



Neural adaptation to faces reveals racial outgroup homogeneity effects in early perception

Brent L. Hughes^{a,1,2}, Nicholas P. Camp^{b,1,2}, Jesse Gomez^{c,d}, Vaidehi S. Natu^b, Kalanit Grill-Spector^{b,e}, and Jennifer L. Eberhardt^{b,2}

^aDepartment of Psychology, University of California, Riverside, CA 92521; ^bDepartment of Psychology, Stanford University, Stanford, CA 94305; ^cDepartment of Psychology, University of California, Berkeley, CA 94720; ^dHelen Wills Neuroscience Institute, University of California, Berkeley, CA 94720; and ^eStanford Neurosciences Institute, Stanford University, Stanford, CA 94305

Contributed by Jennifer L. Eberhardt, May 13, 2019 (sent for review January 2, 2019; reviewed by Anne Krendl and Amy R. Krosch)

A hallmark of intergroup biases is the tendency to individuate members of one's own group but process members of other groups categorically. While the consequences of these biases for stereotyping and discrimination are well-documented, their early perceptual underpinnings remain less understood. Here, we investigated the neural mechanisms of this effect by testing whether high-level visual cortex is differentially tuned in its sensitivity to variation in own-race versus other-race faces. Using a functional MRI adaptation paradigm, we measured White participants' habituation to blocks of White and Black faces that parametrically varied in their groupwise similarity. Participants showed a greater tendency to individuate own-race faces in perception, showing both greater release from adaptation to unique identities and increased sensitivity in the adaptation response to physical difference among faces. These group differences emerge in the tuning of early face-selective cortex and mirror behavioral differences in the memory and perception of own- versus other-race faces. Our results suggest that biases for other-race faces emerge at some of the earliest stages of sensory perception.

intergroup perception | race | neural adaptation | perceptual sensitivity

A core feature of stereotyping is the tendency to view outgroup members in categorical terms, while individuating members of one's own group (1, 2). The inability to distinguish between outgroup members can have immediate real-world effects, ranging from the embarrassing (conflating 2 coworkers of the same ethnicity) to the life-changing (identifying the wrong suspect from a police lineup). Are such mistakes based in errors of recollection and judgment, or do they emerge in the very way that we perceive members of other social groups? We demonstrate that outgroup deindividuation occurs in the earliest stages of face perception, in the form of reduced neural sensitivity to variability among other-race faces.

Humans extract social information at a glance. Race is a particularly salient social category that is detected within fractions of a second (3, 4) and implicates a wide range of cognitive processes, from attentional allocation (5), to memory (6, 7), to estimates of group variability (8). These disparate measures reveal a common effect: the tendency to individuate members of one's own racial or ethnic group and to conflate members of other groups (9, 10). The early perceptual mechanisms of these biases, however, are less understood. When we observe members of another racial group, are their actual physical distinctions blurred in our mind's eye?

Evidence from behavioral research suggests this may be the case. For example, White Americans are more attuned to perceptual differences between White faces than Black faces (11). Such effects are not reducible to low-level differences in stimuli, as they replicate among novel groups and identical racially ambiguous faces labeled as own- versus other-race (12, 13). The precise perceptual mechanisms are unclear but suggest that different outgroup members are perceived as multiple instances of the same category rather than distinct individuals (9). One intriguing possibility is that neural systems responsible for face

perception may be narrowly tuned to distinguish between different identities for ingroup members but more broadly tuned to category membership, rather than individual identity, for outgroup members. Such a mechanism would result in greater perceptual sensitivity toward ingroup identities and reduced perceptual sensitivity toward outgroup identities. However, extant research has not yet examined how face-sensitive neural systems respond to such intragroup variability.

Visual neuroscience provides tools to address this unanswered question. A large body of work highlights a core feature of brain processing: its tendency to habituate to repeated exposures of the same stimuli. Neural populations exhibit reduced activity after repeated exposure to stimuli to which they are tuned, a phenomenon known as neural adaptation (14). For example, face-sensitive neural populations in the fusiform face area (FFA) show a diminished response to blocks where the same face is presented multiple times, relative to blocks of different faces (15). Neural adaptation is further graded by the perceptual similarity across stimuli: more similar faces elicit greater neural adaptation over the course of their presentation, while more dissimilar faces lead to a release from adaptation and greater activity in face-selective regions (16).

Significance

The tendency to view members of social outgroups as interchangeable has long been considered a core component of intergroup bias and a precursor to stereotyping and discrimination. However, the early perceptual nature of these intergroup biases is poorly understood. Here, we used a functional MRI adaptation paradigm to assess how face-selective brain regions respond to variation in physical similarity among racial ingroup (White) and outgroup (Black) faces. We conclude that differences emerge in the different tuning properties of early face-selective cortex for racial ingroup and outgroup faces and mirror behavioral differences in memory and perception of racial ingroup versus outgroup faces. These results suggest that outgroup deindividuation emerges at some of the earliest stages of perception.

Author contributions: B.L.H., N.P.C., K.G.-S., and J.L.E. designed research; B.L.H., N.P.C., and J.G. performed research; B.L.H., N.P.C., J.G., and V.S.N. analyzed data; and B.L.H., N.P.C., and J.L.E. wrote the paper.

Reviewers: A.K., Indiana University; and A.R.K., Cornell University.

