

Modulation of the Fusiform Face Area following Minimal Exposure to Motivationally Relevant Faces: Evidence of In-group Enhancement (Not Out-group Disregard)

Jay J. Van Bavel¹, Dominic J. Packer², and William A. Cunningham³

Abstract

■ Studies have shown that fusiform face area (FFA) activity increases with visual expertise. We present an fMRI study showing that faces from a social category made relevant by an experimental manipulation (members of an experimentally created in-group) preferentially recruited the FFA even when they were matched in exposure to face stimuli from a less significant social category (members of an experimentally created out-group). Faces were randomly assigned to groups and fully counterbalanced so that no perceptual cues allowed participants to visually distinguish category membership. The results revealed a pattern of in-group enhancement (not out-group disregard), such that the FFA was selectively engaged following the presentation of in-group compared with out-group or unaffiliated

control faces even when the intergroup distinction was arbitrary, and exposure to in-group and out-group faces was equivalent and brief. In addition, individual differences in FFA activity for in-group versus out-group faces were correlated with recognition memory differences for in-group and out-group faces. The effects of group membership on the FFA were not affected by task instruction to respond to in-group or out-group members and were functionally dissociated from early visual processing in the primary visual cortex. This study provides evidence that the FFA is sensitive to top-down influences and may be involved in subordinate level (vs. superordinate level) encoding of stimuli in the absence of long-term exposure or explicit task instructions. ■

INTRODUCTION

Functional neuroimaging and lesion studies have identified an area of fusiform gyrus that is involved in facial recognition, labeled the fusiform face area (FFA). This brain region responds preferentially to faces relative to other objects, including scrambled faces, nonface stimuli (e.g., houses), and other body parts (Kanwisher, McDermott, & Chun, 1997; Sergent, Ohta, & MacDonald, 1992). Lesions to this region lead to prosopagnosia, a deficit in face recognition that spares the ability to recognize nonface objects (Ellinwood, 1969; Benton & Van Allen, 1968; De Renzi & Spinnler, 1966). Although research on the FFA has been used to inform theories concerning modular versus distributed models of neural and cognitive processing (Kanwisher & Yovel, 2006), there remains considerable debate about whether this brain region involves a specialized mechanism for recognizing the presence and identity of faces (Kanwisher et al., 1997; Sergent et al., 1992) or is engaged in more general cognitive processes (Grill-Spector, Sayres, & Ress, 2006).

Several recent neuroimaging studies have indicated that FFA activity increases with expertise. This expertise

may be for certain categories of faces, such as members of one's own race (Golby, Gabrieli, Chiao, & Eberhardt, 2001), or nonface stimuli (Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999). For example, car and bird experts have heightened FFA activity while viewing cars and birds, respectively (Gauthier, Skudlarski, Gore, & Anderson, 2000). Further, Gauthier and colleagues experimentally created expertise by training participants to identify specific novel nonface stimuli called "greebles" (which had a small number of parts in a common configuration that allowed for configural processing) and found greater FFA activity during the passive viewing of greebles among trained greeble experts versus nonexperts (Gauthier et al., 1999). These studies led Gauthier and Tarr to suggest that the FFA is better labeled the *Flexible Fusiform Area*, because processing in the region is not limited to predetermined content, such as faces. However, the flexibility of this region may be constrained by the extensive exposure required to develop visual expertise—greeble expertise involved over 3000 trials of intensive training over several days or weeks (Gauthier et al., 1999). These studies suggest that extensive visual experience with faces or other stimulus categories may gradually tune neurons in the FFA to encode stimuli at the subordinate/individual level, that is, to make fine-grained discriminations between

¹New York University, ²Lehigh University, ³The Ohio State University

exemplars within a stimulus category (Tarr & Gauthier, 2000).

Consistent with this perspective, a recent study examined FFA activity when participants were presented with own-race versus other-race faces (Golby et al., 2001). A large body of research has shown that people from a variety of ethnic and racial backgrounds exhibit superior recognition memory for own-race faces and believe that other-race faces look alike—a highly robust effect termed own-race bias (Meissner & Brigham, 2001; Malpass & Kravitz, 1969). Building on this research, Golby et al. (2001) presented Black and White participants with pictures of Black and White faces as well as objects (radios) during neuroimaging. Brain activity to the faces was first contrasted with activity to objects to functionally identify the FFA (Kanwisher et al., 1997). Activity in the FFA was then compared for own-race and other-race faces. As predicted, activity in FFA was greater to own-race than other-race faces for both Black and White participants (see also Lieberman, Hariri, Jarcho, Eisenberger, & Bookheimer, 2005). Moreover, participants with the strongest own-race bias on a subsequent recognition memory test (i.e., superior memory for own-race compared with other-race faces) had the strongest own-race bias in activity in the fusiform gyrus (i.e., greater activity for own-race compared with other-race faces). According to the expertise-based account of FFA function, years of experience with own-race faces may have tuned the FFA to identify own-race faces on the basis of subtle differences in their physiognomic features.

In the current study, we investigate whether activity in the FFA is necessarily contingent on expertise with specific categories or whether the FFA is also sensitive to transient social factors in a top-down fashion. Where previous research has examined the effects of race on FFA activity, we created experimental groups to examine the effects of social group membership on FFA activity, independent of race, in the absence of differential experience with in-group and out-group members and without any visual cues that signify group membership. Studies have now shown that long-term expertise may not be necessary to selectively engage the FFA, and indicate that activity in FFA may have as much to do with the way in which visual stimuli are processed as the class to which they belong (Gauthier, Anderson, Tarr, Skudlarski, & Gore, 1997). For example, participants who completed a task in which they matched nonface stimuli with superordinate categorical (e.g., bird) versus subordinate-level (e.g., pelican) descriptors had greater activity in the ventral visual pathway, including the FFA, during subordinate level judgments (Gauthier, Tarr, et al., 2000; Gauthier et al., 1997). In the current study, we examined whether a social context in which a subset of faces were made motivationally relevant on the basis of their membership in a particular social category (i.e., in-group members) would heighten FFA activity in the absence of prior expertise and explicit task instructions. This hypothesis drew upon

research from the person perception literature in social psychology suggesting that sharing a group membership changes the manner in which members of different social categories are perceived and evaluated.

Models of person perception propose that motivational factors can lead people to move beyond category information and encode others as individuals (Fiske & Neuberg, 1990; Brewer, 1988). In a similar fashion, several influential models of face perception propose that the own-race memory bias stems from a shared racial identity rather than longer-term experience with own-race faces (Hugenberg, Young, Bernstein, & Sacco, 2010; Sporer, 2001). For example, the in-group/out-group model of face processing (Sporer, 2001) argues that all exemplar faces are processed in the same fashion as in-group faces (termed the “default route”), unless the perceiver detects an out-group cue, in which case the face is disregarded. According to this model, in-group faces are encoded as individuals, which increases subsequent recognition for them, whereas out-group faces are encoded at a categorical level, which leads to stereotyping, cognitive disregard and reduces subsequent recognition memory. This pattern of out-group disregard is, thus, predicted to produce the own-race bias and perceptions of out-group homogeneity (Quattrone & Jones, 1980).

