

Race-contingent face aftereffects: A result of perceived racial typicality, not categorization

O. Scott Gwinn

Department of Psychology, Macquarie University,
Sydney, New South Wales, Australia



Kevin R. Brooks

Department of Psychology, Macquarie University,
Sydney, New South Wales, Australia



Opposing aftereffects can be simultaneously induced by adapting to faces of different races distorted in opposite directions, allowing researchers to infer that faces are encoded against race-specific prototypes. This effect also suggests the existence of dissociable pools of neurons sensitive to race, each of which has been differently adapted to cause an opposite aftereffect. More recent studies have suggested that changes in the strength of race-contingent aftereffects reveal evidence of categorical perception, as they are larger when the adapting faces straddle the racial category boundary. We examined whether changes in these effects more closely correspond to a dichotomous categorical judgment, reflecting highly race-selective neural mechanisms, or more continuous perceptions of racial typicality, reflecting visual channels that are more broadly tuned. In Experiment 1, faces with a range of “morph levels” (i.e., relative contributions of Asian/Caucasian faces) were either rated on a continuous scale for Asian/Caucasian typicality, or simply categorized as Asian/Caucasian. As expected, typicality ratings showed a shallow slope (observers were sensitive to morph level over a broad range), while dichotomous racial categorization showed a steep slope (a rapid switch from categorization as Asian-Caucasian). In Experiment 2, race-contingent adaptation was assessed using test faces with various morph levels. Aftereffect size showed a shallow slope, closely resembling the racial typicality ratings, but showing a significant difference to the categorization data. This suggests that although the visual channels processing these faces do show some selectivity to race, they are sensitive to perceptions of racial typicality, showing a gradual transition of activity across a broad range of faces along the racial continuum.

Introduction

Even from the most fleeting glance of an unfamiliar face, a great deal of information about the owner can be gleaned, including robust estimates of their age, gender, emotional state, and race. Perceptions of race and the resulting consequences have fascinated psychologists from a diverse range of subdisciplines for many years. While social and personality psychology has presented evidence that race affects attitudes towards the person (e.g., Blair, Judd, Sadler, & Jenkins, 2002), the field of forensic psychology has investigated the impact this has on the legal system (e.g., Eberhardt, Davies, Purdie-Vaughns, & Johnson, 2006). Cognitive psychologists have also identified differences in participants’ ability to memorize faces of their own race compared to other races—a phenomenon known as the “other race effect” (for review see Meissner & Brigham, 2001). Although studies from these disciplines have taught us much about the consequences of race perception, the underlying mechanisms by which observers determine the racial origin of a face had, for many years, remained obscure.

More recently, techniques traditionally employed in areas of perception and psychophysics have been applied to the study of face perception. Examples include the decomposition of faces to isolate low-level visual properties such as spatial frequency (Keil, Lapedriza, Masip, & Vitria, 2008) and orientation (Dakin & Watt, 2009), and the use of 2AFC procedures to establish psychophysical thresholds (e.g., Brooks & Kemp, 2007; MacLin, MacLin, Peterson, Chowdhry, & Joshi, 2009). Perhaps one of the most informative of these techniques has involved the use of adaptation and the observation of aftereffects (Clifford & Rhodes, 2005; Webster & MacLeod, 2011).

Citation: Gwinn, O. S., & Brooks, K. R. (2013). Race-contingent face aftereffects: A result of perceived racial typicality, not categorization. *Journal of Vision*, 13(11):13 1–11, <http://www.journalofvision.org/content/13/11/13>, doi:10.1167/13.11.13.

Sometimes described as the psychologist's micro-electrode (Frisby, 1980), aftereffects are characterized by changes in the perception of a "test" stimulus following an extended period of exposure to an "adaptation" stimulus. For example, continually viewing downward motion results in a stationary scene being perceived as drifting upwards – the well-known Waterfall Illusion (Addams, 1834; Barlow & Hill, 1963). McCollough (1965) was further able to demonstrate that opposing color aftereffects could be simultaneously induced contingent upon gratings of different orientation. That is, after simultaneously adapting to orange vertical and blue horizontal lines, participants reported that subsequently viewed vertical lines appeared blue-green and horizontal lines appeared orange. This was explained as reflecting the differential color adaptation of dissociable pools of cells, each sensitive to different orientations.

Since these experiments, many similar aftereffects have been demonstrated using faces as stimuli. These can involve changes in the perceived size and position of features (Webster & MacLin, 1999), as well as more complex aftereffects of identity (Leopold, O'Toole, Vetter, & Blanz, 2001) and sex (Webster, Kaping, Mizokami, & Duhamel, 2004). In Webster and MacLin (1999), participants adapted to faces that had been expanded (or contracted) around the midpoint of the nose, causing test faces to appear contracted (or expanded). As with other aftereffects, the appearance of the test stimulus is biased away from the adaptor.

Contingent aftereffects have also been shown in the context of face perception. Opposing aftereffects have been induced contingent upon a face's race (Jaquet, Rhodes, & Hayward, 2007, 2008; Little, DeBruine, Jones, & Waitt, 2008), sex (Little, DeBruine, & Jones, 2005; Jaquet & Rhodes, 2008), age and species (Little et al., 2008). In Jaquet et al. (2008) participants adapted to contracted Asian and expanded Caucasian faces (and vice versa in another condition). Subsequently, slightly contracted Asian test faces were judged as more normal whereas preferred Caucasian faces were slightly expanded. As with low-level contingent aftereffects, these findings indicate dissociable neural populations sensitive to different categories of faces.

Levels of dissociation between neural populations are commonly investigated by examining the degree to which aftereffects are selective for or transfer ("cross-adapt") between stimulus categories. Jaquet et al. (2008) showed that although aftereffects were somewhat selective for race, under certain conditions they could also transfer between Asian and Caucasian faces. This finding of cross-adaptation suggests that faces of different races are encoded by both common and selective mechanisms. That is, neurons may fire maximally for faces of a specific race but will retain some degree of activation regardless of race. Parallel

evidence of both common and selective mechanisms has also been found for faces of different sex (Jaquet & Rhodes, 2008).