The authors declare no conflict of interest.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: The data presented in this paper are available through Open Science Framework, <https://osf.io/zsdyf/>.

¹B.L.H. and N.P.C. contributed equally to this work.

²To whom correspondence may be addressed. Email: bhughes@ucr.edu, ncamp@stanford.edu, or jleberhardt@stanford.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1822084116/-DCSupplemental.

Published online July 1, 2019.

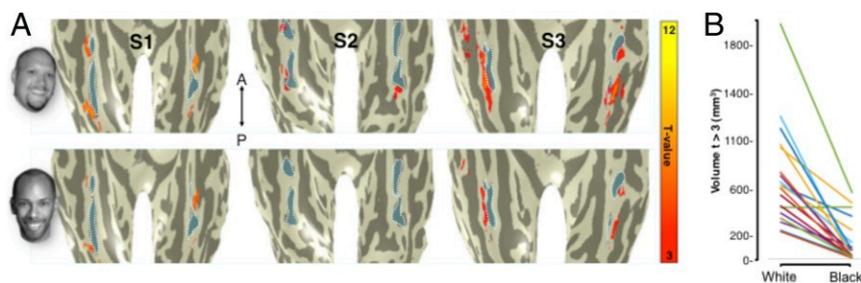


Fig. 1. The volume of face-selectivity in high-level visual cortex is modulated by face race. (A) Thresholded parameter maps show voxelwise t values for the contrast of White faces (Top) or Black faces (Bottom) versus all other stimuli. Data are presented on inflated cortical surfaces of 3 participants (dark regions are sulci; lighter regions are gyri). The region shown is the ventral surface of the temporal lobe, known as the VTC. The blue region outlined in dotted-white is the MFS, whose posterior and anterior tips are anatomical anchors predicting the location of face-selective cortex on the lateral fusiform gyrus. (B) Line plots illustrating the volume of above-threshold (t values > 3) activation in each subject's VTC defined with the contrasts of White or Black faces.

The present research harnesses this approach to examine the perceptual underpinnings of perceived racial homogeneity. The anatomical consistency of face-sensitive regions across humans (17) makes this area of visual cortex an ideal region to test how perceived outgroup homogeneity may reflect differential neural tuning properties. If neural adaptation underlies perceived outgroup homogeneity, then neural activity in the FFA should be less sensitive to physical variation among other- versus own-race faces. Specifically, to the extent that neuronal populations are broadly tuned to outgroup faces, the same population of neurons will respond across distinct outgroup targets, eliciting a greater habituation response across a wider range of physical variability.

To test this hypothesis, we used a functional MRI (fMRI) adaptation design (14, 18) to measure adaptation to own- and other-race faces at varying levels of physical similarity. Self-identified White participants first completed a face-localization task to independently define face-selective regions; in a second adaptation experiment, they responded to an oddball stimulus nested in blocks of own-race (White) and blocks of other-race (Black) faces. Using photo morphing, we parametrically manipulated the physical similarity of faces within each block, ranging from repetitions of the same face, to completely unique identities. By measuring the intragroup similarity of faces within each block against the corresponding adaptation response, we could derive the neural sensitivity to physical variation among faces and thus test the extent to which this sensitivity differed for variation in own- and other-race faces.

Neuroimaging Results

Localizer Task. Participants ($n = 20$) first completed 3 runs of a standard localizer experiment (19). Subjects viewed 4-s blocks of Black or White male faces (20) or of non-face stimuli (bodies, limbs, cars, guitars, houses, corridors, words, and numbers) at a rate of 2 Hz (SI Appendix, Fig. S1). To localize face-selective cortex, t -value parameter maps were produced in each subject contrasting faces versus all nonface stimuli. Data were unsmoothed. Face-selective cortex was defined on the cortical surface in each participant as clusters of voxels on the posterior (pFus-faces) and middle (mFus-faces) fusiform gyrus, just lateral to the mid-fusiform sulcus (MFS): a set of regions of interest (ROIs) commonly referred to as the FFA. To avoid bias in subsequent data analyses, face-selective ROIs were defined using both White and Black faces.

In each subject, we also analyzed the volume of face-selectivity within bilateral ventral temporal cortex (VTC) using contrast maps for White and Black faces, as the VTC plays a causal role in face perception (21). As Fig. 1A illustrates, the volume of activation to Black faces was a fraction of the volume to White

faces, despite being matched for low-level properties. In all but 1 subject, the volume of face-selectivity in VTC was significantly lower for Black faces [$t(19) = 6.5, P < 5 \times 10^{-6}$; Fig. 2B]; in 6 subjects, face-selectivity to Black faces was not observed at all.

