Cognitive Representation of Facial Asymmetry

David White
Department of Psychology
University of Glasgow

Submitted for the Degree of Ph.D. to the Higher Degree Committee of the Faculty of Information and Mathematical Sciences, University of Glasgow.

March 2008

In memory of my grandmothers

Betty and Jessie.

Abstract

The human face displays mild asymmetry, with measurements of facial structure differing from left to right of the meridian by an average of three percent. Presently this source of variation is of theoretical interest primarily to researchers studying the perception of beauty, but a very limited amount of research has addressed the question of how this variation contributes to the cognitive processes underlying face recognition. This is surprising given that measurement of facial asymmetry can reliably distinguish between even the most similar of faces. Furthermore, brain regions responsible for symmetry detection support face-processing regions, and detection of symmetry is superior in upright faces relative to inverted and contrast-reversed face stimuli. In addition, facial asymmetry provides a useful biometric for automatic face recognition systems, and understanding the contribution of facial asymmetry in human face recognition may therefore inform the development of these systems. In this thesis the extent to which facial asymmetry is implicated in the process of recognition in human participants is quantified. By measuring the effect of left-right reversal on various tasks of face processing, the degree to which facial asymmetry is represented by memory is investigated. Marginal sensitivity to mirror reversal is demonstrated in a number of instances, and it is therefore concluded that cognitive representations of faces specify structural asymmetry. Reversal effects are typically slight however and on a number of occasions no reliable effect of this stimulus manipulation is detected. It is likely that a general tendency to treat mirror reversals as equivalent stimuli, in addition to an inability to recall lateral orientation of objects from memory, somewhat obscure the effect of reversal. The findings are discussed in the context of existing literature examining the way in which faces are cognitively represented.

Acknowledgements

I would like to thank Professor Mike Burton for the support, encouragement and guidance he offered throughout the course of this investigation. I would also like to thank Rosemary Wright for her assistance in collecting the data reported in Experiment seven. This research was completed on a part-time whilst I was employed as a research assistant and I thank Mike Burton, The Faculty of Information and Mathematical Sciences and the Department of Psychology for negotiating this arrangement.

Declaration

I declare that this thesis is my own work carried out under the normal terms of supervision.

....

David White

Contents

Chapter 1

General Introduction

1.1 Introduction

In this thesis I investigate the degree to which the subtle morphological asymmetries present in the human face are represented in memory. The form of a human face, in common with the human form and biological mechanisms in general, displays a high degree of bilateral symmetry. Biologically determined factors contribute to the degree and quality of the asymmetry present in a given organism (e.g. Mather, 1953), and this 'breaking of symmetry' is thought to communicate important social information (e.g. Simmons et al. 2004). Furthermore, symmetry is detected more readily in faces than in other objects (Rhodes et al. 2005), and it would appear that the face processing system and perceptual mechanisms responsible for symmetry detection are mutually supportive of one another (Chen et al. 2007). Given the importance of asymmetry to face processing, it was hypothesised that memory for familiar faces would represent structural asymmetry.

In this chapter I will introduce the topic of my research by firstly offering a broad description of the factors that determine asymmetry in biological form: a topic that incorporates both the physical structure of the human face and the human brain. Additionally, asymmetry in the organisation of cognitive function and behavioural asymmetry will be briefly summarised. Following that, I will introduce a number of phenomena associated with the perception of asymmetry before discussing why asymmetry in faces is important to visual systems. I will then summarise literature emphasising the importance of configural information in face recognition and explain why research showing a significant degree of tolerance to geometric transformations of images complicates this issue. Finally, I will describe previous research investigating perceptual sensitivity to the mirror reversal of familiar faces, outline the overall structure of this thesis, and summarise the methodological approach.

1.2 Asymmetry of form and function

Although symmetry is a salient visual property of the natural and man-made worlds, organisms do not conform with any great rigour to the law of symmetry. Asymmetry in biology is observed from molecular (see Johnson, 2005) to multicellular levels (see

Wolpert, 2005) and deviations from perfect bilateral symmetry are common to all organisms. The degree to which an individual organism displays structural asymmetry is determined by a combination of two factors. Firstly, there exist consistent biases within species for greater development on one side of the body than the other such as the positioning of the heart to the left of the axis of symmetry (*Directional Asymmetry*). In addition, environmental factors induce departures from symmetry in individual organisms. The latter source of asymmetry is known as *Fluctuating Asymmetry* and its degree is thought to vary as a function of the individual's genetic health (e.g. Mather, 1953).