Further research into the mechanisms responsible for the encoding of race suggests that the neural populations involved show influences of categorical perception (Jaquet et al., 2007). This study showed that opposing race-contingent aftereffects were larger for average Asian and Caucasian facial images compared to caricatured Asian and average Asian faces. That is, when adapting image pairs straddled racial category boundaries contingent aftereffects were larger compared to when they did not straddle the boundaries. This difference was evident despite the fact that all image pairs featured the same degree of difference to each other, in terms of the low-level properties of the image (i.e., each pair spanned equal distances along a racial morph continuum). It was concluded that the neural populations responsible for this adaptation effect show selectivity for racial category membership, changing their relative activation in a nonlinear manner as the contributions of Asian and Caucasian faces in the stimulus is manipulated. Consistent with this finding is the demonstration that discrimination accuracy was greater for face pairs that fell either side of a racial category boundary, compared to pairs within the same category (Levin & Beale, 2000). Although accuracy is higher near this category boundary, the authors conceded that theirs was not a demonstration of "pure" categorical perception. Observers are not blind to the differences that exist within a category, as discrimination between such stimuli was still possible.

From the evidence currently available it seems clear that (a) the neural populations responsible for the encoding of faces of different races are not completely independent (Jaquet et al., 2008), with each population retaining some degree of activation regardless of the race of the face, and that (b) their relative activation varies in a nonlinear fashion across the race continuum, with a steeper slope near the category boundary than at the extremes (Jaquet et al., 2007). However, no further conclusions can be drawn regarding the shape of the neural response functions of these populations. Pure categorical perception would demand that the designation of a face as belonging to one race or another should be the sole determinant in the difference in activity between the populations, with a sharp reversal in the relative activation at the category boundary. Figure 1a represents this situation in a two-pool system where the steep slope of the lines represents an abrupt change in the activation of selective neurons around the category boundary. Such mechanisms would involve neurons that are maximally activated once a face is categorized as belonging to a given race, and are minimally activated when it is categorized differently.

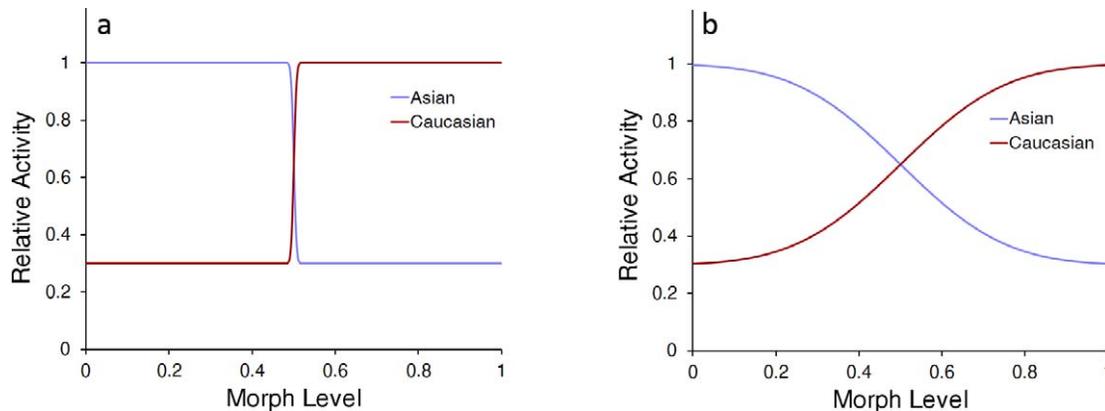


Figure 1. Plot showing potential relative activation levels of mechanisms with (a) an abrupt transition of activity levels around the category boundary, and (b) a gradual transition, representing sensitivity to levels of racial typicality across the range of morph levels. Based on results from Jaquet et al. (2008) both models would involve nonzero levels of activation in channels not tuned to the race being viewed.

Alternatively, faces of different races may be processed by mechanisms whose activity levels change gradually as the racial composition of the target face is manipulated, with only faces high in racial typicality producing a large difference between the activation of cells selective for one race as opposed to another. In Figure 1b this scenario is represented by the lines whose slope represents a more gradual change in activation levels.

The aim of the current study is to determine whether the shape of the response functions for visual channels encoding race show a steep transition at the category boundary, or whether the rate of change of activity is more gradual across the racial morph continuum. To answer this question, we investigated the possible links between either racial categorization or ratings of typicality with the size of race contingent aftereffects as faces cross the racial category boundary. Morphed facial images differing in relative contributions of Asian/Caucasian component faces were rated on a continuous scale for Asian/Caucasian typicality as well as being dichotomously categorized as Asian/Caucasian (Experiment 1). The slopes of curves fit to these data were compared to those obtained using measures of the strength and direction of contingent aftereffects (Experiment 2). If the neural populations involved in processing faces of different races show an abrupt transition as the racial composition of the stimulus is manipulated, then slopes for aftereffect data should closely resemble those for categorical judgments. However, if the selective mechanisms underlying the encoding of race show a more gradual transition of activity levels to faces of differing racial composition, then slopes for the race-contingent aftereffect should be more similar to those for ratings of racial typicality.

Experiment 1

Methods

Participants

Twenty-eight Caucasian undergraduate students (12 male) with a mean age of 23.8 ($SD = 12.6$) enrolled in psychology at Macquarie University took part in the study for course credit. All were instructed to wear optical corrections if required.

