



Face-selective regions show invariance to linear, but not to non-linear, changes in facial images

Heidi A. Baseler, Andrew W. Young, Rob Jenkins, A. Mike Burton, Timothy J. Andrews*

Department of Psychology and York Neuroimaging Centre, University of York, York YO10 5DD, United Kingdom

ARTICLE INFO

Keywords:

FMRI
FFA
Face
Adaptation
Shape
Caricature

ABSTRACT

Familiar face recognition is remarkably invariant across huge image differences, yet little is understood concerning how image-invariant recognition is achieved. To investigate the neural correlates of invariance, we localized the core face-responsive regions and then compared the pattern of fMR-adaptation to different stimulus transformations in each region to behavioural data demonstrating the impact of the same transformations on familiar face recognition. In Experiment 1, we compared linear transformations of size and aspect ratio to a non-linear transformation affecting only part of the face. We found that adaptation to facial identity in face-selective regions showed invariance to linear changes, but there was no invariance to non-linear changes. In Experiment 2, we measured the sensitivity to non-linear changes that fell within the normal range of variation across face images. We found no adaptation to facial identity for any of the non-linear changes in the image, including to faces that varied in different levels of caricature. These results show a compelling difference in the sensitivity to linear compared to non-linear image changes in face-selective regions of the human brain that is only partially consistent with their effect on behavioural judgements of identity. We conclude that while regions such as the FFA may well be *involved* in the recognition of face identity, they are more likely to contribute to some form of normalisation that underpins subsequent recognition than to form the neural substrate of recognition per se.

1. Introduction

A striking property of our ability to recognise familiar faces is that it remains relatively invariant across huge changes in the facial image. Substantial image differences resulting from pose, expression and lighting have little or no effect on familiar face recognition (Hill and Bruce, 1996; O'Toole et al., 1998; Patterson and Baddeley, 1977), even when image quality is severely impoverished (Burton et al., 1999). A critical research question is therefore how this high degree of invariance in responses to very different images is achieved. A useful way to probe the mechanisms underlying our image-invariant responses to familiar faces is to look at transformations of the face image that permit or interfere with recognition. Such transformations can be broadly grouped into those involving linear or non-linear changes.

The evidence for linear transformations is clear. Changing the aspect ratio of a familiar face photograph by stretching or compressing it in the horizontal or vertical dimension has almost no measurable effect on our ability to recognise it (Hole et al., 2002; Sandford and Burton, 2014). The same is true for linear changes affecting both horizontal and vertical dimensions at once, as for a size change. These

linear changes in the image often occur with changes in viewing distance, so some degree of invariance to such changes is useful to everyday recognition. Interestingly, however, invariance to changes in aspect ratio is nonetheless found for changes that fall well outside the limits participants will have encountered (Hole et al., 2002; Sandford and Burton, 2014).

For non-linear transformations the evidence is more mixed. Most of the changes resulting from the facial movements and expressions we make nearly all the time are non-linear, yet they have little impact on familiar face recognition. Furthermore, image non-linearities due to lens properties, such as focal length and perspective, appear to have no effect on viewers' perceptions of familiar faces (Burton et al., 2015; Harper and Latto, 2001). In some cases, it seems that non-linear transformations can *improve* face recognition, as in facial caricaturing (Rhodes et al., 1987, 1996), though evidence for this benefit is strongest when stimuli are degraded in some way (Benson and Perrett, 1991, 1994; Calder et al., 1996; Rhodes et al., 1987). Nevertheless, it is clear that *some* non-linear changes are detrimental to familiar face recognition. Hole et al. (2002) showed that linearly stretching one half of a face image and leaving the remaining half

* Corresponding author.

E-mail address: timothy.andrews@york.ac.uk (T.J. Andrews).

unaltered led to a clear decrement in face recognition. This is rather a non-naturalistic transformation, because it is applied only to part of the face. However, Hole's demonstration shows that some transformations do affect computation of a face's identity – and provides a potential means to investigate the nature of familiar face representations.

The effect of image transformation on face perception has also been examined in ERP studies of face recognition. N250r is an ERP component which has been linked to the representation of facial identity (Schweinberger et al., 2002; Pickering and Schweinberger, 2003). This component is interesting because it shows *both* image-independent sensitivity (successive presentations of different images of a known person affect the response) *and* image-dependent sensitivity (response is maximal if the identical image is repeated). Interestingly, the N250r component is unaffected by linear changes to facial stimuli between repetitions (Bindemann et al., 2008) illustrating an image-independent facial response across these transformations. While this result suggests that some neural responses occur independently of these image transformations, it does not provide evidence about the structures involved.

In the present study, we draw on these effects of linear and non-linear image manipulation to probe the responses of brain regions thought to be critically involved in face perception and recognition. Neuroimaging studies have revealed a network of regions in the occipital and temporal lobes which form a core system for the visual analysis of faces (Sergent and Signoret, 1992; Allison et al., 1994; Kanwisher et al., 1997; Haxby et al., 2000). The widely used neural model by Haxby et al. (2000) identifies the occipital face area (OFA), the fusiform face area (FFA), and the posterior superior temporal sulcus (STS). Within the core face-selective regions, the FFA is thought to be particularly important for the representation of invariant facial characteristics that are important for face recognition (Haxby et al., 2000; Grill-Spector et al., 2004). Consistent with the role of the FFA in processing facial identity, previous fMRI studies have shown a reduced response (adaptation) to repeated images of the same face in the FFA (Grill-Spector et al., 1999; Andrews and Ewbank, 2004). These findings

imply that the identity of the face is represented at some level in the FFA and this representation is being adapted by repeated presentations. However, a stronger test for a direct link between neural activity in a given region and the recognition of facial identity is to determine whether the *pattern* of adaptation in that region corresponds to the behavioural pattern of recognition of identity. That is, whether adaptation is still evident when images of the same identity are transformed along a dimension that is not important for face recognition and release from adaptation is found when the images are transformed in a way that affects recognition.

To achieve these aims, we used functional localiser scans to identify the core face-responsive regions and then compared the pattern of neural adaptation to different stimulus transformations in each region to behavioural data demonstrating the impact of the same transformations on familiar face recognition. In Experiment 1, we compared linear transformations of size and aspect ratio to a non-linear transformation affecting only part of the face (see Hole et al. (2002)). In Experiment 2, we investigated the effects of non-linear transformations of face shape that fell within the range encountered in everyday life. Our prediction was that manipulations of the face that do not affect recognition will not influence adaptation to identity in face regions involved in recognition, whereas manipulations that do affect recognition will reduce the magnitude of adaptation.