Person perception is complicated by the fact that people identify with numerous groups and identification can be highly dynamic, changing to accommodate current social contexts (Turner, Oakes, Haslam, & McGarty, 1994; Turner, Hogg, Oakes, Reicher, & Wetherell, 1987). Indeed, merely assigning people to groups in the absence of factors typically posited to account for intergroup bias (e.g., stereotypes, prior conflict, or competition for resources) is sufficient to evoke identification with fellow in-group members (Tajfel, 1970). Assigning people to a minimal group also increases individuation: participants show superior recognition memory for in-group versus out-group faces, even when group membership is random, recent, and participants have equal exposure to in-group and out-group faces (Bernstein, Young, & Hugenberg, 2007). Thus, the simple categorization of others as “us” or “them” may be sufficient to generate biases in subordinate versus superordinate perception.

In a previous study, we examined the influence of group membership on the neural substrates involved in person perception. We used a variant of the minimal-group paradigm in which White participants were randomly assigned to one of novel, two mixed-race teams—called the Leopards and the Tigers (Van Bavel, Packer, & Cunningham, 2008). Assigning participants to teams allowed us to examine whether group membership would override typical racial biases in social perception. After memorizing the faces of in-group and out-group members, participants then categorized the same faces during fMRI according to their race or team membership. On half the trials, participants categorized the faces according to race (Black and White); on the other half of the trials, participants categorized the

faces according to team membership (Leopards or Tigers). Although the primary purpose of this study was to extend previous research on the relationship between social categories (e.g., race) and amygdala activity (Cunningham et al., 2004), the research design allowed us to make a preliminary examination of the role of minimal group membership on fusiform activity. On the basis of previous research, we hypothesized that in-group faces would be associated with greater activity in the fusiform gyri than out-group faces. However, it was unclear whether fusiform activity would increase only to the more familiar own-race faces (Lieberman et al., 2005; Golby et al., 2001) or to the minimal in-group members (Bernstein et al., 2007). Consistent with the models of person perception described above (Fiske & Neuberg, 1990; Brewer, 1988), we found greater activity in the left and right fusiform gyri to minimal in-group than out-group members, regardless of their race. Thus, although earlier studies found greater fusiform activity to own-race faces (Golby et al., 2001), our study raised the possibility that sharing a social identity may have contributed to those findings—overriding the effects of familiarity. Indeed, our previous study revealed no main effect of race on fusiform activity and the effect of group membership on fusiform activity was not moderated by race. Moreover, the pattern of in-group bias in fusiform activity was not moderated by the categorization task, suggesting that it did not require explicit attention to team membership.

Our previous study provided suggestive evidence that the fusiform gyri may process in-group members differently than out-group members, regardless of their race, placing this bias within the realm of motivated social perception (Balcetis & Dunning, 2006; Bruner, 1957). However, as noted above, any conclusions about the top-down influence of group membership on the FFA require a functional FFA localizer rather than the whole-brain analysis used in our previous study. The FFA is a relatively small area of the fusiform gyri, a large, functionally heterogeneous region of the occipito-temporal cortex, much of which does not necessarily respond preferentially to faces (Kanwisher et al., 1997; Puce, Allison, Gore, & McCarthy, 1995). The FFA is generally in a slightly different area of the fusiform gyri for each participant, which places a limit on any conclusions drawn from a whole-brain analysis; that is, we have no way of knowing whether the previously observed in-group bias was in the FFA proper or in another region of the fusiform gyri. Therefore, the only way to conclusively examine the top-down influence of social identity on person perception in the FFA is to employ a functional FFA localizer.

Overview and Objectives

We used fMRI to explore whether group-based motivational relevance would modulate FFA activity, even when the intergroup distinction was arbitrary and exposure to in-group and out-group faces was equivalent, brief, and

very recent. Following our previous research (Van Bavel et al., 2008), we predicted that in-group faces would elicit greater FFA activity (in-group > out-group), despite participants' limited exposure to members of both categories (and no differential exposure between categories). We created groups through an experimental procedure rather than using an existing intergroup distinction (e.g., race) to examine the FFA in the absence of differential exposure to in-group/out-group members or visual cues that signify group membership (Van Bavel et al., 2008). We randomly assigned participants to one of two arbitrary groups and gave them approximately 15 min to learn the faces in both groups. Critically, faces were randomly assigned to the teams, and assignment was fully counter-balanced so that no perceptual cues allowed participants to visually sort the faces into teams: only the experimental manipulation of group membership could account for differences between in-group and out-group faces.

During fMRI, we presented participants with in-group and out-group faces, as well as faces of individuals who were unaffiliated with the in-group or out-group. As we noted above, the influential in-group/out-group model of face processing (Sporer, 2001) argues that exemplar faces are processed by the “default route,” unless the perceiver detects an out-group cue, in which case the face is processed as an out-group face and, therefore, disregarded. Including unaffiliated control faces was an important extension of previous research because it allowed us to examine whether any relative differences between in-group and out-group faces were due to decreased FFA activity to out-group members (out-group disregard) and/or increased FFA activity to in-group members (in-group enhancement), relative to the control faces. According to the in-group/out-group model (Sporer, 2001), in-group and control faces should be processed in the default route which leads to deeper structural encoding and, therefore, greater FFA activity than out-group faces. However, other models of person perception argue that targets are processed at the categorical level unless the target is motivationally relevant, which leads to deeper encoding and individuation (Fiske & Neuberg, 1990; Brewer, 1988). Including unaffiliated control faces allowed us to directly compare these theories in the context of face perception. Moreover, to our knowledge, no studies comparing in-group versus out-group differences in neural activity have included a similar control group, making it difficult to determine whether relative differences stem from in-group or out-group biases.

Research on the role of the FFA in own-race bias by Golby et al. (2001) has shown a direct correlation between fusiform activity and recognition memory, such that participants with the largest difference in recognition memory for own-race compared with other-race faces also had the largest difference in fusiform activity to own-race compared with other-race faces. Given that the FFA appears to play an important role in subordinate level recognition (Gauthier, Tarr, et al., 2000), this finding may reflect greater

individuation of in-group faces, resulting in heightened memory. Importantly, own-race faces may have been encoded as individuals because of their motivational relevance to perceivers (rather than a long history of exposure), whereas less-relevant other-race faces may be encoded on the basis of their category membership (Hugenberg et al., 2010; Sporer, 2001). We predicted that increased FFA activity to in-group compared with out-group members in the current study would similarly reflect greater subordinate level recognition of these more motivationally relevant faces (Bernstein et al., 2007) and would, thus, correlate with greater recognition memory for in-group relative to out-group faces assessed after scanning during a surprise recognition memory task. This pattern of results would introduce the possibility that the effects of race on the FFA reported by Golby et al. (2001) might be mediated by psychological variables (e.g., motivational relevance) associated with group membership, in addition to expertise.

Tasks that involve subordinate level judgments (Gauthier, Tarr, et al., 2000; Gauthier et al., 1997) or covert attention (Wojciulik, Kanwisher, & Driver, 1998) have been shown to increase FFA activity. We, therefore, varied explicit attention to the identity of in-group or out-group members to investigate whether attention to group membership modulated FFA activity. Importantly, the presence of the unaffiliated control faces ensured that participants could not simply attend to only the in-group (or out-group) identity of the faces when performing the tasks. Thus, the task required visual attention to all the faces but differentially oriented participants to the identity of different faces in different conditions. We predicted that attention to the identity of out-group faces might attenuate any in-group biases in neural activity.