Face Morph Task. Participants next completed a second experiment to measure their perceptual sensitivity to physical variation in Black versus White faces with neural adaptation. In each 3-s block, participants viewed a sequence of Black or White faces sampled from a separate stimulus set (22) and presented at 2 Hz, while responding to periodic oddball stimuli. We used photo morphing software (23) to generate sets of Black and White faces varying in physical dissimilarity from 0% (i.e., the same face presented 6 times) to 100% (6 separate identities), with intermediate levels at 30, 50, and 70% (Fig. 2A). This design let us fit response curves measuring neural adaptation across levels of groupwise variability. Given the causal link between face-selective cortex in the right hemisphere (24, 25), we hypothesized

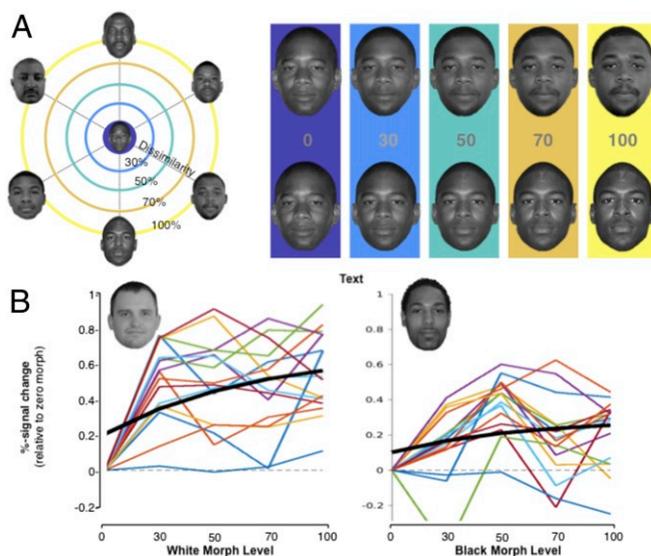


Fig. 2. (A) Experimental design of the adaptation experiment. In each block, 6 faces were presented at 1 of 5 levels of groupwise dissimilarity. An example of the morphing scheme is presented on the Left, and 2 example morph lines are shown on the Right. (B) Plots mapping the percentage of signal change in right hemisphere face-selective cortex in each participant for White (Left) and Black (Right) faces, normalized to the zero-morph level (full adaptation), along with the summary curve in black.

that differences in neural adaptation would be most pronounced in the right hemisphere.

Results demonstrate clear differences in neural adaptation between own- and other-race faces. Plotted in Fig. 2B are response curves extracted from the percentage of signal change in right hemisphere face-selective cortex in response to White and Black faces. We quantify these curves in 2 ways. First, we extract the initial rise of individual adaptation curves, which is the slope from the first (0% morph level) to second morph level. Second, we extract the limit approached by each curve by fitting exponential functions to each individual.

We observed a significant difference in both the rise [$t(16) = 5.13, P = 0.0001$] and limit [$t(16) = 2.7, P = 0.01$] in right face-selective cortex: other-race faces elicited a more gradual and smaller release in adaptation compared with own-race faces. While there was a significant difference in the rise of curve fits in left face-selective cortex [$t(16) = 3.1, P = 0.007$], there was no significant difference in the limit approached by the same curves [$t(16) = 0.89, P = 0.38$], suggesting a smaller contribution from left hemisphere to racial differences in neural responses. The significant difference in rise for own- versus other-race faces remains significant when excluding the participant with a negative response at the 30% morph level for other-race faces [$t = 6.57, P = 0.00013$]. We next asked if these racial differences in neural adaptation were mirrored in behavioral measures of perceptual sensitivity.

Behavioral Results

Participants completed 3 behavioral experiments following the fMRI session to assess individuation of racial ingroup and outgroup faces at different stages of processing: perceived similarity, perceptual discrimination, and recognition memory. Within each measure, we tested whether the threshold for differentiating White faces was lower than for Black faces. These measures complement the neuroimaging findings to pinpoint perceptual sensitivity mechanisms for the other-race effect (*SI Appendix, Fig. S2*).

Perceived Similarity. First, participants judged perceived similarity of groups of Black and White faces that varied in groupwise physical difference from 0 (i.e., completely identical), 30, 50, 70, and 100% (i.e., separate identities). On each trial, 6 faces from a corresponding level of dissimilarity were presented in an array until the participant rated how similar the faces were on a scale from 1 (“Completely Different”) to 7 (“Completely Identical”).

We tested whether the perceptual threshold for similarity differed for Black and White faces by fitting 2 local linear fit functions relating each participant’s perceptions of similarity to the objective similarity of the faces: 1 corresponding to their perception of Black faces and 1 for their ratings of White faces. The point at which these functions cross the y axis midpoint represents the point of subjective equality: the threshold at which groups of faces are seen as more similar than different. The level of physical difference to cross this threshold was higher for Black faces [mean (M) = 65.23, $SD = 15.12$] than for White faces [$M = 62.2, SD = 13.63, t(21) = 2.32, P = 0.03, d = 0.50$].

Perceptual Discrimination. In the perceptual discrimination task, participants judged whether pairs of faces were identical or different. Black and White morph continua were created in 10% increments. On each trial, 2 faces from a morph pair were presented for 500 ms on the screen: the 50% morph of the pair and a second face from the pair that was either identical (i.e., the 50% morph), or different. After 500 ms, participants indicated whether the 2 faces were identical or different. We used signal detection to analyze the threshold for participants’ judgments (c -criterion). While overall sensitivity (d') did not differ across

race [$t(23) = 0.31, P = 0.76$], participants set a lower criterion for identifying Black faces as identical ($M = -0.47, SD = 0.70$) than they used for White faces [$M = -0.18, SD = 0.67, t(23) = -2.66, P = 0.01, d = 0.54$].

Recognition Memory. To assess memory for own and other-race faces, we computed signal-detection measures for Black and White faces in an old/new recognition task. Repeated-measures t tests revealed that overall sensitivity did not differ by race [$t(23) = -1.50, P = 0.15$], but participants set a lower criterion for identifying Black faces as previously seen ($M = 0.04, SD = 0.38$) than for White faces [$M = 0.34, SD = 0.32, t(23) = -4.05, P < 0.001, d = -0.83$].