Stimuli and apparatus

Twenty-four male and 24 female frontal images of Asian and Caucasian faces were collected (96 images in total).¹ Keeping race and sex separate, images were randomly divided into groups of four and combined using Abrosoft FantaMorph (California, USA), resulting in a total of 24 averaged facial images with typical male Asian, female Asian, male Caucasian, and female Caucasian characteristics (six in each set). Averaged images were used to reduce idiosyncratic variations between faces and ensure each was representative of their race. Using Adobe Photoshop CS5 (Adobe, San Jose, CA), images were cropped following the jaw and hairline to remove all information outside of the face. Images were also formatted to a standard width of 255 pixels, measured across the midline of the eyes and placed on grey backgrounds with a width and height of 377×490 pixels. Each Caucasian image was then paired with an Asian image of the same sex to create 12 pairs. Using Abrosoft FantaMorph, each of these pairs was used to construct continua containing 11 faces with different morph levels representing the amount of input being contributed by the Caucasian component face. Given our particular interest in faces

near the likely location of the category boundary, this included seven faces with Caucasian contributions from 0.35 to 0.65 in 0.05 increments, along with the two extreme morph levels (0 and 1) and those at 0.2 and 0.8. Images were presented on a Toshiba Tecra laptop (Tokyo, Japan) with a screen size 30 cm × 23 cm and resolution of 1024 × 768 pixels. This meant that all faces had a standard width of 7.65 cm and from a viewing distance of 50 cm subtended a visual angle of 8.75°.

Procedure

Half of the participants were instructed that they would be shown a series of faces and for each one would be required to assign them a number between 0 and 10 to indicate how typically Asian or Caucasian they appeared, with higher numbers indicating a more typically Caucasian appearance. The other half were instructed to simply categorize each face as either Asian or Caucasian by assigning them a 0 or 1 respectively. A 3 s delay from the image onset to requesting a rating was included in the program to ensure participants viewed each image for an adequate amount of time before responding. Faces were presented in blocks of male and female, the order of which was counterbalanced between participants. The order of images within each block was randomized. Participants in the categorization condition saw each image five times and an average of their ratings was taken, allowing for uncertainty or inconsistent criteria for categorization to be appropriately represented by fractional values, while certainty is indicated by extremes of 0 or 1.

Results

Data were analyzed in two ways: as aggregate data, and as individual data. In the first analysis, (aggregate data) ratings for each morph level were averaged across continua within each participant and then averaged across participants. This resulted in a single mean score (with associated variance) for each morph level. These data were subject to curve-fitting principally to characterize the rate of change of values with morph level. Data were scaled to lie between zero and one to facilitate comparisons between typicality ratings and categorization data. Both typicality ratings and racial categorization data appeared sigmoidal in shape, representing the familiar psychometric functions (see Figure 2a and b). A cumulative Gaussian curve was fitted to typicality and categorization data using GraphPad Prism (GraphPad Software, La Jolla, CA) 5.² Both categorization and typicality data produced curve fits with $R^2 = 0.9$ and 0.87 respectively. While the horizontal position of a cumulative Gaussian is given

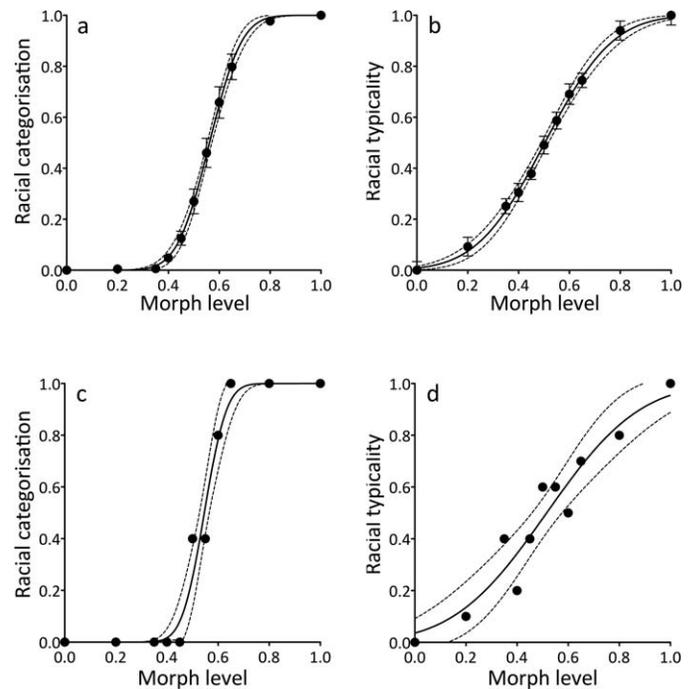


Figure 2. Curves fit to (a) aggregate categorization data, (b) aggregate typicality data, (c) one participant’s individual categorization data and (d) another’s individual typicality ratings. Lower numbers on x-axes indicate images that contained more input from Asian component faces. Lower numbers on y-axes indicate images that were more often categorized as Asian (a, c) or rated as appearing more typically Asian (b, d). Curves shown in (c) and (d) are for the same images. Error bars show standard error of mean (SEM). Dashed lines show a 95% confidence band.

by its mean, the slope—the principle dependent variable in this study—is given by its standard deviation. We will refer to this as the “slope” to avoid confusion with issues of variance between data values from different participants, and will not consider the horizontal position of the curve further. The slope for the curve fit to typicality ratings ($slope = 0.21$) appeared to be steeper than that for categorization ($slope = 0.09$). An extra sum-of-squares F -test confirmed that these values were significantly different, $F(1, 304) = 60.5$, $p < 0.0001$. This result indicates that ratings of typicality produce a relatively shallow slope with participants sensitive to differences in morph level over a broad range. As expected, the slope for categorization is much steeper and indicates a more rapid switch from classification as Asian to Caucasian within a relatively narrow range of intermediate morph levels, showing that consistent criteria were used across repetitions.