In addition, we extracted a functionally defined region of the occipital lobe sensitive to faces—termed the occipital face area (OFA)—and an anatomically defined region of the primary visual cortex (calcarine sulcus) to assess whether the effects of group membership and/or task on FFA activity were associated with other regions of the face-processing network or regions involved in earlier visual perception (Posner & Gilbert, 1999), respectively. The OFA is thought to form an initial perceptual representation of a face and to project this representation to the FFA and other face-processing regions, such as the STS (Haxby, Hoffman, & Gobbini, 2000). On the other hand, the calcarine sulcus is involved in very low-level aspects of processing visual information. We included the calcarine sulcus in the present analysis to determine whether any effects in the FFA were mediated by more low-level perceptual or attentional processes in the primary visual cortex. If the effect of group membership on the FFA is merely due to low-level attentional processes (e.g., simply looking more at in-group faces), we should find the same pattern of results in the FFA and the calcarine sulcus; if not, we should find a dissociation between activity in the FFA and the calcarine sulcus. Such

a dissociation would provide discriminant evidence that group membership can exert a top-down influence on the encoding of faces over-and-above low-level attentional biases.

METHODS

Participants

Nineteen White participants (11 women, mean age = 20.1 years) were recruited from the community at Queen's University and paid \$40 for completing the study. Participants reported no history of neurological problems, were native English speakers, and had normal or corrected-to-normal vision. One participant failed a manipulation check at the end of the experiment, and a computer error disrupted the task during neuroimaging for another participant. The data from these participants were omitted from analysis, leaving 17 participants. Each participant gave written informed consent to participate in the study.

Procedure

Group Assignment

Participants arrived at the neuroimaging center and were taken to a behavioral testing room. They were informed that they were in a study investigating learning about groups and that they had been assigned to one of two teams: the Leopards or the Tigers (Van Bavel & Cunningham, 2009; Van Bavel et al., 2008). Participants were told that it was important for them to learn the members of both teams. Participants then completed two learning tasks lasting approximately 15 min on a Dell laptop. During the first learning task, 16 male faces divided into two teams of eight (Leopards and Tigers) appeared simultaneously on the computer screen and participants spent 3 min memorizing the faces associated with each team. Face stimuli were color images created in Photoshop and presented as 2 × 2.5 in. at 72 pixels/in. (all shot from the same angle and roughly the same position). All faces presented in the study had a neutral expression. Critically, faces were randomly assigned to the teams, and assignment was fully counter-balanced so that no perceptual cues allowed participants to visually sort the faces into teams. This design ensured that participants were equally likely to see each face as an in-group or out-group member. Participants did not interact with actual members of either team nor were they informed that they would interact with members of either team. Their only exposure to team members was when they saw the faces on the computer.

The second learning task contained two blocks; in both blocks, the faces were shown one at a time, and participants categorized each according to whether it was affiliated with the Leopards or the Tigers. To ensure that participants identified with their team, each participant also categorized a digital photograph of his or her own face as part of this task. The participant's own face was

randomly interspersed at three points within each block of faces during this learning task. This learning task was the only part of the study in which each participant viewed his or her own face: Participants did not see their face during neuroimaging or the postneuroimaging recognition memory task. During the first block of learning trials, a label reminded participants whether each face was a Leopard or Tiger. Participants categorized each of the 16 in-group and out-group faces once and his or her own face three times during the first block, for a total of 19 trials. During the second block of trials, the team label was removed so that participants needed to rely on their memory to categorize the faces. Following each trial in the second block, feedback indicated whether the response was correct and listed the correct team affiliation for each face. Participants categorized each in-group and out-group face three times and his or her own face three times during the second block, for a total of 51 trials. Importantly, participants were run in individual sessions and never interacted with in-group or out-group members at any time (nor were they led to believe that they would interact with group members).

Following group assignment and learning, participants completed a Face Categorization Task and Fusiform Face Area Localizer Task in a Siemens 3T Tim Trio scanner. All stimuli presented during fMRI were back projected from an LCD projector to a clear screen at the back of the scanner bore. Participants were able to see these stimuli using

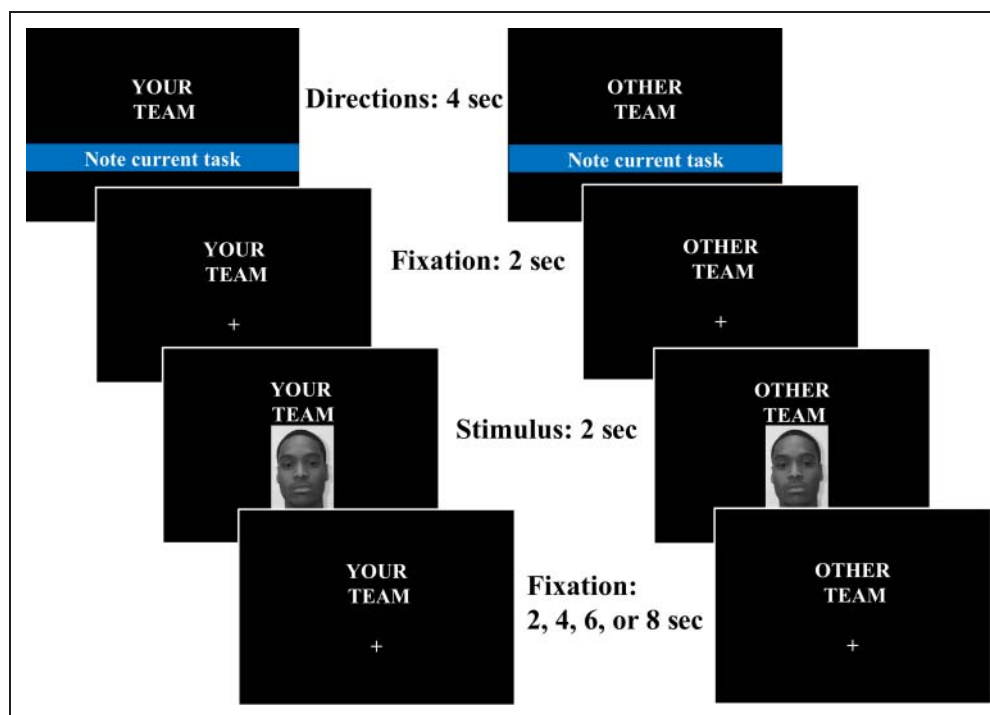
a mirror mounted on top of the head coil (the visual angle of the stimuli was approximately $8^\circ \times 6^\circ$). Stimuli and directions were presented in the center of an otherwise black screen.

