Vision Research 98 (2014) 1–13

Contents lists available at ScienceDirect

Vision Research

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

Face identity aftereffects increase monotonically with adaptor extremity over, but not beyond, the range of natural faces

Elinor McKone [a,*], Linda Jeffery [b], Alexandra Boeing [b], Colin W.G. Clifford [c], Gillian Rhodes [b]

[a] Research School of Psychology, Australian National University & ARC Centre of Excellence in Cognition and its Disorders, Australia
[b] ARC Centre of Excellence in Cognition and its Disorders, School of Psychology, University of Western Australia, Australia
[c] School of Psychology, The University of Sydney, and Australian Centre of Excellence in Vision Science, Australia

ARTICLE INFO

Article history:
Received 13 June 2013
Received in revised form 16 January 2014
Available online 28 February 2014

Keywords:
Face perception
Face identification
Face adaptation
Face aftereffects

ABSTRACT

Face identity aftereffects have been used to test theories of the neural coding underlying expert face recognition. Previous studies reported larger aftereffects for adaptors that are morphed further from the average face than for adaptors closer to the average, which appeared to support opponent coding along face-identity dimensions. However, only two levels were tested and it is not clear where they were located relative to the range of naturally occurring faces. This range is of interest given the functional need of the visual system both to produce good discrimination of real everyday faces and to process novel kinds of faces that we may encounter. Here, Experiment 1 establishes the boundary of faces judged as being able to occur in everyday life. Experiment 2 then shows that aftereffects increase with adaptor extremity up to this natural-range boundary, drop significantly immediately outside the boundary, and then remain stable with no drop towards zero even for highly distorted adaptors far beyond the boundary. Computational modelling shows that this unexpected pattern cannot be explained either by a simple opponent or by a classic multichannel model. However, its qualitative features can be captured either by a combination of opponent and multichannel coding (raising the possibility that not all identity-related face dimensions are opponent coded), or by a 3-pool model containing two S-shaped-response channels and a central bell-shaped channel around the average face (raising the possibility of unexpected similarities with coding of eye and head direction).

1. Introduction

We can discriminate and recognize thousands of faces despite their similarity as visual patterns. This expertise seems to rely on norm-based coding of identity, where identity-related dimensions of facial appearance are coded relative to average values that function as norms (for a review, see Rhodes & Leopold, 2011). Norm-based coding offers an efficient way to focus processing resources on distinctive information, which is what matters for recognition. Moreover, the updating of norms by experience allows face-coding mechanisms to be finely calibrated to our diet of faces (for reviews see Armann et al., 2011; Rhodes & Leopold, 2011; Webster & MacLeod, 2011).

The coding of face identity has been widely studied using face identity aftereffects, in which viewing a face for a few seconds selectively biases us to see the opposite identity in a subsequently presented face (Anderson & Wilson, 2005; Armann et al., 2011; Jeffery et al., 2011; Leopold et al., 2001, 2005; Rhodes & Jeffery, 2006; Rhodes et al., 2007; Rhodes & Leopold, 2011). For example, viewing antiDan, who lies opposite Dan in face space, biases us to identify the average face as Dan (Fig. 1). Identity aftereffects survive changes in retinal position between adapt and test faces (see Rhodes & Leopold, 2011 for a review) and are larger for upright than inverted faces (Rhodes, Evangelista, & Jeffery, 2009), indicating that they reflect, at least partially, adaptation of higher-level face-coding mechanisms.

An important feature of face identity aftereffects is that they are much larger for opposite than non-opposite adapt-test pairs, even when these are matched on perceived dissimilarity (Rhodes & Jeffery, 2006). This selectivity of the perceptual bias to see an identity that lies opposite the average face in face-space indicates a special status for the average face in identity coding. Moreover, it is harder to perceive the component identities in a blend of two opposite identities (face and its antiface), which results in an average face, than in a blend of two non-opposite identities, which produces a non-average face (Rhodes & Jeffery, 2006). This finding

* Corresponding author. Address: Research School of Psychology, Australian National University, ACT 0200, Australia. Fax: +61 2 6125 0499.
E-mail address: elinor.mckone@anu.edu.au (E. McKone).

http://dx.doi.org/10.1016/j.visres.2014.01.007
0042-6989/

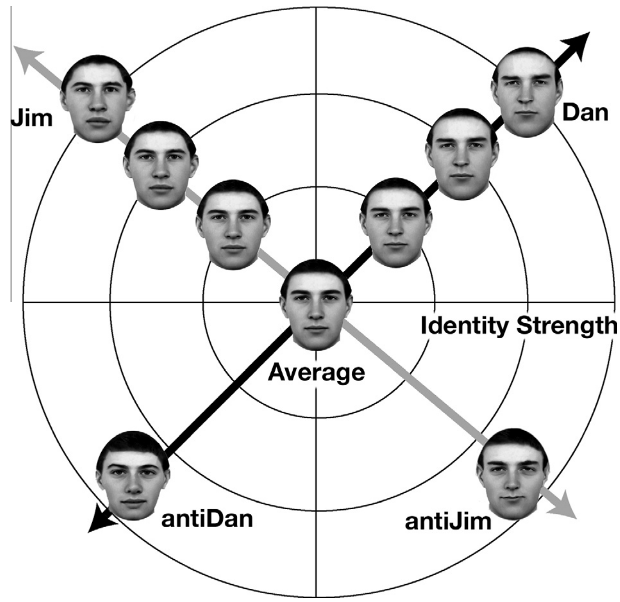

Fig. 1. A simplified face space with two faces, Dan and Jim, an Average face, created by morphing 20 male, Caucasian faces, and two antifaces, antiDan and antiJim. An antiface is made by morphing a face towards, and beyond, the Average, and has opposite properties to that face. Reduced identity strength versions (anticaricatures) of Dan and Jim, created by morphing those identities towards the Average, are also shown. Identity aftereffects occur when exposure to a face biases subsequent perception towards a face with opposite properties. For example, after viewing antiDan for a few seconds, we are biased (briefly) to see Dan.

also highlights a special status of the average face as a neutral point in face-space, from which deviations signal unique identities.