Directionally asymmetric traits exhibit normally distributed right-left differences but have a mean that is either significantly greater or less then zero. The human face is considered to show mild directional asymmetry, with the right side being generally larger than the right (e.g. Farkas & Cheung, 1981), however some studies have failed to replicate these directional patterns (e.g. Hardie et al. 2005). The degree of directional asymmetry present in a typical face is small in relation to the degree of fluctuating asymmetry (i.e. normally distributed right-left differences with a mean of zero), which is implicated in the perception of beauty (e.g. Rhodes et al. 1998), and our perceptual systems are highly sensitive to this latter source of variance (Simmons et al. 2004).

The appearance of the human brain is characterised by the longitudinal fissure that divides the organ into two halves, and the two hemispheres appear, on a superficial level at least, to be mirror images of one another. In common with faces however, subtle anatomical asymmetries have been reported, with the plannum temporale (found under the sylvian fissure) commonly lager on the right side of right-handers and the left side of left-handers. This asymmetry is thought to be associated with left-hemisphere dominance in language (Geschwind & Levitsky, 1968; Petty, 1999; Shapleske et al. 1999). In addition, the neurochemistry of the brain is asymmetric (e.g. Glick et al. 1982), and disruption of normal brain asymmetries can underlie psychopathologies. Abnormal cerebral asymmetry in both neurochemistry (e.g. Reynolds, 1983; Shirakawa et al. 2001) and anatomy (e.g. Petty, 1999; Falkai et al. 2002) has been shown in schizophrenic patients. Abnormal lateralisation has also been observed in children with learning disabilities (Cornish & McManus, 1996) and cognitive disorders (Hugdahl, 1998).

The development of directional asymmetry during embryonic development is non-random and is dependent upon firstly specifying the midline of the organism before specifying which side is the left and which side is the right. This process of specifying the left-right axis of the body leads to asymmetry in morphology and positioning of visceral organs such as the heart, which is invariantly positioned to the left of the body in vertebrates and this process is driven by molecular mechanisms (Ramsdell & Yost, 1998). Normal directional asymmetry in the visceral organs can occasionally be disrupted however. In a bizarre and very rare condition known as *situs inversus*, disorder of these molecular mechanisms (the nodal cilia) can cause the entire physiology of the body to be reversed (e.g. McManus, 2005). This disorder also causes anatomical asymmetries in the brain to be reversed, but crucially these same patients do not show reversed handedness or language dominance (Kennedy et al. 1999). In a more recent study it was shown that 7 out of 46 individuals with s*itus inversus* were found to be left-handed, which is approximately equivalent to normal prevalence of left-handedness, making it highly likely that different mechanisms are responsible for determining anatomical and functional asymmetries (McManus et al. 2004; see also McManus, 2005).

Though systems responsible for the development of functional and anatomical asymmetries may be dissociable, they nevertheless produce comparably asymmetric organisation. Functional organisation of the human brain is asymmetric (see Hugdahl & Davidson, 2002) and functional asymmetries have been observed in a variety of animal species (Corballis, Funnell & Gazzaniga, 2000). Although classical research has shown that the two cerebral hemispheres continue to function independently after nerves connecting them are severed, the two hemispheres continue to perform separate functions (Sperry, 1961). In fact, since Paul Broca's initial observation that a lesion in the left hemisphere caused impaired language function (Broca, 1861), there have been shown to be many a great differences in cognitive function between the two hemispheres (Gazzaniga, 2005). Additionally, functional asymmetry is apparent from an early age and certain asymmetries would appear to be innately determined (e.g. Trevarthen, 1996).