Brain-Behavior Correlations. To determine whether the race differences in the above behavioral measures correlated with neural adaptation responses (exponent fit and limit), we conducted Pearson correlations. For example, in each participant, we computed the difference between the exponent fit from right face-selective cortex for Black and White faces and correlated that value with the difference in behavioral sensitivity for Black and White faces. We did not observe any significant correlation between behavioral measures of outgroup deindividuation and neural measures of differential adaptation as a function of race (*SI Appendix, Table S2*). Whereas repeated-measure fMRI designs provide stable estimates of within-subject variability, the samples used here provide low power to reliably detect between-subjects effects. As such, we interpret these null results with caution but discuss their implications below.

Discussion

The present research used a fMRI adaptation design to test the sensitivity of face-selective cortex to variation in own- and other-race faces. By measuring signal change across blocks of faces varying in similarity, we fit exponential functions to each participant’s adaptation response. We could thus describe the rise, or release from adaptation, with increasing face dissimilarity, as well as the limit, or the drive in blood oxygenation level-dependent response approached by each participant’s response curve. Other-race faces elicited a more gradual and smaller release from adaptation and a reduced overall limit of activation relative to own-race faces: a slower and less pronounced recovery from adaptation. This effect was more pronounced in the right hemisphere, consistent with previous findings on face-selective cortex (22, 24, 25). In the process of localizing these regions, we further found that the volume of face-selectivity in bilateral VTC to other-race faces was a fraction of the response to own-race faces. The differences in neural adaptation and volume of face-selectivity to own- versus other-race faces cannot be explained by differences in low-level stimulus properties nor to attentional demand effects elicited by our tasks.

We collected complementary behavioral data to capture differences in individuation of own- and other-race faces. In each of these tasks, participants individuated own-race faces to a greater extent than other-race faces. While these findings parallel the reduced neural sensitivity to other-race faces in our adaptation and localization tasks, we did not observe significant correlations between behavioral and neural measures, findings that diverge from past work on differences in FFA activity and biases in memory (22, 26). One possibility for this discrepancy is that neuronal populations in visual cortex may be finely tuned to linear gradations in face identity, whereas subjective experience of these gradations is categorical. Indeed, past work shows that although visual cortex responds linearly to faces that vary in gender-typicality, subjective evaluations of gender remain categorical (27). Our findings point to one potential strength of using an fMRI adaptation design: it allows us to pinpoint a specific mechanism in visual cortex that may precede subjective

experience and would be impossible to access via traditional measures alone.

We conclude that racial disparities in perceptual individuation emerge in the tuning of face-selective cortex and mirror behavioral differences in memory and perception. These results suggest that biases for other-race faces emerge at some of the earliest stages of sensory processing and add further nuance and specificity to the mechanisms guiding intergroup perception. For example, past work documents differential activation in FFA and other face-sensitive regions for encoding own- versus other-race (22, 28, 29) and group (26, 30) targets. Such research has generally relied on a coarse neural measure of individuation: overall activity in FFA collapsed across all own-versus other-group targets (cf. ref. 31). While overall FFA activation suggests individual versus group-level processing (32), the adaptation paradigm used here allowed us to measure a more fine-grained mechanism of outgroup perception: neural sensitivity to graded variations in dissimilarity across Black and White faces. A further limitation of past work is that, since participants were explicitly instructed to memorize (22) or categorize (26) targets, group differences in FFA activity could have stemmed from cognitive or motivational regulation evoked by experimental demand, rather than perceptual sensitivity. In contrast, our design allowed us to assess perceptual sensitivity by removing such demand effects and instead having participants complete an oddball task orthogonal to the encoding of social targets.

The increased sensitivity to variation among own- versus other-race faces in face-selective cortex observed here likely represents the sharpening of neural tuning to own-race exemplars. That is, individuals are more sensitive to physical variation among own-race faces and, conversely, have broader tuning to other-race faces, habituating to them as repeated instances of the same social category rather than distinct individuals. This account is consistent with different theoretical explanations of the other-race effect (9, 10, 33). First, these biases could stem, in part, from perceptual expertise with own-race faces, resulting from a lifetime of experience interacting with racial ingroup members. Recent research consistent with this account finds that structural and functional properties of face-selective cortex are tuned by one's visual experience across development (18, 34). Alongside these experiential accounts, top-down processes could also contribute to the biases we observed. For example, recent work finds that patterns of FFA activity can accurately predict the race of face identities but only among more prejudiced individuals (35). Racial prejudice, then, may lead to divergent representations of own- versus other-race faces. Our current findings suggest one way these representations differ: populations of neurons in FFA are more narrowly tuned toward the individual identity of racial ingroup members and more broadly tuned toward social category information for outgroup members.

The fMRI adaptation approach applied here opens avenues for examining the flexibility of neural tuning across perceivers and their social context. Individuating experiences with racial outgroup members, for example, attenuate other-race biases in perception and memory (4, 36), and lack of exposure to racial outgroups can exacerbate them (37). It is possible, then, that cross-group contact over the course of development could lead to enhanced neural sensitivity to outgroup members. Future research can also test how the social and motivational context in which others are perceived influences neural tuning. For instance, perceivers are more likely to individuate other-race targets with whom they share a superordinate group identity (26, 38, 39). This would suggest that identical targets would elicit broader neural tuning if they were categorized as members of a rival team or political party than if they were designated as members of one's ingroup. In a similar vein, intergroup competition often leads individuals to view outgroup members in a more

categorical or stereotyped manner (40, 41); changes in this motivational context could accompany differences in adaptation to outgroup members.

