.1073/pnas.1102433108.

Olson, I.R., McCoy, D., Klobusicky, E., Ross, L.A., 2013. Social cognition and the anterior temporal lobes: a review and theoretical framework. Soc. Cogn. Affect. Neurosci. 8 (2), 123–133. http://dx.doi.org/10.1093/scan/nss119.

Olson, I.R., Ezzyat, Y., Plotzker, A., Chatterjee, A., 2014. The end point of the ventral visual stream: face and non-face perceptual deficits following unilateral anterior temporal lobe damage. Neurocase 1–9 http://dx.doi.org/10.1080/13554794.2014.959025.

Pitcher, D., Walsh, V., Yovel, G., Duchaine, B., 2007. TMS evidence for the involvement of the right occipital face area in early face processing. Curr. Biol. 17 (18), 1568–1573. http://dx.doi.org/10.1016/j.cub.2007.07.063.

Pitcher, D., Walsh, V., Duchaine, B., 2011. The role of the occipital face area in the cortical face perception network. Exp. Brain Res. 209 (4), 481–493 (Retrieved from http://www.springerlink.com/index/78482075p4563582.pdf).

Rissman, J., Gazzaley, A., D'Esposito, M., 2004. Measuring functional connectivity during distinct stages of a cognitive task. NeuroImage 23 (2), 752–763. http://dx.doi.org/10.1016/j.neuroimage.2004.06.035.

Ritchey, M., Wing, E.A., LaBar, K.S., Cabeza, R., 2012. Neural similarity between encoding and retrieval is related to memory via hippocampal interactions. Cereb. Cortex bhs258. http://dx.doi.org/10.1093/cercor/bhs258.

Ross, L.A., Olson, I.R., 2010. Social cognition and the anterior temporal lobes. NeuroImage 49 (4), 3452–3462. http://dx.doi.org/10.1016/j.neuroimage.2009.11.012.

Sederberg, P.B., Howard, M.W., Kahana, M.J., 2008. A context-based theory of recency and contiguity in free recall. Psychol. Rev. 115 (4), 893–912. http://dx.doi.org/10.1037/a0013396.

Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage 23 (Suppl. 1), S208–S219. http://dx.doi.org/10.1016/j.neuroimage.2004.07.051.

Smith, T.A., Hasinski, A.E., Sederberg, P.B., 2013. The context repetition effect: predicted events are remembered better, even when they don't happen. J. Exp. Psychol. Gen. 142 (4), 1298–1308. http://dx.doi.org/10.1037/a0034067.

Thomaz, C.E., Giraldi, G.A., 2010. A new ranking method for principal components analysis and its application to face image analysis. Image Vis. Comput. 28 (6), 902–913. http://dx.doi.org/10.1016/j.imavis.2009.11.005.

Tistarelli, M., Bicego, M., Grosso, E., 2009. Dynamic face recognition: from human to machine vision. Image Vis. Comput. 27 (3), 222–232. http://dx.doi.org/10.1016/j.imavis.2007.05.006.

Tsao, D.Y., Freiwald, W.A., Tootell, R.B.H., Livingstone, M.S., 2006. A cortical region consisting entirely of face-selective cells. Science 311 (5761), 670–674. http://dx.doi.org/10.1126/science.1119983.

Turk-Browne, N.B., Simon, M.G., Sederberg, P.B., 2012. Scene representations in parahippocampal cortex depend on temporal context. J. Neurosci. 32 (21), 7202–7207. http://dx.doi.org/10.1523/JNEUROSCI.0942-12.2012.

Turk-Browne, N.B., Golomb, J.D., Chun, M.M., 2013. Complementary attentional components of successful memory encoding. NeuroImage 66, 553–562. http://dx.doi.org/10.1016/j.neuroimage.2012.10.053.

Van Bavel, J.J., Packer, D.J., Cunningham, W.A., 2008. The neural substrates of in-group bias a functional magnetic resonance imaging investigation. Psychol. Sci. 19 (11), 1131–1139. http://dx.doi.org/10.1111/j.1467-9280.2008.02214.x.

Verosky, S.C., Todorov, A., Turk-Browne, N.B., 2013. Representations of individuals in ventral temporal cortex defined by faces and biographies. Neuropsychologia 51 (11), 2100–2108. http://dx.doi.org/10.1016/j.neuropsychologia.2013.07.006.

Visser, M., Jefferies, E., Lambon Ralph, M.A., 2009. Semantic processing in the anterior temporal lobes: a meta-analysis of the functional neuroimaging literature. J. Cogn. Neurosci. 22 (6), 1083–1094. http://dx.doi.org/10.1162/jocn.2009.21309.

Ward, B., 2002. Deconvolution Analysis of fMRI Time Series Data. Biophysics Research Institute, Medical College of Wisconsin, Milwaukee, WI.

Xue, G., Dong, Q., Chen, C., Lu, Z., Mumford, J.A., Poldrack, R.A., 2010. Greater neural pattern similarity across repetitions is associated with better memory. Science 330 (6000), 97–101. http://dx.doi.org/10.1126/science.1193125.