Neurally, it has been proposed that norm-based coding could be implemented by opponent coding of identity-related face dimensions of the form shown in Fig. 2A or 2B (Rhodes & Jeffery, 2006; Rhodes et al., 2005; Robbins, McKone, & Edwards, 2007; Tsao & Freiwald, 2006). In this case, each face dimension (e.g., eye size) would be coded by a pair of neural populations with monotonically increasing or decreasing response functions, one tuned to high (i.e., above-average) values and the other tuned to low (i.e., below-average) values on that dimension. The average value or norm is signalled by balanced activation in the two channels, and unbalanced activity signals low or high values on the dimension. This type of coding is used for other visual attributes that are coded relative to perceptual norms, such as color and aspect ratio (Regan & Hamstra, 1992; Suzuki, 2005; Webster & MacLeod, 2011). Fig. 2 shows two variants of opponent coding, both of which are consistent with neurophysiological data: Single cell recordings of face-selective cells in monkeys have produced both S-shaped (Fig. 2A) and linear (Fig. 2B) monotonic tuning functions (Freiwald, Tsao, & Livingstone, 2009).

In current literature, opponent coding of identity-related face information has been contrasted with non-norm-based models in which each dimension is coded by multiple channels with bell-shaped response functions tuned to different values along the dimension (Fig. 2C). In multichannel coding the average value on a dimension has no special status, does not function as a perceptual norm and need not have any channel specifically tuned to it. Multichannel coding is used for several basic visual attributes, including spatial frequency and tilt (Blakemore & Sutton, 1969; Clifford, Wenderoth, & Spehar, 2000). However, if face identity coding is norm-based, as argued above, then we would not expect identity-related dimensions to be coded using a non-norm-based, multi-channel system (at least not all of them).

Previous studies have sought to test whether face identity is indeed opponent coded by examining how the identity aftereffect changes in size as the adapting faces become more extreme or distinctive (i.e., further from average) (Fiorentini et al., 2012; Jeffery et al., 2010; Jeffery et al., 2011). The opponent coding models illustrated in Fig. 2 predict that the aftereffects will increase with increasing adaptor extremity over the range in which the response functions are increasing. The increase in aftereffect occurs because more extreme adaptors activate their preferred channel more strongly (and their non-preferred channel more weakly) than less extreme adaptors, producing a stronger reduction in response with adaptation, and thus a larger aftereffect (larger shift in the cross-over point at which the two pools are responding equally strongly). Therefore, using an average face as the test image, the bias to see the identity opposite the adaptor, i.e., the aftereffect, should increase as the extremity of adaptors increases.

In contrast, multichannel coding with bell-shaped tuning curves (Fig. 2C) predicts that aftereffects will initially increase as adaptors move away from the test image, but will then reach a maximum and decrease to zero for more extreme adaptors. The decrease occurs because more extreme adaptors will have less impact on channels that respond to the average test face than will less extreme adaptors. This pattern has been reported for both spatial frequency and tilt, which are multichannel coded in V1 (Blakemore & Sutton, 1969; Clifford, Wenderoth, & Spehar, 2000). The precise location of the maximum, and of the decay to zero adaptation, will depend on the range tiled by the channels, the breadth (i.e., full-width-half-maximum) of the channels, and the breadth of spread of the adaptation (i.e., how similar in tuning other channels have to be to the adapted value for their responses to be reduced).

What does previous research suggest happens to identity-related face aftereffects as adaptors become more extreme? For a simple change in the position of a single face feature (eye- or mouth-height) results match the predictions of opponent coding with linear response functions (Fig 2B): aftereffects increase monotonically with adaptor extremity across multiple adaptor values, even up to very extreme values (e.g., eyes almost touching the hairline) (Robbins, McKone, & Edwards, 2007; Susilo, McKone, & Edwards, 2010). It is only when the eyes move outside of the head, thus violating the first-order face configuration, that the aftereffects drop to zero (McKone & Edwards, 2011).

Face identity aftereffects, in which multiple attributes of the face vary simultaneously, also show larger aftereffects for "far" than "near" adaptors (40% vs 80% identity strength) (Fiorentini et al., 2012; Jeffery et al., 2010; Jeffery et al., 2011). These identity aftereffect studies, however, have two limitations. First, only two adaptor levels were used. Both opponent and multichannel models predict an initial increase in aftereffects, and thus the earlier results for identity aftereffects (Fiorentini et al., 2012; Jeffery et al., 2010) are potentially consistent with either model: the pattern of increase from near to far adaptors is as predicted by opponent coding, but alternatively it might be the case that an insufficiently large range of adaptor values was tested to see a later turnaround indicative of multichannel coding.

Second, it is of interest to know where adaptors fall with respect to the range of faces that occur in the natural world. Within observers' perceptual face-space, faces differ in their distinctiveness, with more typical real-world individuals lying closer to the average face and real-world individuals with more unusual facial appearance lying further from the average (Johnston, Milnes, Williams, & Hosie, 1997; Valentine, 1991). Whatever system of neural tuning along face dimensions is used, it needs to be able to provide good discrimination of subtle differences in facial appearance across this full natural range; that is, it needs to provide good coverage of the full 'diet' of faces that we see in everyday life. In addition, there is a question about the extent to which detailed neural coding might continue along face-space dimensions outside the everyday range of faces; that is, for faces that