Face Categorization Task

During fMRI, participants completed a mixed block/event-related design involving five runs of four blocks of 12 trials for a total of 240 trials of the Face Categorization Task. During the task, participants were presented with 24 faces: eight in-group faces, eight out-group faces, and eight novel faces of individuals who were unaffiliated with the in-group or out-group. Stimuli were presented one at a time in the center of an otherwise black screen. Participants saw the unaffiliated faces for the first time during fMRI scanning. Faces were racially diverse such that half of the faces were White and half were Black (i.e., race was orthogonal to team membership). Replicating our previous research (Van Bavel et al., 2008), there was no main effect of race on FFA activity [$F_{(1, 14)} = 1.11, p = .31$] and this variable is not discussed any further for the sake of brevity.

On each trial, participants identified one of the 24 faces in one of two ways (see Figure 1). During in-group categorization blocks, participants pressed a button only if the face was an in-group member. During out-group categorization blocks, participants pressed a button only if the face was an out-group member. Direction screens

Figure 1. Sample trials in the in-group categorization block (left) and out-group categorization block (right) during fMRI. Each block started with a directions screen (the top screen in the figure). After the directions screen, participants completed 12 trials. On each trial, participants hit a button if a randomly presented face (the third screen in the figure) was an in-group member (the left screens in the figure) or out-group member (the right screens in the figure) and then saw a fixation cross (the bottom screen in the figure). Each face appeared for 2 sec, during which time participants responded with a button box in their right hand. To allow for estimation of the hemodynamic signal, fixation crosses appeared between names for 2, 4, or 6 sec (in pseudorandom order). After the completion of each block, directions for the next block appeared. Each of five runs contained two in-group categorization blocks and two out-group categorization blocks (counterbalanced). The face in the figure was not used in the actual study and is reproduced with permission from Minear & Park (2004). (Note: figures are not shown to scale.)



were presented for 4 sec before each block of 12 trials. In-group and out-group blocks were counterbalanced within runs, creating four randomized blocks within each run. Each of the 24 faces was categorized twice in each run (once in the in-group block and once in the out-group block). Each face appeared for 2 sec, during which time participants responded with a button box in their right hand. To allow for modeling of the hemodynamic signal, fixation crosses appeared between names for 2, 4, or 6 sec (in pseudorandom order). Faces were presented in random order within each run.

Fusiform Face Area Localizer Task

Following the Face Categorization Task, participants completed a one-run Fusiform Face Area Localizer Task consisting of four blocks of 12 trials for a total of 48 trials. During the task, participants were presented with 24 face and 24 nonface stimuli (e.g., cars, houses). Stimuli were presented one at a time in the center of an otherwise black screen in random order. On each trial, participants were instructed to press a button when each stimulus appeared on the screen. Direction screens were presented for 4 sec before each block of 12 trials. Each of the 24 face and nonface stimuli was presented in the localizer task. Each face appeared for 2 sec, during which time participants responded with a button box in their right hand. To allow for estimation of the hemodynamic signal, fixation crosses appeared between names for 2, 4, or 6 sec (in pseudorandom order). Faces were presented in random order within the localizer task.

Face Recognition Memory Task

After neuroimaging, participants completed a behavioral Face Recognition Memory Task on a Dell laptop. During the task, participants saw each of the 24 faces they had seen during neuroimaging (eight in-group faces, eight out-group faces, and eight faces of individuals who were unaffiliated with the in-group or out-group). The faces were presented one-at-a-time in the center of the screen in random order and participants were instructed to indicate with a button press whether each face was a member of the (a) Leopards, (b) Tigers, or (c) neither team. Response accuracy was recorded for each face.

Neuroimaging Parameters, Acquisition, and Analysis

Functional scanning was prescribed parallel to the AC–PC line, and nearly isotropic functional images were acquired from inferior to superior using a single-shot gradient-echo planar pulse sequence (32 axial slices, 3.5-mm thick, 0.5-mm skip, echo time = 25 msec, repetition time = 2000 msec, in-plane resolution = 3.5×3.5 mm, matrix size = 64×64 , field of view = 224 mm). Data were pre-


processed and analyzed in SPM8 (Wellcome Department of Cognitive Neurology, London, United Kingdom). Data were corrected for slice acquisition time and motion, transformed to conform to the default EPI Montreal Neurological Institute brain, smoothed using a 9-mm FWHM kernel, corrected for artifacts and detrended. BOLD signal was modeled as a function of a canonical hemodynamic response function and its temporal derivative with a 128-sec high-pass filter.

We extracted ROIs for the FFA, OFA, and the calcarine sulcus using the Marsbar ROI Toolbox for SPM8. FFA ROIs were defined individually for each participant as the voxel clusters in or partially overlapping the left or right fusiform gyrus that were more active while viewing faces (both in-group and out-group faces) compared with nonface objects in the FFA localizer task. Voxel clusters that were more active to faces than objects ($p < .01$, uncorrected) in the right or left fusiform gyrus (defined by MRIcron; Rorden, Karnath, & Bonilha, 2007) during the FFA localizer task were extracted as ROI masks for each participant. We identified an FFA in 15 of 17 participants by this criterion (14 participants had FFA ROIs in the right hemisphere and six participants had FFA ROIs in the left hemisphere). FFA ROIs were combined in participants who had bilateral FFA activations that exceeded the statistical threshold. Liberal thresholds are often used for defining well-known functional ROIs, including the FFA (Rhodes, Michie, Hughes, & Byatt, 2009; Maurer et al., 2007; Anderson, Christoff, Panitz, De Rosa, & Gabrieli, 2003; Gauthier, Tarr, et al., 2000). However, we also report analyses with more conservative FFA ROI thresholds ($p < .001$, $p < .0001$, uncorrected; see Figure 2 for sample ROIs) in Results to examine the generality of the main effect(s). See supplement for information about OFA ROIs. The mean activity was averaged across all contiguous voxels in each ROI and compared across conditions for each participant.

ROIs of the primary visual cortex were defined for each participant using an anatomical mask of the calcarine sulcus available in the Marsbar ROI Toolbox for SPM8 (Brett, Anton, Valabregue, & Poline, 2002) based on anatomical ROIs (Tzourio-Mazoyer et al., 2002). Specifically, the voxels surrounding the left and right calcarine fissure and its branches were extracted and averaged for each participant. ROIs of the calcarine sulcus were structurally independent from the OFA and FFA ROIs.

ROI masks for each participant were used to extract signal change within the FFA, OFA and calcarine sulcus for the first five functional runs. First level analyses were implemented in SPM8. In each analysis, a series of regressors was constructed to examine BOLD brain activity for each of the trial types. The BOLD signal was modeled as a function of a canonical hemodynamic response plus the temporal derivative. For each ROI, BOLD signal was computed across all voxels within the ROI at each repetition time for each trial type. Mean BOLD signal was exported to SAS for each ROI and analyzed using a 3 (Group: in-group,

Figure 2. Activation maps from six participants who had functionally defined FFA at the statistical threshold of $p < .0001$. Voxels more active while viewing faces compared with objects that reached a statistical threshold were defined as the face-responsive ROIs. Signal change within these ROIs was extracted for all conditions during the first five functional runs and analyzed in SAS. MNI coordinates in parentheses (X, Y, Z).



out-group, unaffiliated) \times 2 (Task: in-group, out-group) repeated measures ANOVA. The mixed block/event-related design allowed us to examine the phasic event-related effect of group membership and the interaction of phasic events in the context of different tonic influences (i.e., tasks). A statistical threshold of $p < .05$ was used to analyze extracted ROI data.