2. Methods

2.1. Participants

The study was approved by the York Neuroimaging Centre (YNiC) Research Ethics Committee. Informed written consent was obtained from all participants before each experiment, and all participants had normal or corrected-to-normal vision. In addition, a screening test was used to ensure familiarity with the identities of the familiar faces. Images of the identities used in the screening test were different from the images used in the actual experiments. In Experiment 1, 21

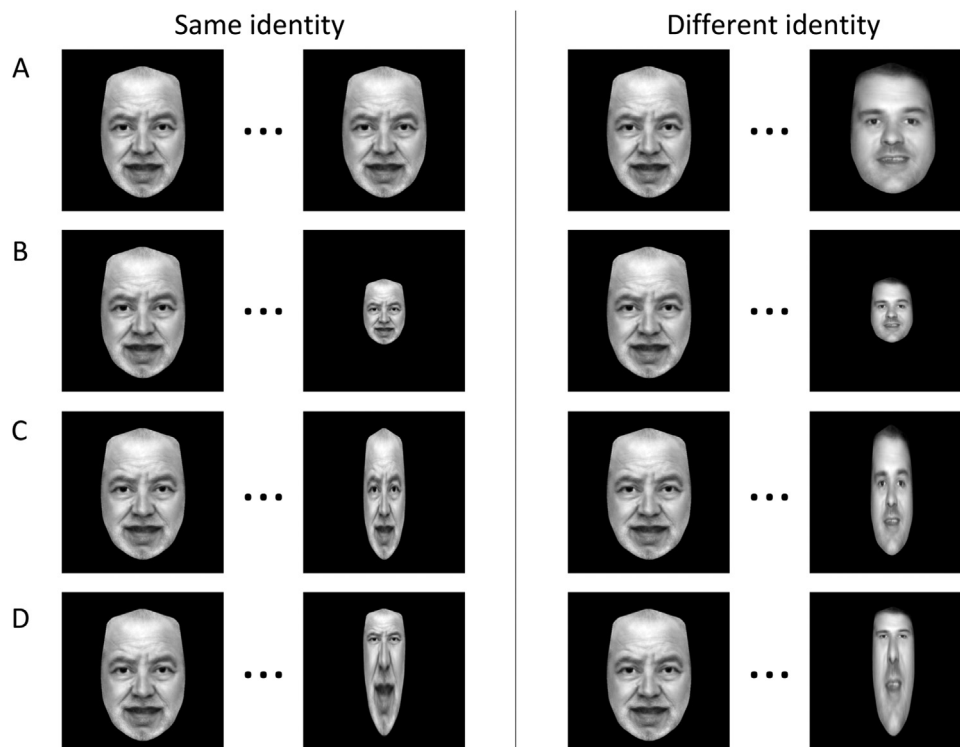


Fig. 1. Examples of familiar face stimuli used in Experiment 1. Examples of same identity trials are shown on the left, different identity trials shown on the right. (A) no change, (B) size change, (C) one-dimensional linear change, (D) non-linear change.

participants (female=13; mean age=25.8, SE=1.5) took part in the behavioural experiment and 23 participants took part in the fMRI experiment (female=15; mean age=22.1, SE=1.4). In Experiment 2, 25 participants took part in the fMRI experiment (female=14; mean age=25.8, SE=1.7).

2.2. Experiment 1

Stimuli consisted of averaged faces of British celebrities that were familiar to the participants (Chris Moyles, Gary Lineker, Jeremy Paxman, Louis Walsh, Jonathan Ross, Alan Sugar). Stimuli were generated by averaging 12 different greyscale images of each identity (see Burton et al. (2005) and Jenkins and Burton (2011), for full details of this procedure). The reason for using averaged images is that averaging gets rid of variation in lighting, picture quality, camera characteristics and so on that would otherwise introduce identity-irrelevant low-level variation.

Faces in the original configuration subtended $\sim 6 \times 8^\circ$. The average face for each identity was then changed by applying various linear and non-linear spatial transformations, as illustrated in Fig. 1. Images for Experiment 1 underwent 4 different spatial transformations: (A) no change, (B) size change (an equal linear change of 50% reduction in image height and width), (C) a one-dimensional linear change (image width reduction of 50%, with no change in height) or (D) non-linear transformation (top half of face unchanged, bottom half stretched vertically by 100%, then the resulting image resized to its original height so that the overall height and width were the same as for the

one-dimensional linear changes – see Fig. 1).

A behavioural study was performed to measure the ability of participants to discriminate face identities across these four different spatial transformations. Examples of trials are shown in Fig. 1. In each trial, two face images (the first of which was always an untransformed average image) were presented sequentially for 900 ms each with a 100 ms inter-stimulus interval, and participants pressed a button to indicate whether the two faces were of the same identity or different identities. The eight types of trials (conditions - same or different face identity pairs across four types of transform) were presented in random order, but the same order was used for all participants. Participants viewed a total of 168 trials (21 trials per condition); the first 24 trials were considered practice trials (3 trials per condition) and not included in the final analysis. For each participant, the median reaction time was calculated for correct responses for each condition, and the mean and standard error of the median reaction times were calculated across participants. Mean accuracy rates were also calculated for each participant and condition. Responses were considered correct or incorrect based on the identities of the faces from which the transformed stimuli were created. Participants viewed face images at a distance of 57 cm.

An fMR-adaptation paradigm was then used to measure neural responses to the different spatial transformations. Images were presented in a blocked design. During each block, 6 images were presented in sequence, each for 900 ms with a 100 ms black screen between each image, followed by a 9 s gray mean luminance screen with a fixation cross, resulting in a total block length of 15 s. The same four image



Fig. 2. Examples of stimuli from Experiment 2. There were 6 image conditions (A) no change, (B) caricature, (C) shape change, (D) different identity (3), (E) different identity (6).

transformations were used as in the behavioural experiment (unchanged, size change, one-dimensional linear change and non-linear change), but to create variation between images in a block sequence an additional (75%) transform was included that fell within the range used in the behavioural experiment. In the unchanged condition, images in each block were shown at the original size. In the size change condition, the size of the images in each block was shown at 50, 75 or 100% of the original image dimensions. In the linear stretch condition, the height of the images in each block was unchanged, but the width of images was 50, 75 or 100% of the original image. In the non-linear condition, the size of the image was 50, 75 or 100% of the original image, then the bottom half of the image was stretched and the image resized so that the height of the face was the same as the original face. Participants viewed face images at a distance of 57 cm.

For each of the four transformation types, identity was either held constant (same identity), or varied (different identity) throughout the block. This gave 8 different stimulus conditions. Each condition was repeated in six blocks, and conditions and blocks were presented in pseudorandom order (i.e. the condition order was randomised, but held constant across participants). To maintain attention throughout the stimulus presentation, participants were required to press a button when detecting a small, red dot (0.1° visual angle), which appeared in a random position on one or two of the gray scale face images within each block.

2.3. Experiment 2

Experiment 2 investigated behavioural and neural sensitivity to non-linear changes in spatial configuration that fall within the normal range of variation. Images were derived from the original averaged images used in Experiment 1. These images underwent two types of non-linear transformations: shape change (by mapping the average surface properties of each identity onto the average spatial position of facial features for a different identity) and caricature (exaggerating individual spatial feature position differences from an average position of these features across the faces).

The shape change transformations were achieved by defining 82 key fiducial points (e.g., corners of the mouth, of the eyes, etc.) for each image in a graphics program. These fiducial positions were then connected to form a grid representing the shape of the image (for details see Burton et al. (2015)). This allowed the surface properties of each face to be mapped onto the spatial fiducial positions of a different face (Burton et al., 2005). For the caricature condition, an average spatial configuration was calculated from the fiducial positions of the 6 original images. Two levels of caricature were then generated for each original image by exaggerating the deviation from the average face by 25% and 50%. These two levels were chosen such that the average shape deviation for caricatures within identities was equivalent to the average shape deviation between images in the shape condition. In this way, we ensured that the changes in feature positions for each type of transform fell within the range of variability of the fiducial locations in the original set of images. As is clear in Fig. 2, shape or caricature changes within this range (rows B and C in Fig. 2) have little impact on the perception of identity of familiar faces (that is, the images in rows B and C of Fig. 2 all look like the same person). This is consistent with previous studies of such transforms (Burton et al., 2005).