More broadly, fMRI adaptation can shed light on higher-level social cognitive processes and their interaction with lower-level perception. Indeed, extant work has applied this technique to test whether higher-order representations of social targets draw on overlapping neuronal populations as the self (42) or members of one's own group (43, 44). The extent to which these social categories affect neural adaptation in early visual perception, however, had not been tested. We suggest that the failure to deeply process racial outgroup members as individuals could drive neural adaptation in response to dissimilar outgroup targets. This perspective is consistent with research demonstrating more shallow processing of individuating information about outgroup members and reduced prefrontal neural activity in response to such information (45, 46). Future research can formally test the connections between neural adaptation in low-level perception observed here and in higher-order representations of social groups. Recent work raises the possibility that these processes may be connected: stereotypic associations across race, gender, and emotion may influence the representation of social category exemplars in VTC (47).

The extent to which people individuate or conflate members of other social groups has a range of adverse consequences. Inaccuracies in distinguishing outgroup members can lead to mix-ups in the classroom or faulty eyewitness testimony. Given the role that categorical processing plays in intergroup biases more broadly, the biased perceptual mechanisms we document here could cascade into a range of harmful outcomes, from the application of group stereotypes to the spillover of fear or other responses across individuals.

Materials and Methods

Protocols were approved by the Stanford Institutional Review Board for Human Subjects Research, and informed consent was obtained from all participants; 24 participants (mean age, 19.7 y; 14 female) were recruited through a paid participant pool. We determined our sample size in advance of data collection, based on the robustness of visual stimulus-driven effects and from past work on neural adaptation to faces (18, 19, 48). All participants self-identified as White in a prescreening eligibility survey and had normal or corrected-to-normal vision. Participants completed a scanning sequence consisting of a T1-weighted anatomical scan, 3 runs of a face-localizer task (5 min, 24 s per run), and 4 runs of the fMRI adaptation experiment (4 min per run). Following the scanning protocol, participants completed behavioral tasks outside the scanner. Of the 24 subjects, 3 were unable to complete the full set of scans and 1 subject's data could not be properly reconstructed after collection. 3 of the remaining 20 subjects had atypical data from the face morph experiment (resulting from either fatigue in the scanner or venous artifacts inducing signal variability), such that typical fMRI adaptation to face similarity could not be observed. Of these 17 participants, face-selectivity could be localized on the fusiform gyrus for all participants in the left hemisphere and all but 1 for the right.

Localizer Task. Face-selective ROIs were identified for each participant in a functional localizer task adapted from prior work (19). In each run, participants viewed seventy-eight 4-s blocks of gray-scaled images randomly positioned in the visual field, presented against a scrambled noise texture at a rate of 2 Hz. In each block, participants viewed stimuli from 1 of 7 categories: Black male faces, White male faces, character strings (numbers and words), bodies (limbs and headless figures), objects (cars and guitars), scenes (buildings and corridors), and baseline blocks of background texture. Facial stimuli for the localizer were selected from the Bainbridge Face Database (20) and Mind, Culture, and Society laboratory database (49) to capture Black and White faces at various angles (direct, oblique) and expressions (neutral, smiling) to capture voxels that respond selectively to faces. Angles and expressions of face stimuli were balanced across race.

Face Morph Experiment. Stimuli for the adaptation experiment were drawn from a separate set of Black and White male faces used in previous imaging research (22). We created 6 White and 6 Black sets of faces. In each set, 1 "source" face was morphed with 6 remaining "target" faces in a

pairwise fashion, aligning n points around features of the source face and each target (eyes, nose, mouth, and external contour) and then interpolating between targets to make a continuum of images that varied in their similarity to the source face. Stimuli were converted to grayscale and matched for luminance using the SHINE toolbox (50).

To confirm that race was not confounded with low-level similarity, we calculated the Structural Similarity Index (SSIM) (51), an image-processing measure of the similarity between images. We computed the pairwise SSIM between the 6 target images in each morph line and between each source face and its associated targets. Black and White blocks did not significantly differ in the average SSIM of source and target faces [$t(10) = -0.71, P = 0.48$], or among target faces [$t(10) = 0.31, P = 0.76$], or in their variability in low-level similarity among target faces [$F(1,10) = 0.01, P = 0.93$] and between source and target faces [$F(1,10) = 0.72, P = 0.41$].

Images from the 6 morph lines within a set and at the same level of dissimilarity to the source face were grouped together into blocks of 6 faces. Participants viewed these stimuli in 4 runs of 4 min. Each run began and ended with 12 s of fixation and contained 18 blocks of White faces, 18 blocks of Black faces, and 18 fixation blocks. Each block lasted 3 s and consisted of 6 faces of a single morph level presented at a rate of 2 Hz, followed by 1 s of fixation. Each of the 5 morph levels (0, 30, 50, 70, 100%) occurred 3 times during a run, for each of the Black and White face blocks, and contained a different morph set. The block order was further counterbalanced so consecutive blocks never displayed stimuli from the same morph set or morph level. The appearance of images from the same morph set were pseudorandomized such that the next appearance of the same morph set would be in a block that was different by at least 2 morph steps from the preceding one (e.g., 0 to 50%, 100 to 30%). In each block, subjects were instructed to fixate on a central dot and to respond to an oddball image (a face outline containing a noise pattern) that appeared randomly within one-third of the blocks. Participants detected the oddball presence around half of the trials ($M = 50\%$, $SD = 23.2\%$).