RESULTS

Neuroimaging Data

Fusiform Face Area

We examined whether presentations of in-group and out-group faces would modulate activity in the FFA relative to unaffiliated faces. Specifically, we were interested in whether these biases would be characterized by *in-group enhancement* (in-group $>$ unaffiliated $=$ out-group) or *out-group disregard* (in-group $=$ unaffiliated $>$ out-group). The FFA for each participant was extracted and interrogated for differential responses to in-group, out-group, and unaffiliated faces. As predicted, group membership modulated mean BOLD activity within the functionally defined FFA [$F_{(2, 13)} = 6.20, p < .01$; see Figure 3A). To identify whether this effect was driven by in-group enhancement or out-group disregard, we compared in-group and out-group faces to unaffiliated faces. A series of planned

comparisons revealed greater FFA activity to in-group than out-group [$t_{(14)} = 2.24, p = .04$] and unaffiliated [$t_{(14)} = 3.65, p < .01$] faces and marginally greater FFA activity to out-group than unaffiliated faces [$t_{(14)} = 1.99, p < .07$]. Analyses using more stringent thresholds for defining the FFA (Golby et al., 2001) also indicated greater activity for in-group versus out-group faces at $p < .001$ [$p = .07$] and $p < .0001$ [$p = .04$]. These results indicate that activity in the FFA was heightened following the presentation of in-group faces compared with out-group or unaffiliated faces; activity in FFA was not, however, decreased to out-group compared with unaffiliated faces. In previous research, tasks that increase subordinate level judgments (Gauthier, Tarr, et al., 2000; Gauthier et al., 1997) or covert attention (Wojciulik et al., 1998) have sometimes been shown to increase FFA activity. However, the effect of group membership on FFA activity was not moderated by task [$F_{(2, 13)} = 2.62, p = .11$]. These results indicate that explicit attention to in-group versus out-group membership did not moderate FFA activity. OFA ROIs followed the same pattern of results (see Supplementary Data).

Primary Visual Cortex

The primary visual cortex (operationalized here as the calcarine sulcus) is involved in relatively low level aspects of visual information processing and was analyzed to

determine whether the patterns of in-group enhancement reported above were mediated by more basic perceptual processes in the primary visual cortex. Alternatively, if the FFA/OFA were functionally dissociated from the calcarine sulcus, we would expect a different pattern of effects in both sets of regions. For example, the more novel, unaffiliated faces—seen for the first time during the neuroimaging task—may have captured attention and led to greater activity in the calcarine sulcus. An anatomically defined region of the calcarine sulcus for each participant was extracted and interrogated for differential responses to in-group, out-group, and unaffiliated faces. Consistent with the pattern of activity in FFA and OFA, there was greater mean BOLD activity within the anatomically defined region of the calcarine sulcus for in-group than out-group faces [$F_{(2, 13)} = 6.20, p < .01$]. To identify whether this

effect was similar to the results in FFA/OFA, we compared in-group and out-group faces to unaffiliated faces (see Figure 3B). Similar to the results in FFA and OFA, there was greater activity in the calcarine sulcus to in-group than out-group faces [$t_{(14)} = 5.31, p < .01$]. However, in contrast to the pattern of activity in the FFA and OFA, mean activity in the calcarine sulcus to unaffiliated faces was almost identical to in-group faces [$t_{(14)} = 0.02, p = .99$] faces and was significantly greater than activity to out-group faces [$t_{(14)} = 6.20, p < .01$]. Unlike the results in FFA and OFA, in-group, and unaffiliated faces evoked greater activity in the calcarine sulcus compared with out-group faces (see below for a formal test directly comparing these regions). This pattern of results suggests that the primary visual cortex might be better characterized by out-group disregard.

The dissociation in activity between face processing regions (FFA/OFA) and primary visual cortex (the calcarine sulcus) raises the possibility that attention influenced early visual processing in this task. Consistent with the notion that attention played a role, the effect of group membership on calcarine sulcus activity was moderated by task, which varied explicit attention to group identity [$F_{(2, 13)} = 11.12, p < .01$]. To decompose this effect, we compared the effects of group membership separately during the in-group and out-group tasks. A series of planned comparisons indicated greater calcarine sulcus activity to in-group than out-group faces during the in-group attention task [$t_{(14)} = 6.38, p < .01$] but not during the out-group attention task [$t_{(14)} = 0.36, p > .72$; see Figure 3]. These results indicate that explicit attention to in-group versus out-group identity moderated activity in the calcarine sulcus and provide additional evidence that basic perceptual and attentional processes in the primary visual cortex were functionally dissociated from processes in the face processing regions.

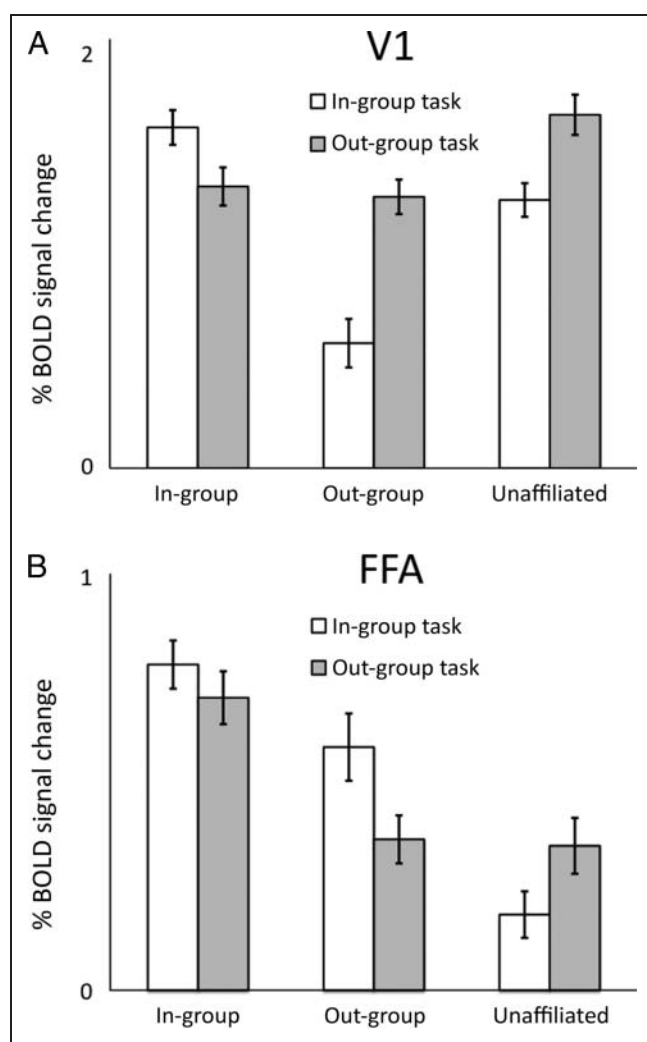


Figure 3. The effect of Group (in-group, out-group, unaffiliated) and Task (in-group, out-group attention) on the functionally defined FFA and the anatomically defined calcarine sulcus. (A) Mean BOLD signal (in percent signal change) within the functionally defined FFA at $p < .01$. (B) Mean BOLD signal (in percent signal change) within the calcarine sulcus (V1; $n = 15$). Error bars reflect within-subject standard errors.