A behavioural study was performed to measure the ability of participants to discriminate face identities across these four different spatial transformations. In each trial, two face images (the first of which was always an untransformed average image) were presented sequentially for 900 ms each with a 100 ms inter-stimulus interval, and participants pressed a button to indicate whether the two faces were of the same identity or different identities. The six types of trials (conditions - same or different face identity pairs across three types of transform - no change, shape change and caricature) were presented in random order, but the same order was used for all participants.

Participants viewed a total of 168 trials (28 trials per condition); the first 24 trials were considered practice trials (4 trials per condition) and not included in the final analysis. For each participant, the median reaction time was calculated for correct responses for each condition, and the mean and standard error of the median reaction times were calculated across participants. Mean accuracy rates were also calculated for each participant and condition. Responses were considered correct or incorrect based on the identities of the faces from which the transformed stimuli were created. Participants viewed face images at a distance of 57 cm.

Next, we investigated the sensitivity of face-selective regions to these non-linear changes in face image. Images were presented in a blocked fMRI-adaptation design. During each block, 6 images were presented in sequence, each for 900 ms with a 100 ms black screen between each image, followed by a 9 s gray mean luminance screen with a fixation cross, resulting in a total block length of 15 s. There were five experimental conditions: (A) no change (same identity, same image); (B) caricature (two examples at each of 0%, 25%, 50% caricature); (C) shape change (two examples of each of 3 shapes, original shape and two others); (D) identity change (3); (E) identity change (6). The identity change (3) condition had three identities repeated twice. This matched the level of image repetition in the caricature and shape change conditions. We also included the identity change (6) condition, which had all six identities shown once. Our rationale for including this condition was to measure the maximal release from adaptation to identity. Examples of face blocks used in Experiment 2 are shown in Fig. 2. A red dot detection task was used to maintain participants' attention throughout the scan. Participants viewed face images at a distance of 57 cm.

2.4. fMRI analysis

Data from the fMRI scans were collected using a GE 3 Tesla HD Excite MRI scanner at the York Neuroimaging Centre at the University of York. A T1-weighted structural MRI (1×1.13×1.13 mm voxel) and a gradient-echo EPI were acquired for each participant. A gradient-echo EPI sequence with a radio-frequency coil tuned to 127.4 MHz was used to acquire 38 axial slices (TR=3 s, TE=33 ms, flip angle=90°, FOV=260 mm, matrix size=128×128, slice thickness=3 mm, voxel size: 2.25×2.25×3 mm). Data were analysed with FEAT version 4.1 (<http://www.fmrib.ox.ac.uk/fsl>). The first 9 s (3 volumes) from each scan were discarded, and MCFLIRT motion correction, spatial smoothing (Gaussian, FWHM 6 mm), and temporal high-pass filtering (cutoff 0.0093 Hz) were applied. The BOLD response for each condition was modelled with a boxcar function convolved with a standard haemodynamic response function.

In both experiments, a separate scan was run to localise face selective regions within the brain of each individual participant. This localiser used a block design with 7 different conditions: same identity faces, different identity faces, bodies, inanimate objects, places, and scrambled images. Each block consisted of 10 images from each condition; each image was presented for 700 ms and followed by a 200-ms blank screen. Stimulus blocks were separated by a 9-s gray screen with a central fixation cross. Each condition was repeated 4 times in a counterbalanced design. Face-selective regions were defined in each individual using the localiser scan by contrasting the response to faces with the response to each non-face condition. These statistical maps were averaged and thresholded at $p < 0.001$ (uncorrected). Contiguous clusters of voxels located within the occipital and temporal lobes were defined as the FFA, OFA and posterior STS in each participant (Supplementary Fig. 1). The MNI coordinates of the centre of gravity of each region of interest were determined for the face localiser scans for each participant. Supplementary Tables 1–2 lists the mean MNI coordinate positions across participants for Experiments 1 & 2.

To analyse the data from the experimental scans, the time-course of

response from each voxel within each ROI was converted from units of image intensity to percentage signal change. Voxels were then averaged to give a single time series for each ROI in each participant. Individual stimulus blocks were corrected to the same baseline by subtracting the response at the start of the block from the response at every time point and then averaged to produce a mean time series for each condition for each participant. The peak response was at 9 s post-stimulus onset. fMR-adaptation to identity was determined for each condition by comparing the same identity and different identity conditions for each image manipulation.

3. Results

3.1. Experiment 1

A behavioural experiment was used to determine the effect of linear and non-linear manipulations on the ability to discriminate facial identity by sequentially presenting pairs of images derived from the same or a different identity (see Fig. 1). Fig. 3 shows the mean accuracy and reaction time on same identity and different identity trials. A repeated measures ANOVA performed with two factors of Identity (same, different) and Transformation (no change, size change, one-dimensional linear change, non-linear change) revealed significant main effects of transformation on both accuracy ($F(3,60)=5.27$, $p < 0.005$) and reaction time ($F(3,60)=10.50$, $p < 0.001$), and a significant interaction between identity and transformation for both accuracy ($F(3,60)=3.18$, $p < 0.05$) and reaction time ($F(3,60)=4.71$, $p < 0.005$). Post-hoc paired t -tests were performed on behavioural data, using a Holm-Bonferroni correction for multiple comparisons. On same identity trials, participants performed significantly worse on the non-linear trials compared to the other conditions. Reaction times were longer for non-linear transformations compared to all other conditions (no change: $t(20)=4.18$, $p < 0.001$; non-linear: $t(20)=3.52$, $p=0.002$; linear: $t(20)=4.37$, $p < 0.001$). Participants were also significantly less accurate when making judgments about non-linearly transformed faces (no change: $t(20)=3.85$, $p=0.001$; linear: $t(20)=2.97$, $p=0.008$). On different identity trials, there was no effect of spatial transformation on reaction time or accuracy. These findings show that participants were less accurate and took more time on trials comparing original images with images that had been transformed non-linearly if they had the same identity.

An fMR-adaptation paradigm was then used to determine the sensitivity of the core face-selective regions (see Supplementary Fig. 1 and Tables 1 & 2) to linear and non-linear spatial transformations. There was no significant difference between left and right hemisphere responses in either the FFA or OFA, so the responses from regions in the two hemispheres were averaged together. The posterior STS data represents responses from the right posterior STS only, as the left posterior STS could not be reliably identified in the majority of participants. Fig. 4 shows examples of stimulus blocks used in the fMRI scan and the average time-course of response in the core face-selective regions to each spatial transformation for same identity or different identity images.

A 2-way ANOVA with Identity (same, different) and Transformation (no change, size change, unidimensional linear change, non-linear change) was performed for each region of interest. In the FFA, there were significant main effects for identity ($F(1,22)=55.3$, $p < 0.001$) and transformation ($F(3,66)=5.81$, $p < 0.001$) as well as a significant interaction ($F(3,66)=6.12$, $p < 0.001$). In the OFA, there was a significant main effect of identity ($F(1,21)=28.4$, $p < 0.001$), no effect of transformation, but a significant interaction between transformation and identity ($F(3,63)=3.11$, $p=0.033$). There were no significant effects in the STS.