In the second analysis (individual data) curves were fit to each participant’s ratings for each of the continua (see Figure 2c and d). The mean of the 12 slopes (six male and six female continua) was calculated for each individual and then averaged across participants

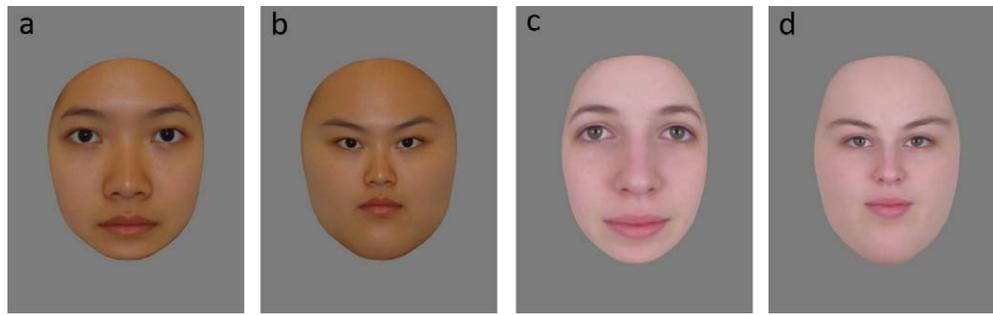


Figure 3. Examples of expanded Asian (a), contracted Asian (b), expanded Caucasian (c) and contracted Caucasian (d) adaptors used in the aftereffects experiment.

resulting in a single value (with associated variance) representing the average slope of all typicality ratings and an average slope of all categorization data. One participant's typicality ratings for one continuum produced a poor fit ($R^2 = 0.24$) and this slope value was excluded from the analysis. Remaining curves for categorization and typicality data all produced fits with $R^2 > 0.74$ and > 0.51 , respectively). The mean of the typicality slopes ($M = 0.22$, $SD = 0.05$)³ appeared to be larger than that for categorization ($M = 0.07$, $SD = 0.01$). An independent samples t -test (two-tailed) confirmed that these values were significantly different ($t_{26} = 10.55$, $p < 0.0001$). This result again indicates a relatively shallow slope for typicality ratings while the slope for categorization is much steeper.

Experiment 2

Methods

Participants

One of the authors and 30 Caucasian undergraduate students (eight male) with a mean age of 21.2 ($SD = 5.9$) enrolled in psychology at Macquarie University took part in the study for course credit. All were instructed to wear optical corrections if required.

Stimuli and apparatus

Adapting images: Images used in adaptation were the endpoints of each of the continua. Two versions of each face were created in which their internal features had been either contracted or expanded by 50% using the “spherize” function in Adobe Photoshop (see Figure 3). Images were presented on a Dell P1130 monitor (Dell, Round Rock, TX) with a screen size 40 cm × 30 cm and resolution of 1024 × 768 pixels. This meant that all the faces had a standard width of 9.95 cm and from a viewing distance of 65 cm subtended a visual angle of 8.75°.

Test images: Seven morph levels were selected to test for a race-contingent aftereffect. These included the two endpoint (0 and 1) and midpoint (0.5) images as well as those with values of 0.2, 0.35, 0.65, and 0.8 (see Figure 4). Thirteen versions of each face were created ranging from –30% contraction to +30% expansion with 5% steps in between. Test images were made $\frac{3}{4}$ the size of adapting images. Each participant only saw test images from one continuum and adapting images were never from the same continuum as test images.

Procedure

The experiment was run over three blocks defined by the morph level of the test faces involved. Block one comprised 0 and 1 images, block two 0.2 and 0.8, block three 0.35, 0.65, and 0.5. The order of blocks was randomized across participants. Participants were only shown face stimuli from one sex and this was also randomized across participants.

Adaptation: Each block began with a 2 min adaptation phase. During adaptation, participants viewed three Asian and three Caucasian faces presented individually and sequentially for 2 s at a time on a continual loop. Each participant only saw expanded versions of Asian images and contracted versions of Caucasian images, or vice versa. Participants also saw 6 s of “top-up” adaptation in between each test trial. This consisted of three Caucasian and three Asian faces each shown for 1 s. The presentation order of images was randomized but the same face was never presented twice in a row.

Test phase: The presence of a race contingent aftereffect was measured using a double interleaved staircase. Following the initial adaptation phase participants were shown a test image and were required to rate whether the face appeared expanded or contracted. If a participant rated an image as appearing expanded, the next time that face was seen a more contracted version was presented. This continued until a reversal was obtained (participants rating the face as appearing contracted) at which point a more expanded version was presented. Beginning with a 50% adjustment in

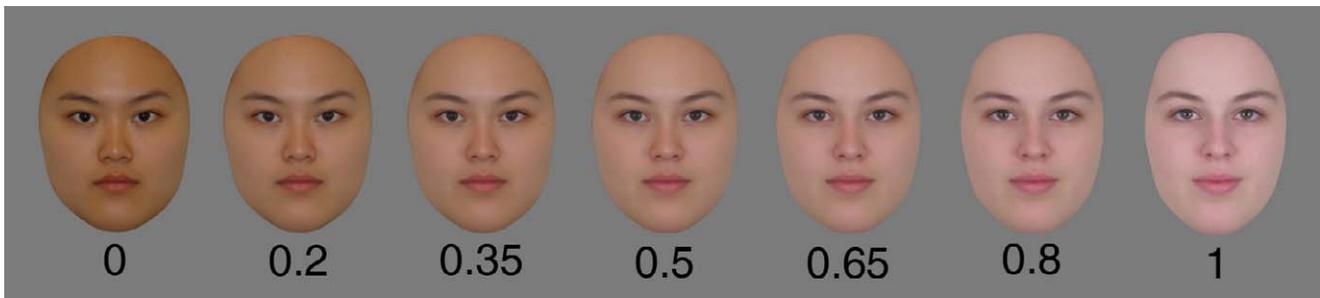


Figure 4. One continuum of images used in the aftereffects experiment. Values show input contributed by the Caucasian component face.

expansion/contraction (e.g., going from -30% contraction to $+20\%$ expansion), this step size was reduced by 15% after each reversal so the minimum 5% step size was reached after the third reversal. Each staircase progressed until it reached a maximum of eight reversals (or 28 trials), at which point the average distortion level of each participant's last four reversals was calculated. This calculation represents an estimate of the point of subjective equality (PSE) indicating the level of expansion/contraction an observer perceives as being undistorted. Three participants who failed to reach eight reversals within 28 trials were not included in analyses. Two interleaves were run for each morph level with one starting at fully expanded and the other fully contracted. The order of interleave presentation was randomized, and PSE estimates from the two interleaves were averaged to reach a single PSE value for each participant in each condition.