Behavioral Measures. Participants completed 3 behavioral experiments following the MRI session: a measure of groupwise perceptual similarity, a perceptual discrimination task, and a memory recognition test. All tasks were presented using Medialab and DirectRT software on a 13-inch laptop computer.

Stimuli for the perceived similarity task consisted of the 6 morph groups presented in the scanner, as well as 6 additional groups created in the same manner, at 0, 30, 50, 70, and 100% levels of dissimilarity. On each trial, all 6 faces were presented in an array on the screen until the participant rated the similarity of the faces in the group on a scale of 1 ("Completely Different") to 7 ("Completely Identical"). 2 participants whose similarity judgments were not correlated with the objective similarity of the stimuli ($r = 0.48 [-0.41, 0.90]$) were excluded from analysis on this measure.

Stimuli for the perceptual discrimination task were created from 8 Black and 8 White faces from NimStim (52) and an internal laboratory database. Pictures were cropped to an oval window, and facial hair was removed with photo editing software. 4 Black and 4 White morph continua were created from these images in 10% increments. Participants completed 112 trials of the task. On each trial, 2 faces from a morph pair were presented for 500 ms on the screen: the 50% morph of the pair and a second face from the pair that was either identical (i.e., the 50% morph, 32 trials) or different (ranging from 10 to 50% different, 80 trials). After 500 ms, the faces were replaced by a fixation cross, and participants indicated via key press whether the 2 faces were identical or different.

In the recognition memory task, 80 faces (40 White, 40 Black) were selected from a separate stimulus set (49). In the encoding phase, participants were shown a sequence of 40 faces selected at random from this pool (20 White, 20 Black), at a rate of 2 s per face. Participants were instructed to memorize the faces, as they would answer questions about them later. Following a 3-min distractor task, participants proceeded to the test phase, in which they were presented with each of the 80 faces in a random order. Participants indicated via keyboard press whether or not they had seen each target during the encoding phase.

MRI Acquisition and Segmentation. Neuroimaging data were acquired from a 3T GE Signa scanner using a custom-built phase array 32-channel, receive-only head coil. During MRI scanning, participants lay supine inside the magnet. Visual stimuli were projected onto a monitor and were viewed through an angled mirror mounted to the head coil.

To obtain anatomical scans, we collected a whole-brain, anatomical volume [T1-weighted Brain Volume imaging pulse sequence; resolution: $1 \times 1 \times 1$ mm, time of inversion = 450 ms, flip angle = 12° , n of excitations = 1, field of view (FOV) = 240 mm]. Anatomical data were aligned to the anterior commissure–posterior commissure plane. Functional scans were obtained with the same scanner and coil using a T2*-sensitive gradient echo spiral pulse sequence with a resolution of $2.4 \times 2.4 \times 2.4$ mm, repetition time = 1,000 ms, echo time = 78.6 ms, flip angle = 76° , FOV = 192 mm. We collected 48 oblique slices, oriented parallel to the superior temporal sulcus, using a multiplexing technique allowing whole-brain coverage of functional data. The same prescription was used to obtain whole-brain anatomical T1-weighted images (inplane scan), which were used to align functional data with the high-resolution anatomical volume of each participant.

T1 anatomical images were segmented into white and gray matter using the FreeSurfer segmentation tool (<http://surfer.nmr.mgh.harvard.edu>). White matter surfaces were inspected and manually fixed for missing or mislabeled white matter voxels using ITK-SNAP (<http://www.itksnap.org/>). Reconstruction of cortical surface was generated from the boundary of white and gray matter. This surface was inflated for visualization of activity within sulcal folds.

MRI Data Analysis. Imaging data were analyzed using the MATLAB-based mrVista toolbox (<http://github.com/vistalab>) in each subject's native brain space. Data were corrected for within- and between-run head motion. In both localizer and face morph experiments, only datasets with motion of less than 2 voxels were included. Time courses for each voxel were converted from arbitrary scanner units into units of percentage of signal change. For each subject's data, we ran a generalized linear model (GLM) to model each voxel's time course. The experimental design matrix was convolved with the hemodynamic response function (<https://www.fil.ion.ucl.ac.uk/spm>) to generate predictors. Using a GLM to fit the predictors to the data, we estimated the response amplitudes for each condition (betas) and residual variance of each voxel's time course. We used the beta values and residual variance from the GLM to generate contrast maps comparing responses in different conditions.

Face-selective voxels were defined as voxels that responded significantly more to faces than images of other categories ($t > 3$, voxel level). We defined in each subject 2 face-selective regions in ventral stream (19), collectively known as the FFA: a region in posterior fusiform gyrus (pFus-faces), also referred to as FFA-1, and a region in mid-fusiform gyrus (mFus-faces), also referred to as FFA-2. ROIs were defined in both hemispheres and used for analyses in the face-morph experiment. For purposes of localizing face-selectivity, both Black and White faces were used in contrast maps. We also analyzed the volume of $t > 3$ voxels elicited by each race separately.