To determine whether adaptation to facial identity was evident for each spatial transformation, responses to same and different identity conditions were compared within each ROI with post-hoc t -tests, using

a Holm-Bonferroni correction for multiple comparisons. The FFA and OFA showed significant adaptation (different identity > same identity) for the no change, size change and one-dimensional linear change transformations (FFA – no change: $t(22)=5.91$, $p < 0.001$; size: $t(22)=5.97$, $p < 0.001$; linear: $t(22)=3.61$, $p=0.002$; OFA – no change: $t(21)=3.48$, $p=0.002$; size: $t(21)=4.27$, $p < 0.001$; linear: $t(21)=2.85$, $p=0.010$). Only the non-linearly transformed faces produced a complete release from adaptation (i.e. there was no difference between same and different identities) in the FFA and OFA. No significant differences between responses to same and different identity faces were found in the STS.

Previous studies have defined a region in the anterior temporal lobe (aTL) that is selective for faces (Rajimehr et al., 2009). Although we were unable to define this region at the individual participant level, it was possible to locate a region from a group analysis (56, –12, –16). The coordinates of this region were back-projected to individual participants' data to determine the response in this region to conditions in Experiment 1. An ANOVA of the responses in this region failed to show any effect of Identity ($F(1,22)=0.15$, $p=0.70$) or any interaction between Identity and Transformation ($F(2.17,47.8)=0.59$, $p=0.57$).

Finally, we measured responses in two control regions. The parahippocampal place area (PPA) was defined by contrasting the response to places and faces. This region is located near to the FFA, but is not

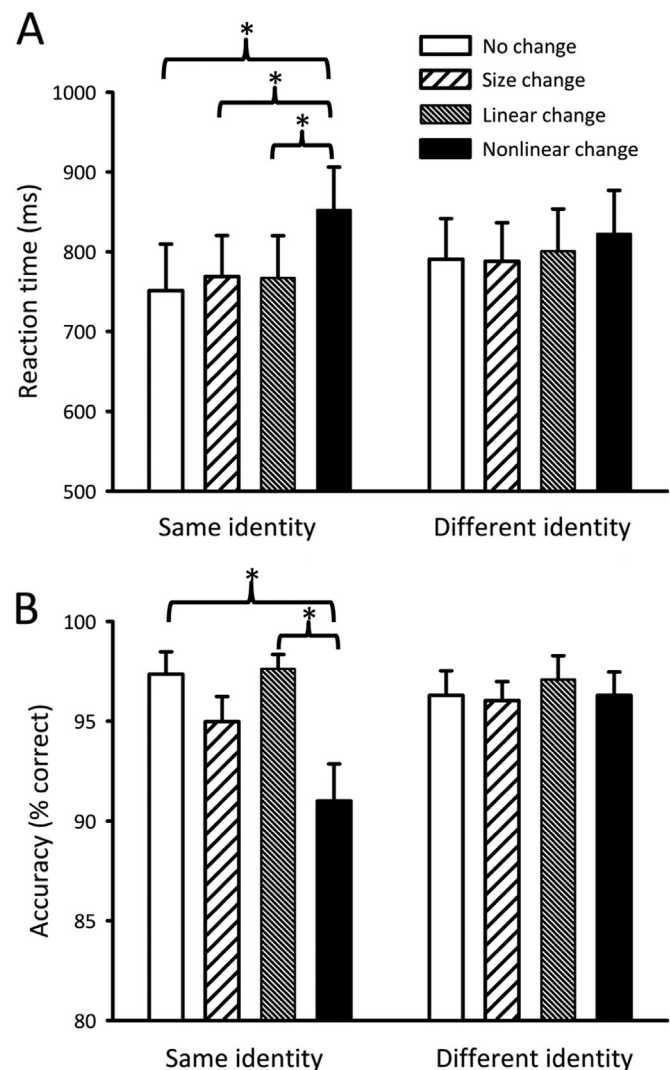


Fig. 3. Results of behavioural responses from Experiment 1. (A) Median reaction time, averaged across participants ($N=21$) and (B) mean accuracy across all participants. Error bars represent ± 1 standard error of the mean. * $p < 0.01$.