Results

Following the same procedure described in Experiment 1, PSEs were separately analyzed as aggregate and individual data and fit with cumulative Gaussian curves. Aggregate results for the size of the contingent face aftereffect are plotted in Figure 5. Here, higher values represent a perception of normality for expanded stimuli, indicating that a perceptual aftereffect of contraction has occurred. Conversely, lower values indicate that contracted faces appear undistorted, indicating an expansion aftereffect. In Figure 5 it can be seen that the majority of PSEs are negative. This represents a general tendency for our participants to consider undistorted faces to be slightly expanded (by approximately 5%), irrespective of the test condition. More importantly, PSEs recorded by our participants differed markedly between adaptation and test conditions. After adapting to expanded Asian and contracted Caucasian images (Asian+Caucasian-), participants recorded high PSEs for typical Asian test images (0), and low values for typical Caucasian faces

(1). This demonstrates a race-contingent aftereffect, where opposite adaptation effects have been induced on stimuli from different racial groups. Between the two extremes, intermediate degrees of adaptation were demonstrated as morph level changed. The opposite pattern of results was found for participants who adapted to contracted Asian and expanded Caucasian images (Asian-Caucasian+), demonstrating that this effect is due to the adaptation manipulation, and is not an inherent perceptual bias that applies to these groups of stimuli. A 2×7 ANOVA with the between subjects variable "adaptation condition" and the within subject variable "morph level" respectively, confirmed the presence of a significant interaction, $F(1, 29) = 8.166$, $p = 0.008$, and hence a race-contingent aftereffect. This statistic represents a lower-bound estimate.

Each participant's PSEs were then ordered in terms of their racial similarity to the expanded adaptor (i.e., data from the Asian+Caucasian- condition were flipped left-to-right), to be fit by a single cumulative Gaussian curve, as shown in Figure 6a ($R^2 = 0.12$). In Figure 7a it can be seen that the slope produced by the aftereffect data more closely resembles the typicality

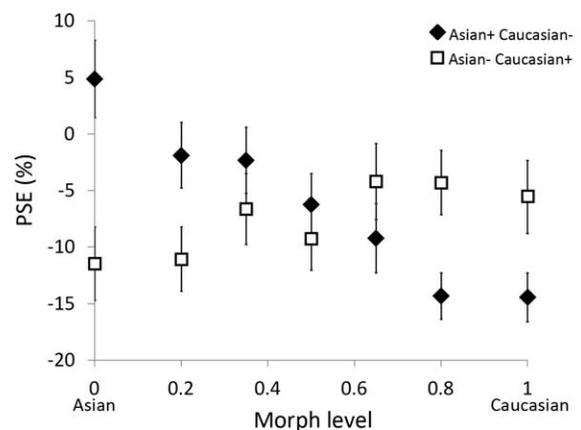


Figure 5. Mean PSEs for facial images following adaptation to expanded Asian and contracted Caucasian images (black diamonds), or to contracted Asian and expanded Caucasian images (open squares). Error bars show standard error of mean.

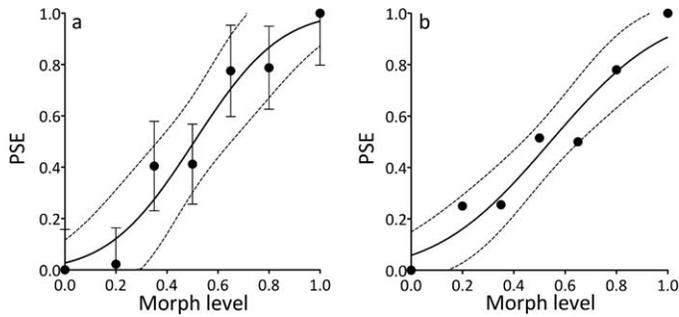


Figure 6. Curves fit to (a) aggregate aftereffect data and (b) one participant's individual aftereffect data. Lower numbers on x-axes indicate images more similar in racial appear to the contracted adaptors. Lower numbers on y-axes indicates PSEs corresponding to more contracted images. Curve shown in (b) is data for the same images represented in Figure 2. Error bars show *SEM*. Dashed lines show a 95% confidence band.

ratings than the categorization data. Extra sum-of-squares *F*-tests confirmed that the aftereffect slope (slope = 0.26) was significantly different to the categorization slope, $F(1, 367) = 3.98, p = 0.046$, but not the typicality slope, $F(1, 367) = .25, p = 0.58$, indicating that the rate of change of race-contingent aftereffects more closely matches the pattern describing a face's perceived racial typicality, rather than its categorization as belonging to one racial group or the other.

Figure 6b presents an example of a curve fit to one participant's individual aftereffect data. Although curve fits were successful in terms of goodness of fit for 16 of our observers, 15 of our participants' data showed variations not compatible with a strict cumulative Gaussian shape. Slopes from these curve-fits were deemed unreliable and were excluded from further analysis with all remaining fits producing $R^2 > 0.5$. In Figure 7b it can be seen that the average of all participants' slopes produced by the individual aftereffect data more closely resembles the typicality ratings than categorization. A pair of independent samples *t*-

tests (two-tailed) confirmed that the aftereffect slope ($M = 0.23, SD = 0.12$) was significantly different to the categorization slope ($t_{29} = 5.37, p < 0.0001$) but was not significantly different from the typicality slope ($t_{29} = 0.48, p = 0.632$), again indicating that the rate of change of race-contingent aftereffects more closely matches the pattern describing a face's perceived racial typicality, rather than its categorization as belonging to one racial group or the other.

Discussion

The aim of the current study was to examine whether the pattern of aftereffects across morph level best matches the pattern of responding for typicality or categorization, and to infer the shape of the response functions of neural mechanisms underlying the encoding of race. Faces ranging in racial appearance were rated on their typicality and also categorized as either Asian or Caucasian. Curves were fit to these data and the magnitudes of the slopes produced were compared to those representing the strength and direction of race-contingent aftereffects. Rather than showing a sharp transition in aftereffect direction as faces went from being categorized as one race to the other, the degree of aftereffect changed slowly as the racial typicality of test faces changed. Results remained consistent whether curves were fit to the aggregate of participants' data or to each participant's individual data.