Comparing the Fusiform Face Area to the Primary Visual Cortex

These analyses indicate that the calcarine sulcus is more sensitive to the relatively novel, unaffiliated faces and effects of our task than the FFA and OFA, suggesting a dissociation between early visual perception and face processing regions. We conducted a 3 (Group: in-group, out-group, unaffiliated) \times 2 (Task: in-group, out-group) \times 2 (Region: FFA, calcarine sulcus) repeated measures ANOVA to directly assess whether the effects of group membership and task on FFA activity were dissociated from earlier visual perception regions. Consistent with the results reported above, an interaction between group and brain region provides statistical evidence that activity in FFA versus the calcarine sulcus is dissociated in terms of processing group membership [$F_{(2, 13)} = 8.12, p < .01$]. The calcarine sulcus responded more strongly to the relatively novel unaffiliated faces than did the FFA. There was also a three-way interaction between Group, Task, and Brain region [$F_{(2, 13)} = 7.72,$

$p < .01$], indicating that the effects of Task on processing in-group versus out-group members was stronger in the calcarine sulcus compared with FFA. Taken together, these results suggest a clear dissociation between activity in the FFA and this area of the primary visual cortex. BOLD activity in FFA (and OFA) was characterized by in-group enhancement (in-group > unaffiliated = out-group) whereas BOLD activity in the calcarine sulcus was characterized by out-group disregard (in-group = unaffiliated > out-group). Further, the calcarine sulcus was more sensitive to the effects of explicit attention to in-group versus out-group members.

Memory Data

Faces encoded at the subordinate level should be correctly recognized more accurately than faces encoded at the superordinate level (Sporer, 2001). Indeed, recent research showing that group membership leads to superior recognition memory for in-group faces (Bernstein et al., 2007) led us to hypothesize that participants would have superior recognition memory for in-group compared with out-group faces. To examine this hypothesis, we analyzed response accuracy (i.e., hit rates) on the recognition memory task that participants completed after scanning using a 3 (Group: in-group, out-group, unaffiliated) one-way ANOVA.¹ The analysis indicated that participants had superior recognition memory for more familiar in-group ($M = 0.76 \pm 0.03$) and out-group ($M = 0.68 \pm 0.04$) faces relative to the less familiar unaffiliated faces ($M = 0.48 \pm 0.04$ [$F_{(2, 13)} = 10.36, p < .01$]). A planned comparison between in-group and out-group faces indicated that participants had superior recognition memory for in-group compared with out-group faces [$t_{(14)} = 1.23, p < .12$, one-tailed]. Although this effect was in the right direction and moderately large ($d = .67$), it did not reach statistical significance. However, it has been replicated several times with larger samples (see Hugenberg et al., 2010, for a review). These results are consistent with the idea that participants successfully encoded the identity of individuals during the learning phase and were slightly more likely to encode in-group than out-group faces in the absence of instructions to attend preferentially to either group.

Brain–Behavior Correlations

More importantly, we investigated whether there was a correlation between the recognition memory differences for in-group and out-group faces and the brain activation differences in the FFA for in-group and out-group faces. Individual differences in mean recognition memory for in-group versus out-group faces were calculated and correlated with individual differences in mean FFA activity for in-group and out-group faces for the 15 participants with behavioral data and FFA ROIs. As shown in Figure 4, we found a marginally significant correlation between in-group biases in recognition memory and FFA activity [$r =$

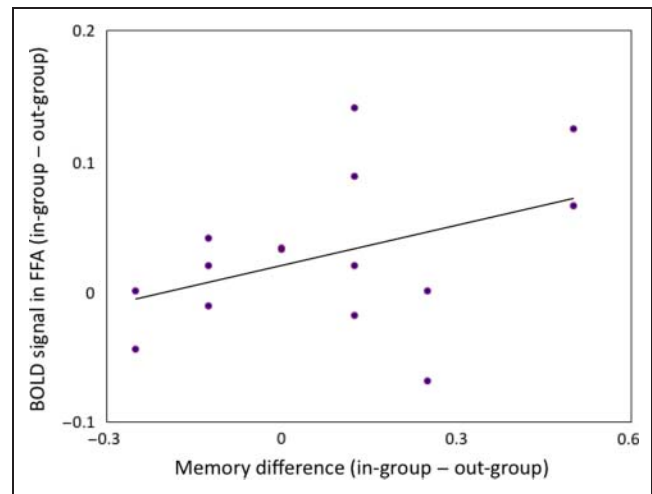


Figure 4. Correlation between memory differences for in-group versus out-group faces with mean BOLD signal in the FFA to in-group versus out-group faces ($r = .49$). The scatterplot illustrates the association between relative signal change in the functionally defined FFA for in-group minus out-group faces during five functional runs and differences in mean accuracy for in-group minus out-group faces during the recognition memory task administered after neuroimaging ($n = 15$).

$0.49, p = .06$]. These results indicate that participants with greater in-group bias (in-group > out-group) in recognition memory had greater in-group bias (in-group > out-group) in FFA activity. More generally, these results suggest that the FFA may be associated with encoding the more motivationally relevant in-group faces at the subordinate level relative to out-group faces.

DISCUSSION

The current study provides the strongest evidence to date that the FFA is sensitive to shifts in social context, responding selectively to face stimuli that are imbued with psychological significance by virtue of their group membership. Specifically, the functionally defined FFA was selectively engaged following the presentation of in-group versus out-group faces even when the intergroup distinction was arbitrary, there were no visual cues to distinguish group membership, and exposure to in-group and out-group faces was equivalent, brief (~15 min), and very recent. Indeed, our group assignment manipulation ensured that no perceptual cues allowed participants to visually sort the faces into teams. Replicating our previous research (Van Bavel et al., 2008), there was no effect of race—the social category presumably associated with greater visual expertise—on FFA activity. Thus, only the experimental manipulation of group membership could account for the difference in FFA activity between in-group and out-group faces. To identify whether this effect was driven by heightened activity to in-group members or decreased activity to out-group members, we compared

in-group and out-group faces to unaffiliated faces shown to participants for the very first time during neuroimaging. The results revealed a pattern of in-group enhancement: FFA activity was heightened following the presentation of in-group faces compared with out-group and unaffiliated faces. If anything, FFA activity to out-group faces was slightly greater than to unaffiliated faces.

This research extends our previous research on the role of group membership on the neural substrates of in-group bias in several important ways (Van Bavel et al., 2008). First, including a functional FFA localizer increases our confidence that group membership modulates the FFA, as opposed to other regions of the fusiform gyri. Second, the positive correlation between in-group bias in FFA activity and in-group bias in recognition memory is consistent with the idea that the FFA may play a role in the subordinate level encoding of in-group members. Third, including a third set of faces that were unaffiliated with the in-group and out-group supports our assertion that the difference in FFA activity between in-group and out-group faces could be characterized as in-group enhancement (i.e., enhanced activity for the more motivationally relevant in-group members). Taken together, these results increase our confidence that group membership has a top-down influence on the FFA and may be enhancing the subordinate level encoding of minimal in-group members.