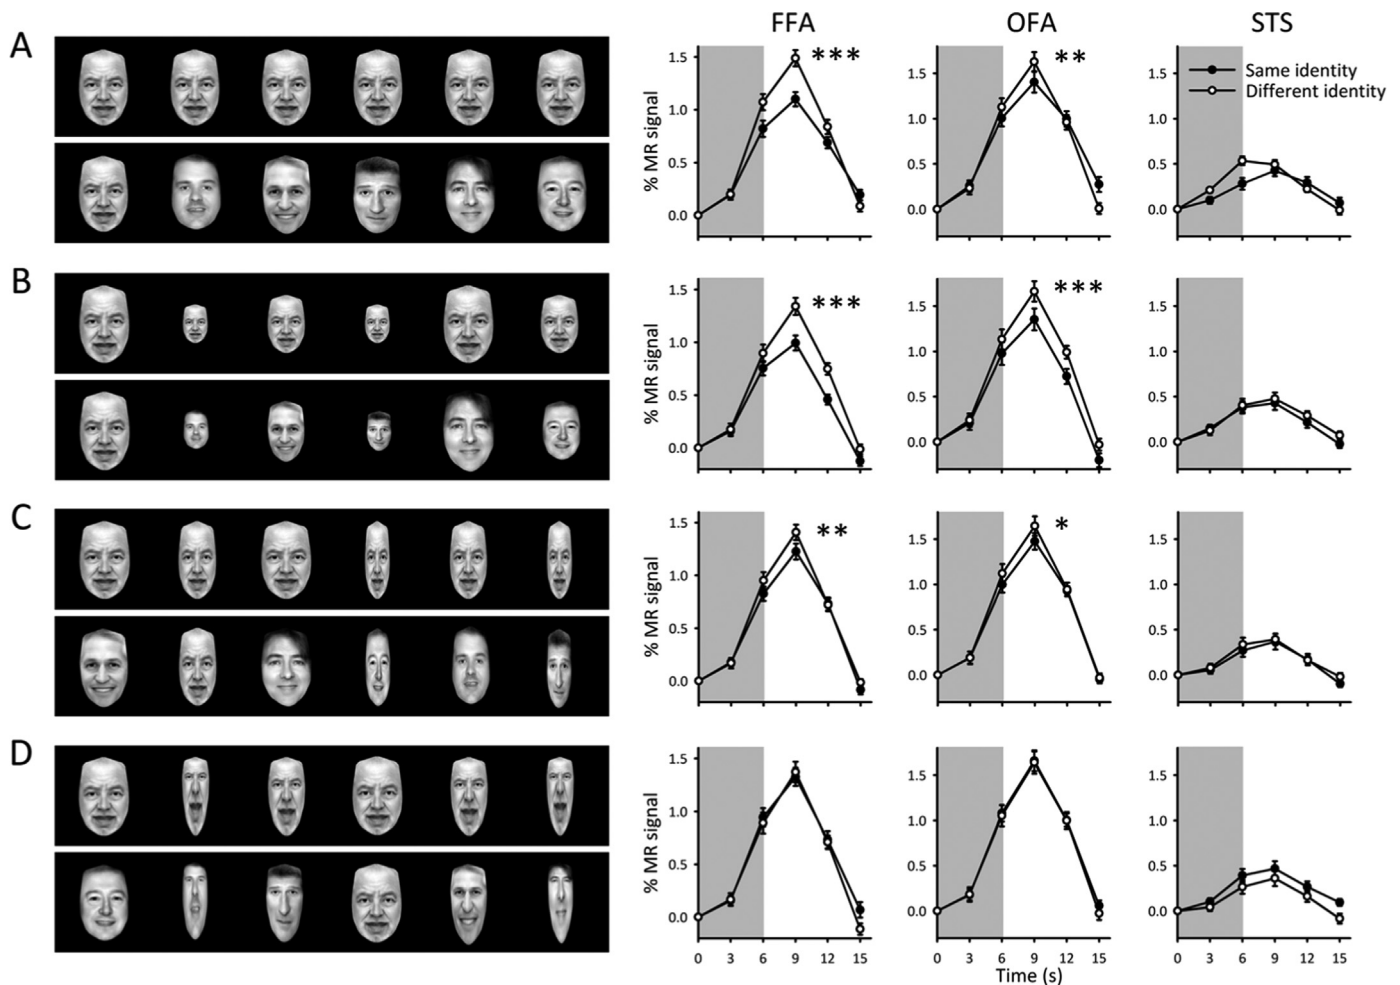


Fig. 4. Experiment 1 – fMRI responses in face-selective regions to different spatial transformations (A – no change, B – size change, C – one-dimensional linear change, D – non-linear change) in which identity was either kept constant (same) or varied (different). fMRI responses for each of the conditions in three face-selective regions are shown for same identity faces (closed circles) and different identity faces (open circles). Responses are averaged across repetition blocks and across participants. Error bars represent ± 1 standard error of the mean response across participants. Face stimuli were presented during the first 6 s of the block (shaded gray region). Comparing peak response amplitude between same and different identity faces for each condition: *** $p < 0.001$; ** $p < 0.005$; * $p < 0.01$.

selective for faces. The PPA showed no effect of Identity ($F(1,22)=0.35$, $p=0.56$) nor any interaction between Identity and Transformation ($F(1.8,40.6)=2.7$, $p=0.08$) in Experiment 1. We also used an anatomical mask (occipital pole – Harvard Oxford Atlas) to measure the responses in early visual areas. The occipital pole showed no effect of Identity ($F(1,22)=0.29$, $p=0.59$), nor any interaction between Identity and Transformation ($F(3,66)=0.82$, $p=0.49$).

3.2. Experiment 2

A behavioural experiment was used to determine the effect of small non-linear manipulations on the ability to discriminate facial identity by sequentially presenting pairs of images derived from the same or a different identity. Fig. 5 shows the mean accuracy and reaction time on same identity and different identity trials. A repeated measures ANOVA performed with two factors of Identity (same, different) and Transformation (no change, size change, one-dimensional linear change, non-linear change) revealed significant main effects of Transformation in both accuracy ($F(1.4,27.4)=4.58$, $p=0.03$) and reaction time ($F(1.5,29.5)=11.09$, $p=0.001$), and a significant interaction between Identity and Transformation in both accuracy ($F(1.5,28.0)=7.44$, $p=0.005$) and reaction time ($F(2,38)=10.23$, $p < 0.001$). The interactions between Identity and Transformation were due to lower accuracy and slower reaction times for the shape change compared to the no change (accuracy: $t(19)=3.29$, $p=0.004$; RT: $t(19)$

$=-5.26$, $p < 0.001$) or caricature (accuracy: $t(19)=-2.71$, $p=0.014$; RT: $t(19)=3.04$, $p=0.007$) in the same identity trials. There was no difference in accuracy between the no change and caricature conditions. There was, however, a slightly faster RT for the caricature compared to the no change ($t(19)=2.73$, $p=0.013$) and shape change ($t(19)=2.82$, $p=0.011$) conditions in the different identity trials.

The aim of Experiment 2 was to use fMR-adaptation to determine the sensitivity of face-selective regions to non-linear changes in spatial feature positions that are within the normal range of facial variation. There were five experimental conditions (see Fig. 2): (A) no change (same identity, same image); (B) caricature; (C) shape change; (D) identity change (3); (E) identity change (6). There was no significant difference between left and right hemisphere responses in either the FFA or OFA, so the responses of regions from the two hemispheres were averaged together. Again, the posterior STS data represents responses from the right posterior STS only, as the left posterior STS could not be reliably identified in all participants.

Fig. 6 shows the peak response to each condition in the core face-selective regions. One-way repeated measures ANOVAs indicated a significant effect of Transformation in the FFA ($F(4,92)=4.93$, $p=0.001$) and STS ($F(4,84)=3.04$, $p=0.022$), but not in the OFA. Next, release from adaptation was measured by comparing responses in each condition to the no-change condition within each ROI with post-hoc t -tests, using a Holm-Bonferroni correction for multiple comparisons. Within the FFA, there was a significantly greater

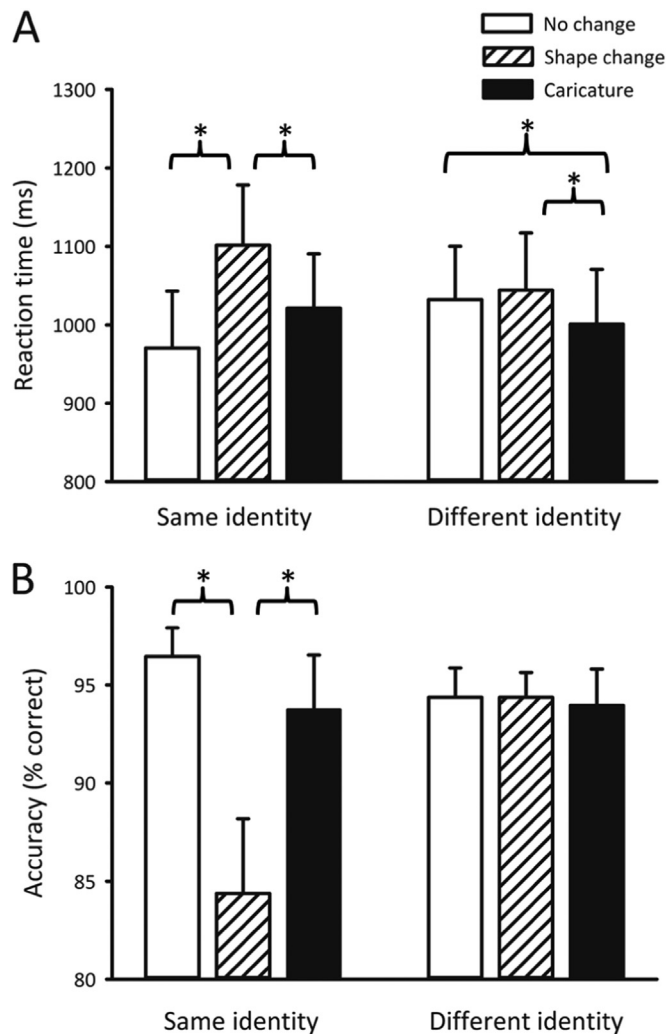


Fig. 5. Results of behavioural responses from Experiment 2. (A) Median reaction time, averaged across participants ($N=20$) and (B) mean accuracy across all participants. Error bars represent ± 1 standard error of the mean. * $p < 0.05$.