The most parsimonious account for these data is a two-channel model in which neurons are selective for race, as suggested by previous work showing race-contingent adaptation (Jaquet et al., 2007, 2008; Little et al., 2008). Previous evidence suggests that these mechanisms are not entirely race selective, as cross-adaptation of a distortion aftereffect was demonstrated between faces of different races (Jaquet et al., 2008). This implies that the pools of neurons retain a certain level of activity regardless of the race of the face being

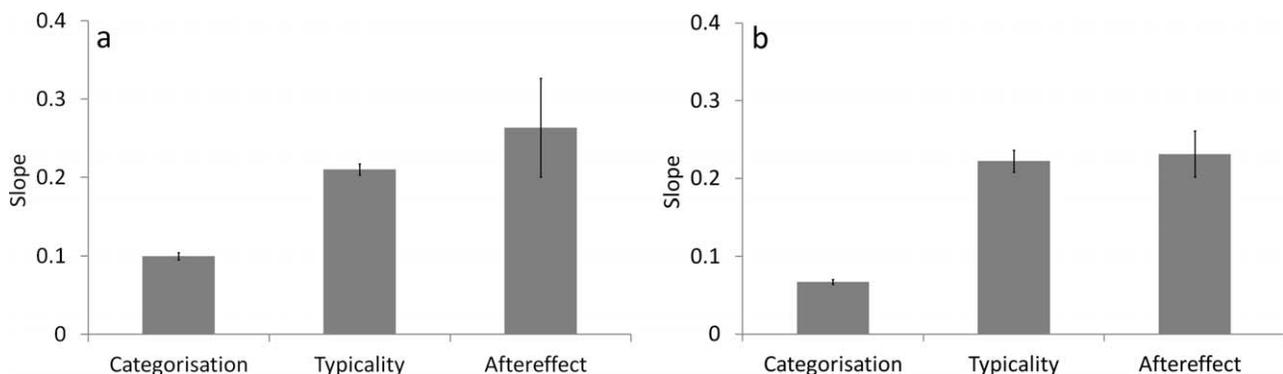


Figure 7. Slope values for (a) aggregate and (b) individual categorization, typicality, and aftereffect data. Error bars show *SEM*.

viewed. Although this was already known before the current study, details of the shape of their response functions were not. The current study has extended this work by revealing the tuning properties of these channels, suggesting that levels of activity change gradually as the race of the stimulus changes between typical Asian and typical Caucasian, showing substantial overlap in their dynamic ranges. Rather than the extent of activity in each pool being determined solely by the category that the stimulus fits into (i.e., Asian or Caucasian), it changes in a manner more consistent with the perceived racial typicality of the face. As such, we prefer a model akin to that depicted in Figure 1b, over 1a.

Although adaptation stimuli were racially typical (averages of several 100% Asian or of 100% Caucasian faces), causing the same degree of adaptation to each pool of neurons in all conditions, the various morph levels of the test stimuli allowed different aftereffects to be expressed as they recruited adapted neural populations with different racial selectivities. For those stimuli at the extremes, activation would be at its highest for the pools of neurons selective for the relevant race, and at its lowest for those selective for the other race, resulting in relatively large aftereffects. For stimuli with intermediate morph levels, the smaller aftereffects that were measured are likely to be the net result of differing degrees of competition between two aftereffects: expansion produced by neurons selective for one race, and contraction in neurons selective for the other.

Although there is a long and rich history of researchers using the paradigm of adaptation to reveal the characteristics of neural mechanisms underlying perceptual tasks (see Frisby, 1980; Clifford & Rhodes, 2005), it should be remembered that our conclusions are inferences, rather than assertions. Although the broadly-tuned two-channel model may be the most parsimonious account of our data, the current investigation is unable to provide information on several aspects of race coding that may be of interest, such as the possibility that the effects may be accounted for by a multichannel model and the possibility that not all channels have equivalent tuning properties. It is conceivable that the effects reported here could be explained by a system comprising many channels each tuned to a specific point on the racial continuum. If this were the case, each of these channels would have to be broadly tuned to account for the effects of racially typical (0 and 1) adaptors on mixed race test faces (e.g., 0.35). Within each of these channels, a different extent of expansion/contraction aftereffect would manifest dependent upon the degree to which their tuning curves overlapped with the adaptor. However, some may consider this unlikely. While we frequently encounter racially typical individuals, those with mixed race heritage are less prevalent.⁴ It may be considered

surprising if the human visual system had developed channels specifically tuned to such faces.

Several studies have attempted to parse two-channel and multiple-channel models of face adaptation aftereffects, by examining the relationship between aftereffect strength and the position of the adaptor on the dimension being adapted (Robbins, McKone, & Edwards, 2007; Jeffery et al. 2010). An alternative technique has been suggested by Storrs & Arnold (2012), who used a ternary classification task to probe the same issue. Interestingly, results suggested that different aspects of face processing may be underpinned by a two-channel system (facial distortion aftereffects), while others (gender aftereffects) are more consistent with a multichannel explanation.