Humans also display asymmetry in their outward behaviour. Roughly 90% of humans prefer to use their right hand for manual tasks, a pattern that is consistent across culture, and this ratio is far greater than has been reported in any other mammal (Corballis, 2003). In addition, behavioural asymmetries are displayed very early in development, with embryos showing a tendency to turn their head to the right (Ververs et al. 1994a, 1994b) and newborns typically displaying greater motor development on their right (Grattan et al. 1992; Tan et al. 1992; Rönnqvist, 1995). The former tendency, established in embryonic development, is likely to contribute to a bias later in life to turn the head towards the right (Güntürkün, 2003). This tendency is further evidenced in portraiture, where subjects are more often shown with their left-cheek facing the viewer (McManus &Humphrey, 1973; Grüsser et al. 1988; Nicholls et al. 1999; ten Cate 2002). The causes of both anatomical and functional asymmetries are themselves difficult to determine (Hobert et al. 2002), however as functional and behavioural asymmetries are observed remarkably early in development, it is likely that brain asymmetry and lateralised behaviours are genetically determined (Sun & Walsh, 2006). As such, the distinction between the asymmetry of form and function, in the human brain at least, is not straightforward.

Mechanisms responsible for processing human faces are functionally asymmetric, with a greater reliance on the right hemisphere than the left (e.g. De Renzi et al. 1994). A number of small functional areas have been identified as responding more to faces than to other objects suggesting that these areas have developed specifically for the task of perceiving human faces. Two regions in particular have been repeatedly shown to respond selectively to faces (i.e. the 'Fusiform Face Area' and the 'Occipital Face Area') and activity in these two areas is often stronger in the right hemisphere (Halgren et al. 1999; Haxby et al. 1999; Rossion et al. 2000 & 2003). In addition, prosopagnosia is most commonly caused by damage to the right hemisphere (see Sorger et al. 2007). This right hemisphere bias in face processing is accompanied by a number of leftward perceptual biases. When viewing faces people more commonly attend to the side appearing to the left (Yarbus, 1967; Butler et al. 2005), and the left side of faces contributes more to perception of identity (Wolff, 1933) and emotional expression (Heller & Levy, 1981; David, 1993; Ferber & Murray, 2005).

1.3 The perception and representation of asymmetry and lateral orientation

Although examples of perfect bilateral symmetry rarely occur in nature, humans display *a priori* knowledge of what it means for something to be symmetrical. In fact, even organisms with relatively small nervous systems such as bees show an ability to discriminate between symmetrical and asymmetrical stimuli and abstract this understanding to novel patterns (Giurfa et al. 1996). Sensitivity to symmetry has also been found in insects and birds (Lehrer et al. 1994; Swaddle & Cuthill 1994; Møller 1995; Møller & Sorci, 1998) and infants as young as four months show enhanced processing of vertically symmetrical stimuli relative to horizontally symmetrical or asymmetrical stimuli (Bornstein et al. 1981). The latter advantage has also been demonstrated when asking adults to reproduce visual patterns from memory (Deregowski, 1971), and when detecting symmetry in random dot displays humans show a preference for a vertically oriented axis of symmetry (Barlow & Reeves, 1979).

The importance of bilateral symmetry in visual perception was recognised by Gestalt psychologists, who posited symmetry as one of key principles in the law of *prägnanz* (Wertheimer, 1923; Koffka, 1935). The rapid detection of symmetry is likely to aid early visual processing in tasks such as figure-ground seperation (Rock, 1983) and efficient processing of bilateral symmetry may also contribute to object constancy across changes in viewpoint (e.g. Vetter et al. 1994; Troje & Bülthoff, 1998). It has also been argued that symmetry contributes to object recognition by helping to establish an object-centred axis against which coordinates can be represented (Marr & Nishihara, 1978). Perhaps most importantly however, symmetry allows objects to be represented economically (Barlow & Reeves, 1979). For bilaterally symmetric objects, only one half of the object needs to be represented as it is physically identical to the other half, and this could provide the visual system with a basis on which to store representations in a reduced form. Furthermore, when viewing symmetrical stimuli participants eye movements tend to cluster on one half on the stimulus, suggesting that perception is optimally efficient when perceiving symmetrical objects (Locher & Nodine, 1973).