Fig. 6. Experiment 2 – Peak amplitude fMRI response, averaged across participants for three face-selective regions of interest. There were 5 different conditions: No change, Caricature, Shape change, Identity change (3) and Identity change (6). Examples of stimuli are illustrated in Fig. 2. Comparing peak response amplitude between No change and all other conditions: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

response (i.e. release from adaptation) for all conditions compared to the no change condition (caricature: $t(23)=2.65$, $p=0.014$; shape change: $t(23)=4.13$, $p < 0.001$; identity change (3): $t(23)=2.41$, $p=0.024$; identity change (6): $t(23)=3.52$, $p=0.002$). There were no

other differences between any of the other conditions in the FFA.

In the STS, there were no differences between any of the conditions and the no change condition. However, there was a larger response to the shape change condition compared to the different identity conditions (identity change (6): $t(21)=2.69$, $p=0.014$; identity change (3): $t(21)=2.48$, $p=0.022$) and between the caricature and the different identity conditions (identity change (6): $t(21)=2.53$, $p=0.019$; identity change (3): $t(21)=1.90$, $p=0.072$).

Next, we measured the responses in the aTL face-selective region defined by our group analysis. An ANOVA on this region failed to show a significant effect of Transformation ($F(2.8,67.2)=0.38$, $p=0.82$).

Finally, we measured responses in two control regions. The place-selective PPA showed no effect of Transformation ($F(4,96)=1.21$, $p=0.31$). The occipital pole showed an effect of Condition ($F(3,72)=2.99$, $p=0.04$). This was due to smaller responses in the shape change condition compared to the caricature ($t(24)=2.35$, $p=0.027$) and identity(3) ($t(24)=3.25$, $p < 0.005$) conditions and smaller responses to the identity(6) condition compared to the caricature ($t(24)=2.78$, $p=0.01$) and identity(3) ($t(24)=2.50$, $p=0.02$) condition. However, it is important to note that in contrast to the FFA, there was no adaptation to the no change condition which was not significantly different to the shape ($t(24)=-1.86$, $p=0.07$), caricature ($t(24)=0.32$, $p=0.75$), identity(3) ($t(24)=-0.68$, $p=0.50$) and identity(6) ($t(24)=-1.27$, $p=0.22$) conditions.

4. Discussion

The aim of this study was to investigate the sensitivity of face-selective regions to linear and non-linear changes in face images. We found adaptation to identity in face-selective regions such as the OFA and FFA, which showed invariance to linear changes in the spatial configuration of the image. This held both for size changes and for the more distorting one-dimensional linear changes that fell well outside the range of everyday image variability. In contrast, we found that adaptation to identity was not invariant to non-linear changes in the spatial configuration of the image, even when these changes fell within the normal range of variation across face images. In contrast to the OFA and FFA, there was no adaptation to facial identity in the STS. However, the STS was sensitive to non-linear changes that fell within the normal range of image variability.

4.1. The effect of linear and non-linear changes on face recognition

Global linear changes in the image can occur in a range of natural viewing conditions. For example, two-dimensional (2D) linear transformations (i.e., size changes) of faces in the real world typically signal changes in distance from the viewer. One-dimensional linear changes (i.e., aspect ratio) can also occur in the real world when we look at an incorrectly reproduced photograph or badly adjusted television. In such circumstances, the incorrect aspect ratio is often noticed, but the face is still easily recognised. Similarly, we found that participants were able to match the identity of two images derived from the same face when the spatial configuration of the face was changed linearly. This was true not only for two-dimensional linear transformations, but also when the image was distorted by one-dimensional linear transformations. These findings are consistent with previous studies, both in adults (Hole, 2002; Frowd et al., 2014) and infants (Yamashita et al., 2013), as well as non-human primates (Taubert and Parr, 2011). However, when the face shape was changed by substantial non-linear transformation in Experiment 1, both accuracy and reaction times were attenuated.

4.2. Neural response to linear and non-linear changes in the face

The fMRI experiment built on these behavioural results by using the logic that a neural region that is directly responsible for the

recognition of facial identity should show a corresponding differential sensitivity to non-linear over linear transformations. This was achieved with a block design fMR-adaptation paradigm in which neural responses to faces with the same identity were compared to faces with different identity. Consistent with our behavioural findings, we found adaptation (lower response to same identity faces) in the FFA and OFA was tolerant of linear changes in the image. For example, adaptation to faces that changed in size was very similar to adaptation to faces that were unchanged in size (see also, Grill-Spector et al. (1999) and Andrews and Ewbank (2004)). We also found significant adaptation to facial identity when a one-dimensional linear transformation was applied to the images, so that they varied in width. However, the magnitude of this adaptation was smaller than for unchanged images suggesting that the responses in these regions are not fully invariant to this manipulation. Nonetheless, these results are consistent with a previous EEG study that also showed that adaptation of the N250 potential to identity was invariant to one-dimensional linear transformations of the images (Bindemann et al., 2008).

In marked contrast to the linear transformations, there was no adaptation to identity in any of the face regions when non-linear vertical transformations were applied to the images in Experiment 1. That is, responses to non-linearly transformed faces of the same identity were equivalent to those when identity was varied, indicating a complete release from adaptation. This release from adaptation in the OFA and FFA to large non-linear manipulations of faces shows strong similarities with the impairment in identity judgements shown in the behavioural findings. In this respect, the neural responses of FFA and OFA parallel the behavioural findings. However, the non-linear distortions used in Experiment 1 were substantial and outside the range of everyday variability. To look more carefully at the responses of these core regions, we therefore used Experiment 2 to evaluate the effects of non-linear changes in the image that do not greatly affect facial recognition. For this purpose, we either transposed the shape (fiducial feature positions) of one face onto another (as shown in row B of Fig. 2) or we created photo-realistic caricatures by exaggerating the shape of one face relative to the average shape of the set (as shown in row C of Fig. 2).

We reasoned that if OFA and FFA are the neural substrates for face recognition, when blocks of face images that participants see as having essentially the same identity were presented at different levels of shape change or caricature there should be a lower neural response (adaptation) compared to blocks with different identity faces. This was not what we observed. Instead, we found a release of adaptation in the OFA and FFA under these conditions, with responses as large as those produced when faces changed identity. Therefore, the OFA and FFA register the impact of relatively small non-linear changes that participants often fail to see as changing the face's identity.

4.3. The role of the FFA in face recognition

How do these findings fit with the idea that the FFA is involved in processing the identity of faces? Previous neuroimaging studies using fMR-adaptation paradigms have reported mixed results about whether the FFA has an image-invariant representation of identity. Some studies have reported image invariance (Loffler et al., 2005; Rotshtein et al., 2005; Eger et al., 2005; Ewbank and Andrews, 2008; Davies-Thompson et al., 2013), whereas others have reported image dependence (Grill-Spector et al., 1999; Andrews and Ewbank, 2004; Pourtois et al., 2005; Davies-Thompson et al., 2009; Weibert and Andrews, 2015). Our results show invariance in the neural response to substantial linear changes, but no invariance to non-linear changes, even when these changes do not affect recognition. It seems, then, that the FFA does not represent facial identity at the level where full image invariance is achieved. This is consistent with the idea that interactions between core regions such as the FFA and regions in the extended face processing network are important for the explicit representation of

facial identity (Collins and Olson, 2014).