As we noted in the introduction, several models of person perception propose that motivational factors can lead people to move beyond category information and encode others as individuals (Fiske & Neuberg, 1990; Brewer, 1988), and some researchers have proposed that people engage in deeper encoding of in-group compared with out-group members (Hugenberg et al., 2010; Sporer, 2001). These models differ, however, in what they posit as the default mode of person perception. The in-group/out-group model of face processing (Sporer, 2001) argues that exemplar faces are processed by the “default route,” unless the perceiver detects an out-group cue, in which case the face is processed as an out-group face and is, therefore, disregarded. Other models argue that targets are processed at the categorical level unless the target is motivationally relevant (e.g., in-group members), which leads to deeper encoding and individuation (Fiske & Neuberg, 1990; Brewer, 1988). By using unaffiliated control faces to examine the ostensible default mode of person perception (relative to in-group and out-group faces), we found that the differences between in-group and out-group faces were due to increased FFA activity to in-group members (not decreased FFA activity to out-group members) relative to the unaffiliated control faces. If the FFA reflects deeper encoding of faces, our data are consistent with models of person perception that posit categorical processing as the default (Fiske & Neuberg, 1990; Brewer, 1988). However, we found the opposite pattern of results in the primary visual cortex. Thus, the notion that there is a default mode of person perception may be untenable at the neural level, such

that different brain regions may be differentially sensitive to different information (Van Bavel & Cunningham, 2011; Cunningham, Zelazo, Packer, & Van Bavel, 2007).

In previous research, tasks that increase subordinate level judgments (Gauthier, Tarr, et al., 2000; Gauthier et al., 1997) or covert attention (Wojciulik et al., 1998) have been shown to increase FFA activity. For example, participants who were shown displays of two peripheral faces and two peripheral houses (presented simultaneously) had heightened activity in the FFA when they attended to the faces compared with the houses (Wojciulik et al., 1998). However, the effect of group membership on FFA activity in our study was not moderated by task, suggesting that in-group bias in the FFA did not require explicit attention to in-group membership. Moreover, there was a clear dissociation between activity in the FFA and the calcarine sulcus to unaffiliated faces, and the calcarine sulcus was more sensitive to the effects of explicit attention to in-group versus out-group membership than FFA. The dissociation between FFA and the calcarine sulcus provides discriminant evidence for our contention that group membership likely modulated the encoding of faces and not merely early attention to group membership. However, attention is a multifaceted construct and it remains possible that other aspects of attention may have modulated the effects observed in our study (Knudsen, 2007). For example, participants may have engaged in more sustained attention to in-group than to out-group members and unaffiliated faces in a fashion that selectively increased activity in the FFA but not the primary visual cortex. Future research should explore how different aspects of attention may mediate the influence of group membership on FFA activity.

The current research provides evidence that the motivational relevance of categories, like group membership and social identity, can affect FFA activity in a flexible and dynamic fashion in the absence of long-term experience with the category or explicit task instructions. Previous studies have shown that emotionally significant faces (e.g., fearful faces) can increase FFA activity under certain circumstances (Noesselt, Driver, Heinze, & Dolan, 2005; Vuilleumier, 2005; Winston, Vuilleumier, & Dolan, 2003). The current study extends the previous research by testing whether top-down, motivational factors can increase FFA activity to perceptually indistinguishable stimuli. Moreover, we found a correlation between differences in FFA activity to in-group versus out-group faces and differences in recognition memory for in-group versus out-group faces. Building on previous research showing that faces encoded at the subordinate level are correctly recognized more accurately than faces encoded at the superordinate level (Sporer, 2001), these results provide convergent evidence that the FFA may be engaged in subordinate-level processing for motivationally relevant faces (Kanwisher & Yovel, 2006; Grill-Spector, Knouf, & Kanwisher, 2004; Palmeri & Gauthier, 2004; Winston, Henson, Fine-Goulden, & Dolan, 2004; George et al., 1999). However, we found a similar pattern

of results for both the FFA and OFA, indicating that our conclusions may not be specific to the FFA.

The study of visual perception has made considerable progress in the past half century, yet there remains considerable debate about the relationship between vision and cognition. On the one hand, several theorists have argued that early vision is impervious to the influence of top-down factors, including expectations and motivations (Pylyshyn, 1999). On the other hand, there is evidence that emotionally or motivationally significant stimuli alter perceptual processing (Balcetis & Dunning, 2006; Vuilleumier, 2005; Bruner, 1957) and that the pFC may play a key role in biasing information processing in lower-order brain regions (Miller & Cohen, 2001). For example, recent research has shown that the top-down influence of context and motivation have the potential to modulate a host of lower-order systems (Cunningham, Van Bavel, & Johnsen, 2008; Kim et al., 2004). The current study adds to this debate, providing evidence that the FFA is sensitive to transient, top-down motivational influences.

Acknowledgments

This study was completed as part of Jay Van Bavel's Ph.D. Dissertation at the University of Toronto. We thank Sharon David, Andrew Jahn, Jillian Swencionis, Y. Jenny Xiao, members of the Social Cognition and Affective Neuroscience Laboratory, and several anonymous reviewers for comments on this manuscript. This research was supported by grants from the Social Sciences and Humanities Research Council of Canada to Jay Van Bavel and Dominic Packer and the National Science Foundation (BCS-0819250) to William Cunningham. Data were collected at the Queen's University MRI facility.

Reprint requests should be sent to Jay Van Bavel, New York University Department of Psychology, 6 Washington Place, New York, NY 10012, or via e-mail: jay.vanbavel@nyu.edu.

Note

1. We were unable to compute a measure of discrimination without creating visual cues for in-group and out-group members because visual cues are necessary to compute false alarms during the recognition task. However, we excluded visual cues for group membership after the learning phase to ensure that any differences between in-group and out-group faces reflected the top-down influence of social identity and not the bottom-up influence of visual cues/heuristics.

REFERENCES

Anderson, A. K., Christoff, K., Panitz, D., De Rosa, E., & Gabrieli, J. (2003). Neural correlates of the automatic processing of threat facial signals. *Journal of Neuroscience*, *23*, 5627–5633.

Balcetis, E., & Dunning, D. (2006). See what you want to see: Motivational influences on visual perception. *Journal of Personality and Social Psychology*, *91*, 612–625.

Benton, A. L., & Van Allen, M. W. (1968). Impairment in facial recognition in patients with cerebral disease. *Transactions of the American Neurological Association*, *93*, 38–42.

Bernstein, M., Young, S., & Hugenberg, K. (2007). The cross-category effect: Mere social categorization is sufficient to elicit an own-group bias in face recognition. *Psychological Science*, *18*, 709–712.

Brett, M., Anton, J. L., Valabregue, R., & Poline, J. B. (2002). *Region of interest analysis using an SPM toolbox*. Presented at the 8th International Conference on Functional Mapping of the Human Brain, June 2–6, 2002, Sendai, Japan.

Brewer, M. B. (1988). A dual process model of impression formation. In T. Srull & R. Wyer (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, *64*, 123–152.

Cunningham, W. A., Johnson, M. K., Raye, C. L., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2004). Separable neural components in the processing of black and white faces. *Psychological Science*, *15*, 806–813.