In addition, the adaptation paradigm is silent on the neural locus of the aftereffects shown in the current study. While the differences in size between adapting and test stimuli make low-level retinotopic visual areas unlikely candidates, aftereffects may originate in any one of many higher-level face-responsive regions in a distributed network (Haxby, Hoffman, Gobbini, 2000; Ishai, Schmidt, Boesiger, 2005; Fairhall & Ishai, 2007; Ishai, 2008). Although some of these areas such as the fusiform face area (Golby, Gabrieli, Chiao, & Eberhardt, 2001; Brosch, Bar-David & Phelps, 2013), occipital face area and occipital pole (Brosch et al., 2013) show modulations of activity when race is manipulated, others such as the superior temporal sulcus, caudate, inferior frontal gyrus, insula and orbitofrontal cortex (Brosch et al., 2013) appear insensitive to race. While we might expect the race-contingent aftereffects reported here to have their neural basis in the race sensitive areas, we should remember the evidence for cross-adaptation (Jaquet et al., 2008). This raises the possibility that the face aftereffects may have multiple loci, for example a race-contingent component in one or more race-sensitive neural modules, and a general overall aftereffect originating from one or more race-insensitive face processing areas. Such a system would imply at least three pools—two that are race-selective, and one that responds to faces regardless of race, and is responsible for the cross-adaptation effect. In this case, our inferences about the shape of the response functions would remain relevant for the race-selective channels, but would not apply to the race-independent channel.

Within the framework of Valentine's (1991) model of face processing, race-contingent aftereffects are explained as the encoding of faces against race-specific prototypes, or "norms" that have been differently adapted in terms of their degree of expansion or contraction. While there seems to be little doubt that the human visual system does dissociate the processing of faces of obviously different races, the gradual change in strength and direction of aftereffects found in the

current study suggests that faces may be simultaneously encoded against multiple prototypes with the degree of similarity to a given prototype determining the level of influence that each prototype has on the final face percept. In regards to race these judgments seem to be primarily based on similarity of face morphology rather than skin tone, at least when considering stimuli that span the continuum between typically Caucasian and Black African faces (Brooks & Gwinn, 2010; Willenbockel, Fiset, & Tanaka, 2011; Gwinn & Brooks, 2013).

While the current research has used a bottom-up approach, making specific and quantifiable manipulations to the stimuli to produce perceptual effects, this is not to deny that top-down processes involving social effects on race categorization also have a role to play. A growing body of evidence indicates that the visual mechanisms involved in face perception are sensitive to social judgments of group membership. Deficits in recognition performance similar to those found when viewing other-race faces (Meissner & Brigham, 2001) can be induced by giving a racially ambiguous face a particular hairstyle as a racial label (Maclin & Malpass, 2001) or simply by arbitrarily assigning same-race faces to social in-groups or out-groups (Bernstein, Young, & Hugenberg, 2007). Similarly, reductions in holistic processing typically found when viewing other-race faces (Michel, Rossion, Han, Chung, & Caldara, 2006) can be induced by designating faces as belonging to out-group members (Hugenberg & Corneille, 2009). It should also be noted that contingent face aftereffects have been demonstrated when socially relevant labels (e.g., introvert/extrovert) accompany colored (blue/red) faces, but not when labels were socially irrelevant (born on Monday/Friday) or when they were absent (Little, DeBruine, & Jones, 2011; see also Yamashita, Hardy, De Valois, & Webster, 2005). Although this is referred to as a category contingent aftereffect, it could be accounted for by either of the models presented above in Figure 1. While labels such as introvert or extrovert could encourage an imbalance in activity in pools more selective for one extreme or the other, it is not known whether the size of the contingent aftereffect would be the same regardless of the degree of extroversion/introversion.

When social influences are minimized and held constant in each condition, such as in the current study, results indicate that the degree to which faces of varying racial appearance are processed by selective mechanisms reflects the extent to which they are perceived as typical of a given race. Rather than being encoded in a purely categorical fashion, faces of different races are best conceptualized as forming a continuum of racial typicality.

Keywords: adaptation, aftereffects, faces, face-space, neural mechanisms, race

Acknowledgments

O. Scott Gwinn was supported by a Macquarie University Research Excellence Scholarship. The authors would like to thank Charlotte Rush and Anthony Espinosa for assisting in data collection.

Commercial relationships: none.

Corresponding author: Owen Scott Gwinn.

Email: scott.gwinn@mq.edu.au.

Address: Department of Psychology, Macquarie University, Sydney, New South Wales, Australia.

Footnotes

¹The Asian faces were images of students from Hong Kong University, available at <http://viscog.hku.hk/index.htm> and used with the permission of William G. Hayward. The Caucasian faces were the same as those used in Russell (2009), and were used with the permission of Richard Russell.

²All curve fitting involved automatic outlier elimination using the ROUT method (Motulsky & Brown, 2006), with parameter $Q = 1$.

³Another alternative suggested by an anonymous reviewer would be to average ratings across morph levels within each participant and then fit individual curves. The average of slopes using this method ($SD = 0.21$) is very similar to the one reported.

⁴We are indebted to an anonymous reviewer for this observation.

References

- Addams, R. (1834). An account of a peculiar optical phenomenon seen after having looked at a moving body. *London and Edinburgh Philosophical Magazine and Journal of Science*, 5, 373–374.
- Barlow, H. B., & Hill, R. M. (1963). Evidence for a physiological explanation of the waterfall illusion. *Nature*, 200, 1345–1347.
- Bernstein, M. J., Young, S. G., & Hugenberg, K. (2007). The cross-category effect: Mere social categorization is sufficient to elicit an own-group bias in face recognition. *Psychological Science*, 18, 706–712.
- Blair, I. V., Judd, C. M., Sadler, M. S., & Jenkins, C. (2002). The role of Afrocentric features in person perception: Judging by features and categories. *Journal of Personality and Social Psychology*, 83, 5–25.