The task in our fMRI experiments involved the detection of a red dot rather than a face recognition task. It would not have been possible to use the same behavioural task in the fMRI experiments because of the block design. In any event, we actually think that it is an advantage to use an incidental task in the fMRI part of the study. Otherwise, the clear differences in performance for different image conditions could be a potential confound in the experiment. The aim of the fMRI experiments was to measure the perceptual encoding of the stimuli in the absence of any task-based effects; hence our preference for an incidental task that simply ensures participants attend to the stimuli.

Our findings have important implications for understanding how invariant recognition of different images of familiar faces might be achieved. The most widely discussed hypothesis maintains that the ability to recognise faces can be attributed to specialized mechanisms that process the spatial configuration of the face (Carey and Diamond, 1977; Diamond and Carey, 1986; Carey, 1992). These theories distinguish between the 'first order' spatial configuration, the normal arrangement of features (eyes above the nose and nose above mouth) that allows us to recognise that an object is a face, and the 'second order' spatial configuration created by the relative positions of features within the face itself. They claim that it is this second order configuration that allows us to individuate different face identities (for a review, see Maurer et al. (2002)). Indeed, a number of studies have shown that we are sensitive to subtle changes in the spacing of facial features (Freire et al., 2000; Le Grand et al., 2001; Leder et al., 2001; Leder and Bruce, 2000), and such observations form a necessary precondition for entertaining the second order configuration hypothesis.

However, a major problem with second order configurational accounts of face recognition is that many changes to the second-order configuration do not affect face recognition (Burton et al., 2015; Taschereau-Dumouchel et al., 2010). In fact, the distances between features are not even stable across different images or views of the same person. For example, opening your mouth changes the distance between your lower lip and the tip of your nose, yet has no measurable effect on recognising your face for anyone who knows you reasonably well. Similarly, a change from a full-face to a three-quarter viewing position will alter the distance between your eyes substantially, but won't stop your friends recognising you. Likewise, as we have shown here, huge transformations in the second-order configuration can be ignored both behaviourally and at the neural level if the underlying transforms are linear.

Instead of coding the second-order configuration, what we find in the FFA is a degree of invariance to linear changes, but considerable sensitivity to non-linear changes in facial images, even when these changes do not affect recognition. One way to interpret these findings is that FFA may contribute to some form of normalisation that underpins subsequent recognition. From this perspective, the FFA may well play an important role in the early stages of face recognition even if it does not in itself show all the properties one might expect of a fully-fledged recognition mechanism (Andrews et al., 2016). Indeed, this is consistent with the idea that interactions between core regions such as the FFA and regions in the extended face processing network, particularly in the anterior temporal lobe, are important for the explicit representation of facial identity (Collins and Olson, 2014). To address this question directly, we defined a region in the anterior temporal lobe that responded selectively to faces (see Rajimehr et al. (2009)). However, this region failed to show any significant effect of condition in either Experiment 1 or 2. We do not think that the absence of any effect in our analyses rules out the possibility that anterior temporal regions play an important role in face recognition (Kriegeskorte et al., 2007, PNAS; Nasr and Tootell, 2012, NeuroImage). It may just reflect the fact that we are unable to discriminate the key anterior temporal regions in our data, due to the magnetic susceptibility artifacts of our fMRI data (see Axelrod and Yovel (2013)).

4.4. The role of the STS in face processing

In addition to proposing that the FFA is critical to coding invariant aspects of faces, neural models of human face perception also propose a pathway leading to the posterior STS that is responsible for processing changeable aspects of faces such as gaze and expression (Haxby et al., 2000). In our experiments, the posterior STS did not show any adaptation to identity, but did show a larger response to small non-linear changes in shape that fell within the normal range of variability. One interpretation of these findings is that the changes in shape within a block are being interpreted as dynamic transformations of an individual face (see Lee et al. (2010), Pitcher et al. (2011)). This would be consistent with previous studies that have shown that the posterior STS responds selectively to sequences of same identity faces that vary in gaze and expression (Andrews and Ewbank, 2004; Davies-Thompson et al., 2009; Baseler et al., 2014).

In conclusion, we have shown that core face-selective regions, such as the FFA, tolerate naturally-occurring linear changes in face images. In contrast, the FFA was sensitive to non-linear changes in the face, even when these changes do not affect recognition. These findings show that the neural representation in the FFA does not have a fully invariant representation of facial identity. While the FFA may therefore be *involved* in the recognition of face identity, it is more likely that it contributes to some form of normalisation that underpins subsequent recognition than that it forms the neural substrate of recognition per se.

Acknowledgments

This work was supported by a grant from the Wellcome Trust (WT087720MA).

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.neuropsychologia.2016.10.004.