Response amplitudes for experimental conditions in the adaptation experiment were derived from betas estimated from the GLM. For each subject and ROI, we measured the amplitude of responses as a function of morph level. We fit an exponential model separately for responses to Black faces and White faces. The exponential data were modeled as $r = L(1 - e^{-y \cdot m})$, where m is the morph level step, L is the limit, y is the exponent, and r is the predicted voxel response. This function was fit for each subject's curve using a nonlinear least-squares, minimizing error by sweeping through parameter for L and y . Thus, an exponent and a limit were estimated for each participant per stimulus type (Black or White faces) and per ROI (pFus-faces, and mFus-faces).

ACKNOWLEDGMENTS. This research was supported by Stanford Dean's Awards (to J.L.E.) and by a Center for Cognitive and Neurobiological Imaging Innovation Grant to (to B.L.H. and N.P.C.). We wish to acknowledge Nalini Ambady for her contributions at the earliest stages of this research, although she did not live to witness its progression and completion.

1. G. W. Allport, *The Nature of Prejudice* (Addison-Wesley, Oxford, UK, 1954).
2. S. T. Fiske, S. L. Neuberg, "A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation" in *Advances in Experimental Social Psychology*, M. P. Zanna, Ed. (Academic, New York, NY, 1990), Vol. 23, pp. 1–74.
3. T. A. Ito, E. Thompson, J. T. Cacioppo, Tracking the timecourse of social perception: The effects of racial cues on event-related brain potentials. *Pers. Soc. Psychol. Bull.* 30, 1267–1280 (2004).

4. P. M. Walker, L. Silvert, M. Hewstone, A. C. Nobre, Social contact and other-race face processing in the human brain. *Soc. Cognit. Affective Neurosci.* 3, 16–25 (2007).
5. S. D. Goldinger, Y. He, M. H. Pappas, Deficits in cross-race face learning: Insights from eye movements and pupillometry. *J. Exp. Psychol. Learn. Mem. Cognit.* 35, 1105–1122 (2009).
6. R. S. Malpass, J. Kravitz, Recognition for faces of own and other race. *J. Pers. Soc. Psychol.* 13, 330–334 (1969).