- Brooks, K. R., & Gwinn, O. S. (2010). No role for lightness in the perception of black and white? Simultaneous contrast affects perceived skin tone, but not perceived race. *Perception*, *39*, 1142–1145.
- Brooks, K. R., & Kemp, R. I. (2007). Sensitivity to feature displacement in familiar and unfamiliar faces: Beyond the internal/external feature distinction. *Perception*, *36*, 1646–1659.
- Brosch, T., Bar-David, E., & Phelps, E. A. (2013). Implicit race bias decreases the similarity of neural representations of black and white faces. *Psychological Science*, *24*, 160–166.
- Clifford, C. W., & Rhodes, G. (Eds.). (2005). *Fitting the mind to the world: Adaptation and after-effects in high-level vision (Vol. 2)*. Oxford, UK: Oxford University Press.
- Dakin, S. C., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, *9*(4):2, 1–10, <http://journalofvision.org/content/9/4/2/>, doi:10.1167/9.4.2. [PubMed] [Article]
- Eberhardt, J. L., Davies, P. G., Purdie-Vaughns, V. J., & Johnson, S. L. (2006). Looking deathworthy: Perceived stereotypicality of black defendants predicts capital-sentencing outcomes. *Psychological Science*, *17*, 383–386.
- Fairhall, S. L., & Ishai, A. (2007). Effective connectivity within the distributed cortical network for face perception. *Cerebral Cortex*, *17*, 2400–2406.
- Frisby, J. P. (1980). *Seeing: Illusion, brain and mind*. Oxford, UK: Oxford University Press.
- Golby, A. J., Gabrieli, J. D., Chiao, J. Y., & Eberhardt, J. L. (2001). Differential responses in the fusiform region to same-race and other-race faces. *Nature Neuroscience*, *4*, 845–850.
- Gwinn, O. S., & Brooks, K. R. (2013). No role for lightness in adaptation for Black and White: Race-contingent face aftereffects depend on facial morphology, not skin tone. *Journal of Vision*, *13*(9):860, <http://www.journalofvision.org/content/13/9/860>, doi:10.1167/13.9.860. [Abstract]
- 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., & Corneille, O. (2009). Holistic processing is tuned for in-group faces. *Cognitive Science*, *33*, 1173–1181.
- Ishai, A. (2008). Let’s face it: It’s a cortical network. *NeuroImage*, *40*, 415–419.
- Ishai, A., Schmidt, C. F., & Boesiger, P. (2005). Face perception is mediated by a distributed cortical network. *Brain Research Bulletin*, *67*, 87–93.
- Jaquet, E., & Rhodes, G. (2008). Face aftereffects indicate dissociable, but not distinct, coding of male and female faces. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 101–112.
- Jaquet, E., Rhodes, G., & Hayward, W. G. (2007). Opposite aftereffects for Chinese and Caucasian faces are selective for social category information and not just physical face differences. *The Quarterly Journal of Experimental Psychology*, *11*, 1457–1467.
- Jaquet, E., Rhodes, G., & Hayward, W. G. (2008). Race-contingent aftereffects suggest distinct perceptual norms for different race faces. *Visual Cognition*, *16*, 734–753.
- Jeffery, L., McKone, E., Haynes, R., Firth, E., Pellicano, E., & Rhodes, G. (2010). Four-to-six-year-old children use norm-based coding in face-space. *Journal of Vision*, *10*(5):18, 1–19, <http://journalofvision.org/content/10/5/18>, doi:10.1167/10.5.18. [PubMed] [Article]
- Keil, M. S., Lapedriza, A., Masip, D., & Vitria, J. (2008). Preferred spatial frequencies for human face processing are associated with optimal class discrimination in the machine. *PLoS ONE*, *3*, e2590. doi:10.1371/journal.pone.0002590.
- Leopold, D. A., O’Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, *4*, 89–94.
- Levin, D. T., & Beale, J. M. (2000). Categorical perception occurs in newly learned faces, other-race faces, and inverted faces. *Perception & Psychophysics*, *62*, 386–401.
- Little, A. C., DeBruine, L. M., & Jones, B. C. (2011). Category-contingent face adaptation for novel colour categories: Contingent effects are seen only after social or meaningful labelling. *Cognition*, *118*, 116–122.
- Little, A. C., DeBruine, L. M., & Jones, B. C. (2005). Sex-contingent aftereffects suggest distinct neural populations code male and female faces. *Proceedings of The Royal Society B: Biological Sciences*, *272*, 2283–2287.
- Little, A. C., DeBruine, L. M., Jones, B. C., & Waite, C. (2008). Category contingent aftereffects for faces of different races, ages and species. *Cognition*, *106*, 1537–1547.
- MacLin, O. H., MacLin, M. K., Peterson, D., Chowdhry, O., & Joshi, P. (2009). Social psychophysics: Using psychophysics to answer “social” questions with PsychoPro. *Behavior Research Methods*, *41*, 623–632.
- MacLin, O. H., & Malpass, R. S. (2001). Racial

- categorization of faces. *Psychology, Public Policy & Law*, 7, 98–118.
- McCollough, C. (1965). Adaptation of edge-detectors in the human visual system. *Science*, 149, 1115–1116.
- Meissner, C., & Brigham, J. (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.
- Michel, C., Rossion, B., Han, J., Chung, C. S., & Caldara, R. (2006). Holistic processing is finely tuned for faces of one's own race. *Psychological Science*, 17, 608–615.
- Motulsky, H. M., & Brown, R. E. (2006). Detecting outliers when fitting data with nonlinear regression—A new method based on robust nonlinear regression and the false discovery rate. *BMC Bioinformatics*, 7, 123.
- Robbins, R., McKone, E., & Edwards, M. (2007). Aftereffects for face attributes with different natural variability: Adaptor position effects and neural models. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 570–592.
- Russell, R. (2009). A sex difference in facial pigmentation and its exaggeration by cosmetics. *Perception*, 38, 1211–1219.
- Storrs, K. R., & Arnold, D. H. (2012). Not all face aftereffects are equal. *Vision Research*, 64, 7–16.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 43, 161–204.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428(6982), 557–561.
- Webster, M. A., & Macleod, D. I. A. (2011). Visual adaptation and face perception. *Philosophical Transaction of The Royal Society B: Biological Sciences*, 366, 1702–1725.
- Webster, M. A., & MacLin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin and Review*, 6, 647–653.
- Willenbockel, V., Fiset, D., & Tanaka, J. W. (2011). Relative influences of lightness and facial morphology on perceived race. *Perception*, 40, 621–624.
- Yamashita, J. A., Hardy, J. L., De Valois, K. K., & Webster, M. A. (2005). Stimulus selectivity of figural aftereffects for faces. *Journal of Experimental Psychology-Human Perception and Performance*, 31, 420–437.