References

- Allison, T., Ginter, H., McCarthy, G., Nobre, A.C., Puce, A., Luby, M., Spencer, D.D., 1994. Face recognition in human extrastriate cortex. *J. Neurophysiol.* 71, 821–825.
- Andrews, T.J., Ewbank, M.P., 2004. Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *NeuroImage* 23, 905–913.
- Axelrod, V., Yovel, G., 2013. The challenge of localizing the anterior temporal face area: A possible solution. *NeuroImage* 81, 371–380.
- Baseler, H.A., Harris, R.J., Young, A.W., Andrews, T.J., 2014. Neural responses to expression and gaze in the posterior superior temporal sulcus interact with facial identity. *Cereb. Cortex* 24, 737–744.
- Benson, P.J., Perrett, D.I., 1991. Perception and recognition of photographic quality facial caricatures: implications for the recognition of natural images. *Eur. J. Cogn. Psychol.* 3.
- Benson, P.J., Perrett, D.I., 1994. Visual processing of facial distinctiveness. *Perception* 23, 75–93.
- Bindemann, M., Burton, A.M., Leuthold, H., Schweinberger, S.R., 2008. Brain potential correlates of face recognition: geometric distortions and the N250r brain response to stimulus repetitions. *Psychophysiology* 45, 535–544.
- Burton, A.M., Wilson, S., Cowan, M., Bruce, V., 1999. Face recognition in poor-quality video: evidence from security surveillance. *Psychol. Sci.* 10 (3), 243–248.
- Burton, A.M., Jenkins, R., Hancock, P.J.B., White, D., 2005. Robust representations for face recognition: the power of averages. *Cogn. Psychol.* 51, 256–284.
- Burton, A.M., Schweinberger, S.R., Jenkins, R., Kaufmann, J.M., 2015. Arguments against a configural processing account of familiar face recognition. *Perspect. Psychol. Sci.* 10, 482–496.
- Calder, A.J., Young, A.W., Benson, P.J., Perrett, D.I., 1996. Self priming from distinctive and caricatured faces. *Br. J. Psychol.* 87, 141–162.
- Carey, S., 1992. Becoming a face expert. *Philos. Trans. R. Soc. Lond. B335*, 95–103.
- Carey, S., Diamond, R., 1977. From piecemeal to configurational representation of faces. *Science* 195, 312–314.
- Collins, J.A., Olson, I.R., 2014. Beyond the FFA: the role of the ventral anterior temporal lobes in face processing. *Neuropsychologia* 61, 65–79.
- Davies-Thompson, J., Gouws, A., Andrews, T., 2009. An image-dependent representation of familiar and unfamiliar faces in the human ventral stream. *Neuropsychologia* 47, 1627–1635.
- Diamond, R., Carey, S., 1986. Why faces are and are not special: an effect of expertise. *J. Exp. Psychol. Gen.* 115, 107–117.
- Eger, E., Schweinberger, S.R., Dolan, R.J., Henson, R.N., 2005. Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *NeuroImage* 15, 1128–1139.
- Ewbank, M.P., Andrews, T.J., 2008b. Differential sensitivity for viewpoint between familiar and unfamiliar faces in human visual cortex. *NeuroImage* 40, 1857–1870.
- Freire, A., Lee, K., Symons, L.A., 2000. The face-inversion effect as a deficit in the encoding of configural information: direct evidence. *Perception* 29, 159–170.
- Frowd, C.D., Jones, S., Fodarella, C., Skelton, F., Fields, S., Williams, A., Marsh, J.E., Thorley, R., Nelson, L., Greenwood, L., Date, L., Kearley, K., McIntyre, A.H., Hancock, P.J.B., 2014. Configural and featural information in facial-composite images. *Sci. Justice* 54, 215–227.
- Grill-Spector, K., Knouf, N., Kanwisher, N., 2004. The fusiform face area subserves face perception, not generic within-category identification. *Nat. Neurosci.* 7, 555–562.
- Grill-Spector, K., Kushnir, T., Edelman, Avidan, G., Itzhak, Y., Malach, R., 1999. Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* 24, 187–203.
- Harper, B., Latto, R., 2001. Cyclopean vision, size estimation, and presence in orthostereoscopic images. *Presence* 10 (3), 312–330.
- Haxby, J.V., Hoffman, E.A., Gobbini, M.I., 2000. The distributed human neural system for face perception. *Trends Cogn. Sci.* 4, 223–233.
- Hill, H., Bruce, V., 1996. The effects of lighting on the perception of facial surfaces. *J. Exp. Psychol.: Hum. Percept. Perform.* 22 (4), 986.
- Hole, G.J., George, P.A., Eaves, K., Rasek, A., 2002. Effects of geometric distortions on face-recognition performance. *Perception* 31, 1221–1240.
- Jenkins, R., Burton, A.M., 2011. Stable face representations. *Philos. Trans. R. Soc.* 366, 1671–1683.
- Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* 17, 4302–4311.
- Kriegeskorte, N., Formisano, E., Sorger, B., Goebel, R., 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. *PNAS* 104(51): 20600–5. *Proc. Natl. Acad. Sci.* 104 (51), 20600–20605.
- Le Grand, R., Mondloch, C.J., Maurer, D., Brent, H.P., 2001. Neuroperspective. Early visual experience and face processing. *Nature* 410, 890.
- Leder, H., Bruce, V., 2000. When inverted faces are recognized: the role of configural information in face recognition. *Q. J. Exp. Psychol. A* 53, 513–536.
- Leder, H., Candrian, G., Huber, O., Bruce, V., 2001. Configural features in the context of upright and inverted faces. *Perception* 30, 73–83.
- Loffler, G., Yourganov, G., Wilkinson, F., Wilson, H.R., 2005. fMRI evidence for the neural representation of faces. *Nat. Neurosci.* 8, 1386–1390.
- Maurer, D., Le Grand, R., Mondloch, C.J., 2002. The many faces of configural processing. *Trends Cogn. Sci.* 6, 255–260.
- Nasr, S., Tootell, R.B., 2012. Role of fusiform and anterior temporal cortical areas in facial recognition. *NeuroImage* 63, 1743–1753.
- O'Toole, A.J., Edelman, S., Bulthoff, H.H., 1998. Stimulus-specific effects in face recognition over changes in viewpoint. *Vis. Res.* 38 (15), 2351–2363.
- Patterson, K.E., Baddeley, A.D., 1977. When face recognition fails. *J. Exp. Psychol.: Hum. Learn. Mem.* 3 (4), 406–417.
- Pickering, E.C., Schweinberger, S.R., 2003. N200, N250r, and N400 event-related brain potentials reveal three loci of repetition priming for familiar names. *J. Exp. Psychol. Learn. Mem. Cogn.* 29 (6), 1298–1311.
- Pitcher, D., Dilks, D.D., Saxe, R.R., Triantafyllou, C., Kanwisher, N., 2011. Differential selectivity for dynamic versus static information in face-selective cortical regions. *NeuroImage* 56, 2356–2363.
- Pourtois, G., Schwartz, S., Seghier, M., Lazeyras, F., Vuilleumier, P., 2005. View-independent coding of face identity in frontal and temporal cortices is modulated by familiarity: an event-related fMRI study. *NeuroImage* 24, 1214–1224.
- Rajimehr, R., Young, J.C., Tootell, R.B., 2009. An anterior temporal face patch in human cortex, predicted by macaque maps. *Proc. Natl. Acad. Sci. USA* 106, 1995–2000.
- Rhodes, G., Brennan, S.E., Carey, S., 1987. Identification and ratings of caricatures: implications for mental representations of faces. *Cogn. Psychol.* 19, 473–497.
- Rhodes, G., Byatt, G., Tremewan, T., Kennedy, A., 1996. Facial distinctiveness and the power of caricatures. *Perception* 25, 207–223.
- Rotshtein, P., Henson, R.N., Treves, A., Driver, J., Dolan, R.J., 2005. Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nat. Neurosci.* 8, 107–113.
- Sandford, A., Burton, A.M., 2014. Tolerance for distorted faces: challenges to a configural processing account of familiar face recognition. *Cognition* 132, 262–268.
- Schweinberger, S.R., Pickering, E.C., Jentsch, I., Burton, A.M., Kaufmann, J.M., 2002. Event-related brain potential evidence for a response of inferior temporal cortex to familiar face repetitions. *Cogn. Brain Res.* 14 (3), 398–409.
- Sergent, J., Signoret, J.L., 1992. Functional and anatomical decomposition of face processing: evidence from prosopagnosia and PET study of normal subjects. *Philos. Trans. R. Soc. Lond.* 335, 55–62.
- Taschereau-Dumouchel, V., Rossion, B., Schyns, P.G., Gosselin, F., 2010. Interattribute distances do not represent the identity of real world faces. *Front. Psychol.* 1, 159.
- Taubert, J., Parr, L., 2011. Geometric distortions affect face recognition in chimpanzees (Pan troglodytes) and monkeys (Macaca mulatta). *Anim. Cogn.* 14, 35–43. (Monkey/ape study based on Hole human study).
- Weibert, K., Andrews, T., 2015. Activity in the right fusiform face area predicts the behavioural advantage for the perception of familiar faces. *Neuropsychologia* 75, 588–596.
- Yamashita, W., Kanazawa, S., Yamaguchi, M.K., 2013. Tolerance of geometric distortions in infant's face recognition. *Infant Behav. Dev.* 37, 16–20.