7. C. A. Meissner, J. C. Brigham, Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychol. Public Policy Law* **7**, 3–35 (2001).
8. T. M. Ostrom, C. Sedikides, Out-group homogeneity effects in natural and minimal groups. *Psychol. Bull.* **112**, 536–552 (1992).
9. K. Hugenberg, S. G. Young, M. J. Bernstein, D. F. Sacco, The categorization-individuation model: An integrative account of the other-race recognition deficit. *Psychol. Rev.* **117**, 1168–1187 (2010).
10. S. L. Sporer, The cross-race effect: Beyond recognition of faces in the laboratory. *Psychol. Public Policy Law* **7**, 170–200 (2001).
11. O. Corneille, K. Hugenberg, T. Potter, Applying the attractor field model to social cognition: Perceptual discrimination is facilitated, but memory is impaired for faces displaying evaluatively congruent expressions. *J. Pers. Soc. Psychol.* **93**, 335–352 (2007).
12. O. Corneille, R. L. Goldstone, S. Queller, T. Potter, Asymmetries in categorization, perceptual discrimination, and visual search for reference and nonreference exemplars. *Mem. Cognit.* **34**, 556–567 (2006).
13. C. Michel, O. Corneille, B. Rossion, Holistic face encoding is modulated by perceived face race: Evidence from perceptual adaptation. *Visual Cognit.* **18**, 434–455 (2010).
14. K. Grill-Spector, R. Henson, A. Martin, Repetition and the brain: Neural models of stimulus-specific effects. *Trends Cognit. Sci.* **10**, 14–23 (2006).
15. K. Grill-Spector, R. Malach, fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychol.* **107**, 293–321 (2001).
16. N. Davidenko, D. A. Remus, K. Grill-Spector, Face-likeness and image variability drive responses in human face-selective ventral regions. *Hum. Brain Mapp.* **33**, 2334–2349 (2012).
17. K. S. Weiner, K. Grill-Spector, The improbable simplicity of the fusiform face area. *Trends Cognit. Sci.* **16**, 251–254 (2012).
18. V. S. Natu *et al.*, Development of neural sensitivity to face identity correlates with perceptual discriminability. *J. Neurosci.* **36**, 10893–10907 (2016).
19. A. Stigliani, K. S. Weiner, K. Grill-Spector, Temporal processing capacity in high-level visual cortex is domain specific. *J. Neurosci.* **35**, 12412–12424 (2015).
20. W. A. Bainbridge, P. Isola, A. Oliva, The intrinsic memorability of face photographs. *J. Exp. Psychol. Gen.* **142**, 1323–1334 (2013).
21. J. Parvizi *et al.*, Electrical stimulation of human fusiform face-selective regions distorts face perception. *J. Neurosci.* **32**, 14915–14920 (2012).
22. A. J. Golby, J. D. E. Gabrieli, J. Y. Chiao, J. L. Eberhardt, Differential responses in the fusiform region to same-race and other-race faces. *Nat. Neurosci.* **4**, 845–850 (2001).
23. Morpheus Software, Morpheus photo morpher, Version 3.17. <http://www.morpheussoftware.net> (2010).
24. N. Kanwisher, J. McDermott, M. M. Chun, The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J. Neurosci.* **17**, 4302–4311 (1997).
25. V. Rangarajan, J. Parvizi, Functional asymmetry between the left and right human fusiform gyrus explored through electrical brain stimulation. *Neuropsychologia* **83**, 29–36 (2016).
26. J. J. Van Bavel, D. J. Packer, W. A. Cunningham, Modulation of the fusiform face area following minimal exposure to motivationally relevant faces: Evidence of in-group enhancement (not out-group disregard). *J. Cognit. Neurosci.* **23**, 3343–3354 (2011).
27. J. B. Freeman, N. O. Rule, R. B. Adams Jr, N. Ambady, The neural basis of categorical face perception: Graded representations of face gender in fusiform and orbitofrontal cortices. *Cereb. Cortex* **20**, 1314–1322 (2009).
28. W. A. Cunningham *et al.*, Separable neural components in the processing of black and white faces. *Psychol. Sci.* **15**, 806–813 (2004).
29. J. Ronquillo *et al.*, The effects of skin tone on race-related amygdala activity: An fMRI investigation. *Soc. Cognit. Affective Neurosci.* **2**, 39–44 (2007).
30. M. Cikara, J. J. Van Bavel, The neuroscience of intergroup relations: An integrative review. *Perspect. Psychol. Sci.* **9**, 245–274 (2014).
31. K. G. Ratner, C. Kaul, J. J. Van Bavel, Is race erased? Decoding race from patterns of neural activity when skin color is not diagnostic of group boundaries. *Soc. Cognit. Affective Neurosci.* **8**, 750–755 (2012).
32. I. Gauthier *et al.*, The fusiform “face area” is part of a network that processes faces at the individual level. *J. Cognit. Neurosci.* **12**, 495–504 (2000).
33. J. C. Brigham, R. S. Malpass, The role of experience and contact in the recognition of faces of own-and other-race persons. *J. Soc. Issues* **41**, 139–155 (1985).
34. J. Gomez *et al.*, Microstructural proliferation in human cortex is coupled with the development of face processing. *Science* **355**, 68–71 (2017).
35. T. Brosch, E. Bar-David, E. A. Phelps, Implicit race bias decreases the similarity of neural representations of black and white faces. *Psychol. Sci.* **24**, 160–166 (2013).
36. M. E. Wheeler, S. T. Fiske, Controlling racial prejudice: Social-cognitive goals affect amygdala and stereotype activation. *Psychol. Sci.* **16**, 56–63 (2005).
37. E. H. Telzer *et al.*, Early experience shapes amygdala sensitivity to race: An international adoption design. *J. Neurosci.* **33**, 13484–13488 (2013).
38. C. Kaul, K. G. Ratner, J. J. Van Bavel, Dynamic representations of race: Processing goals shape race decoding in the fusiform gyri. *Soc. Cognit. Affective Neurosci.* **9**, 326–332 (2012).
39. L. Cosmides, J. Tooby, R. Kurzban, Perceptions of race. *Trends Cognit. Sci.* **7**, 173–179 (2003).
40. L. W. Chang, A. R. Krosch, M. Cikara, Effects of intergroup threat on mind, brain, and behavior. *Curr. Opin. Psychol.* **11**:69–73 (2016).
41. A. R. Krosch, D. M. Amodio, Economic scarcity alters the perception of race. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 9079–9084 (2014).
42. A. C. Jenkins, C. N. Macrae, J. P. Mitchell, Repetition suppression of ventromedial prefrontal activity during judgments of self and others. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 4507–4512 (2008).
43. T. Lau, M. Cikara, fMRI repetition suppression during generalized social categorization. *Sci. Rep.* **7**, 4262 (2017).
44. M. Cikara, J. J. Van Bavel, Z. A. Ingbretsen, L. T. Decoding “us” and “them”: Neural representations of generalized group concepts. *J. Exp. Psychol. Gen.* **146**, 621–631 (2017).
45. J. B. Freeman, D. Schiller, N. O. Rule, N. Ambady, The neural origins of superficial and individuated judgments about ingroup and outgroup members. *Hum. Brain Mapp.* **31**, 150–159 (2010).
46. B. L. Hughes, J. Zaki, N. Ambady, Motivation alters impression formation and related neural systems. *Soc. Cognit. Affective Neurosci.* **12**, 49–60 (2017).
47. R. M. Stoller, J. B. Freeman, Neural pattern similarity reveals the inherent intersection of social categories. *Nat. Neurosci.* **19**, 795–797 (2016).
48. A. Stigliani, B. Jeska, K. Grill-Spector, Encoding model of temporal processing in human visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E11047–E11056 (2017).
49. J. L. Eberhardt, P. A. Goff, V. J. Purdie, P. G. Davies, Seeing black: Race, crime, and visual processing. *J. Pers. Soc. Psychol.* **87**, 876–893 (2004).
50. V. Willenbockel *et al.*, Controlling low-level image properties: The shine toolbox. *Behav. Res. Methods* **42**, 671–684 (2010).
51. Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612 (2004).
52. N. Tottenham *et al.*, The nimstim set of facial expressions: Judgments from untrained research participants. *Psychiatry Res.* **168**, 242–249 (2009).