Cunningham, W. A., Van Bavel, J. J., & Johnsen, I. R. (2008). Affective flexibility: Evaluative processing goals shape amygdala activity. *Psychological Science*, *19*, 152–160.

Cunningham, W. A., Zelazo, P. D., Packer, D. J., & Van Bavel, J. J. (2007). The iterative reprocessing model: A multi-level framework for attitudes and evaluation. *Social Cognition*, *25*, 736–760.

De Renzi, E., & Spinnler, H. (1966). Facial recognition in brain-damaged patients: An experimental approach. *Neurology*, *16*, 145–152.

Ellinwood, E. H. J. (1969). Perception of faces: Disorders in organic and psychopathological states. *Psychiatric Quarterly*, *43*, 622–646.

Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1–74). New York: Academic Press.

Gauthier, I., Anderson, A. W., Tarr, M. J., Skudlarski, P., & Gore, J. C. (1997). Levels of categorization in visual recognition studied with functional MRI. *Current Biology*, *7*, 645–651.

Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, *3*, 191–197.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform “face area” increases with expertise in recognizing novel objects. *Nature Neuroscience*, *2*, 568–573.

Gauthier, I., Tarr, M. J., Moylan, J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (2000). Does visual subordinate-level categorization engage the functionally defined face area? *Cognitive Neuropsychology*, *17*, 143–163.

George, N., Dolan, R. J., Fink, G. R., Baylis, G. C., Russell, C., & Driver, J. (1999). Contrast polarity and face recognition in the human fusiform gyrus. *Nature Neuroscience*, *2*, 574–580.

Golby, A. J., Gabrieli, J. D. E., Chiao, J. Y., & Eberhardt, J. L. (2001). Differential fusiform responses to same- and other-race faces. *Nature Neuroscience*, *4*, 845–850.

Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience*, *7*, 555–562.

Grill-Spector, K., Sayres, R. A., & Ress, D. (2006). High-resolution imaging reveals highly selective nonface clusters in the fusiform face area. *Nature Neuroscience*, *9*, 1177–1185.

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

- Hugenberg, K., Young, S. G., Bernstein, M. J., & Sacco, D. F. (2010). The categorization-individuation model: An integrative account of the other-race recognition deficit. *Psychological Review*, *117*, 1168–1187.
- Kanwisher, N., McDermott, J., & Chun, M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for the perception of faces. *Journal of Neuroscience*, *17*, 4302–4311.
- Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *361*, 2109–2128.
- Kim, H., Somerville, L. H., Johnstone, T., Polis, S., Alexander, A. L., Shin, L. M., et al. (2004). Contextual modulation of fMRI responsivity to surprised faces. *Journal of Cognitive Neuroscience*, *16*, 1730–1745.
- Knudsen, E. I. (2007). Fundamental components of attention. *Annual Review of Neuroscience*, *30*, 57–78.
- Lieberman, M. D., Hariri, A., Jarcho, J. M., Eisenberger, N. I., & Bookheimer, S. Y. (2005). An fMRI investigation of race-related amygdala activity in African-American and Caucasian-American individuals. *Nature Neuroscience*, *8*, 720–722.
- Malpass, R. S., & Kravitz, J. (1969). Recognition for faces of own and other “race.” *Journal of Personality and Social Psychology*, *13*, 330–334.
- Maurer, D., O’Craven, K. M., Le Grand, R., Mondloch, C. J., Springer, M. V., Lewis, T. L., et al. (2007). Neural correlates of processing facial identity based on features versus their spacing. *Neuropsychologia*, *45*, 1438–1451.
- Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, and Law*, *7*, 3–35.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.
- Miner, M., & Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments & Computers*, *36*, 630–633.
- Noesselt, T., Driver, J., Heinze, H. J., & Dolan, R. (2005). Asymmetrical activation in the human brain during processing of fearful faces. *Current Biology*, *15*, 424–429.
- Palmeri, T. J., & Gauthier, I. (2004). Visual object understanding. *Nature Reviews Neuroscience*, *5*, 291–303.
- Posner, M. I., & Gilbert, C. D. (1999). Attention and primary visual cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *96*, 2585–2587.
- Puce, A., Allison, T., Gore, J. C., & McCarthy, G. (1995). Face-sensitive regions in human extrastriate cortex studied by functional MRI. *Journal of Neurophysiology*, *74*, 1192–1199.
- Pylshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*, 341–423.
- Quattrone, G. A., & Jones, E. E. (1980). The perception of variability within in-groups and out-groups: Implications for the law of small numbers. *Journal of Personality and Social Psychology*, *38*, 141–152.
- Rhodes, G., Michie, P. T., Hughes, M. E., & Byatt, G. (2009). The fusiform face area and occipital face area show sensitivity to spatial relations in faces. *European Journal of Neuroscience*, *30*, 721–733.
- Rorden, C., Karnath, H. O., & Bonilha, L. (2007). Improving lesion-symptom mapping. *Journal of Cognitive Neuroscience*, *19*, 1081–1088.
- Sergent, J., Ohta, S., & MacDonald, B. (1992). Functional neuroanatomy of face and object processing. A positron emission tomography study. *Brain*, *115*, 15–36.
- Sporer, S. L. (2001). Recognizing faces of other ethnic groups: An integration of theories. *Psychology, Public Policy, and Law*, *7*, 36–97.
- Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American*, *223*, 96–102.
- Tarr, M. J., & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level processing automatized by expertise. *Nature Neuroscience*, *3*, 764–769.
- Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D., & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford, UK: Basil Blackwell.
- Turner, J. C., Oakes, P. J., Haslam, S. A., & McGarty, C. (1994). Self and collective: Cognition and social context. *Personality and Social Psychology Bulletin*, *20*, 454–463.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., et al. (2002). Automated anatomical labelling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single subject brain. *Neuroimage*, *15*, 273–289.
- Van Bavel, J. J., & Cunningham, W. A. (2009). Self-categorization with a novel mixed-race group moderates automatic social and racial biases. *Personality and Social Psychology Bulletin*, *35*, 321–335.
- Van Bavel, J. J., & Cunningham, W. A. (2011). A social neuroscience approach to self and social categorization: A new look at an old issue. *European Review of Social Psychology*, *21*, 237–284.
- Van Bavel, J. J., Packer, D. J., & Cunningham, W. A. (2008). The neural substrates of in-group bias: A functional magnetic resonance imaging investigation. *Psychological Science*, *19*, 1131–1139.
- Vuilleumier, P. (2005). How brains beware: Neural mechanisms of emotional attention. *Trends in Cognitive Sciences*, *9*, 585–594.
- Winston, J. S., Henson, R. N. A., Fine-Goulden, M. R., & Dolan, R. J. (2004). fMRI adaptation reveals dissociable neural representations of identity and expression in face perception. *Journal of Neurophysiology*, *92*, 1830–1839.
- Winston, J. S., Vuilleumier, P., & Dolan, R. J. (2003). Effects of low-spatial frequency components of fearful faces on fusiform cortex activity. *Current Biology*, *13*, 1824–1829.
- Wojciulik, E., Kanwisher, N., & Driver, J. (1998). Modulation of activity in the fusiform face area by covert attention: An fMRI study. *Journal of Neurophysiology*, *79*, 1574–1579.