Although biological visual systems are very sensitive to both symmetry and deviations from symmetry (e.g. Wagemans, 1997) they also, somewhat paradoxically, have particular difficulty discriminating between figures that are mirror images of one other (Bornstein et al. 1978; Davidson, 1935; Rudel &Teuber, 1963; Sutherland, 1961). This difficulty is further evidenced in the course of childhood development with children under seven years of age often confusing letters that are mirror images of one another such as the letters 'b' and 'd' in reading (Mach, 1914; Davidson, 1935) and writing (Cornell, 1985). In addition, individuals commonly report a difficulty in telling their left from right in their everyday lives and reproduce this confusion when tested under laboratory conditions (Snyder, 1991), a phenomenon which has been attributed to the (near) bilateral symmetry of the nervous system (Corballis & Beale, 1976). More recent research has shown that the difficulty in discriminating mirror images is probably caused by neurones responding similarly to a given stimulus and its mirror reflection (Rollenhagen & Olson, 2000).

Reports of difficulties discriminating mirror images in the normal population are paralleled by comparable (but significantly more profound) deficits in brain damaged patients. Davidoff & Warrington (1999, 2001; Warrington & Davidoff, 2000) describe two individuals that, after suffering posterior brain damage, display a complete inability to decide if two simultaneously presented images are in the same or different mirror orientations. This phenomenon had previously been described in two separate case reports (Turnbull et al. 1995; Turnbull & McCarthy, 1996), and in all of these cases selective impairment of mirror discrimination occurred in the context of normal object identification. In fact, Warrington & Davidoff (2000) report one patient who displays reciprocal inhibition of these two abilities (i.e. her ability to discriminate mirror images was dependent upon her *inability* to recognise the presented stimuli). This disorder has been reported on a number of occasions subsequently (Karnath et al. 2000; Cooper & Humphreys, 2000; Harris et al. 2001) and the disorder is thought to result from a deficit in determining object-centred orientation (Priftis et al. 2003).

A similar deficit has been reported by Ramachandran and colleagues (1997) who describe a condition displayed by four right hemisphere stroke patients where their ability to create representation of mirror reflection appears compromised. This condition, which they call 'Mirror Agnosia', is characterised by an inability to alter grasping behaviour when a mirror is placed in the non-neglected hemifield so that the patient can see an object located in the neglected hemifield. When asked to grasp this object patients are observed reaching towards the mirror itself and (somewhat bizarrely) searching behind the mirror, even though they report being fully aware that what they are looking at is a mirror (i.e. based on this knowledge the patient should be able to intellectually deduce the objects' location). One explanation proposed by the authors is that patients cannot form representations of mirror reflection, and although this deficit may be linked to the phenomenon of perceiving mirror reversals as identical images (e.g. Davidoff & Warrington, 2001), there is no literature reporting mirror discrimination deficits in mirror agnostics.

Further case reports describe patients with similar disturbances in the left-right orientation of their visual world. The phenomenon of 'mirror-writing' is observed in normal development (e.g. Cornell, 1985) but may also persist into adulthood and can occur spontaneously as a result of neurological disorder (Critchley, 1926). Furthermore, the tendency to write in mirror reversed script may occur as a result of abnormal development and could be genetically determined, with the prevalence of this disorder being estimated to be 1 in 6500 (Mathewson, 2004). Very rarely, mirror writing occurs in the presence of 'mirror reading' (i.e. written words are processed most efficiently when presented in mirror orientation), and this condition can cause the patients entire visual experience to be reversed in left-right orientation (Jokel & Conn, 1999; Lambon-Ralph et al. 1997; Pflugshaupt et al. 2007).

As mentioned previously, normally functioning humans have difficulty in discriminating between mirror images when the objects are reversed left to right, relative to the same task with top to bottom reversed stimuli (e.g. Bornstein et al. 1978). In addition to difficulties perceiving the left-right orientation of visual stimuli, humans also display a profound difficulty in *remembering* the orientation of objects. In a classical demonstration by Nickerson and Adams (1979) it was shown that when asked to reproduce the orientation that a monarchs' head appears on a coin participants performed at chance level, and subsequent studies have found performance on this task to be equivalently poor (Jones, 1990; Martin & Jones, 1997; Kosslyn & Rabin, 1999; Rubin and Kontis, 1983).

1.4 Why is facial asymmetry important for visual systems?

Faces are asymmetric and generally larger on the right than they are on the left (Farkas & Cheung, 1981; Peck et al. 1991; Sackheim, 1985). This directional asymmetry is not always observed however with some studies showing no difference between average measures of the two hemifaces (Hardie et al. 2005), and others showing that directional asymmetry varies as a function of sex (Ferrario et al. 1993; Smith, 2000), handedness (Hardie et al. 2005) and occupation (Smith, 1998). Although evidence for directional asymmetry is equivocal, asymmetry is a property of all faces, and even faces perceived as being highly symmetrical display significant levels of asymmetry (Peck et al. 1991). The structural asymmetry of faces is actually quite large, with distances from facial landmarks to centre points ranging from 4% to 12% average difference, depending on the landmark measured (Ferrario et al. 2001). In addition, patterns of asymmetry are unique to an individual, to the extent that this information can discriminate between faces of identical twins (Burke & Healy, 1993).

Psychological research on facial asymmetry has focused on the degree to which it affects the perceived attractiveness of a face. Driven by the established link between fluctuating asymmetry (FA) and developmental instability (e.g. Mather, 1953), the covariance of facial symmetry and perceived attractiveness has been repeatedly demonstrated (Jones & Hill, 1993; Grammer & Thornhill, 1994; Langlois et al. 1994; Rhodes et al. 1998, 1999; Scheib et al. 1999; Zebrowitz et al. 1996). Furthermore, it is considered that some of the variance in developmental stability of the face is caused by prenatal exposure to sex hormones (Fink et al. 2004).

Our ability to detect asymmetry in faces is extremely well developed and we are highly sensitive to changes in symmetry (Rhodes, 1999). In addition, it would appear that perceived symmetry reflects actual FA accurately, and that judgements of symmetry are made independently of directional asymmetry, suggesting that developmental instability and perceived attractiveness are closely related (Rhodes, 2004). The importance of symmetry in mate choice is further evidenced by the fact that women show an olfactory preference for more symmetrical men during periods of peak fertility (e.g. Gagenstad & Thornhill, 1998) and detect facial symmetry more accurately during these stages in the menstrual cycle (Oinonen & Mazmanian, 2007).

In addition to morphological asymmetries, faces also display asymmetry when expressing emotion, and this is thought to result from underlying functional asymmetry in brain function (e.g. Sackheim et al. 1978). Furthermore, facial asymmetry during emotional expression is determined in part by the underlying anatomical structure of the face (Schmidt et al. 2006). Paradoxically however, visual information relating to expression is said to be processed in parallel to identity-related information (Bruce & Young, 1986), and therefore this source of asymmetry is unlikely to contribute to the recognition of faces. Facial asymmetry in emotional expression is however known to communicate important social information, with asymmetry being more pronounced during intentional compared to spontaneous displays of emotion (Ekman et al. 1981).

That perceptual systems are so finely tuned to bilateral asymmetry in both face (e.g. Rhodes, 1999) and non-face objects (Wagemans, 1997) implicates the existence of mechanisms specialised for processing and representing spatial relationships in terms of their asymmetry. Indeed, as the ability to detect symmetry is displayed by organisms with relatively small nervous systems (e.g. Møller & Sorci, 1998) it is probable that this ability is driven by low-level visual mechanisms. However, it has been shown that symmetry is more accurately detected when faces are presented in upright versus inverted orientation and inverting the polarity of face images also makes symmetry detection more difficult (Rhodes et al. 2005). This finding suggests that symmetry detection in faces is not entirely reliant on low level mechanisms (as low-level properties of face images are not altered by inversion), implicating higher-level visual mechanisms in facial symmetry detection.

In the same study however, Rhodes and colleagues (2005) found that tilting face stimuli by 45 degrees disrupted symmetry detection. They argue that mechanisms responsible for face recognition can therefore be dissociated from those processing facial symmetry, as this same manipulation does not disrupt normal face processing (e.g. McKone et al. 2001). Little and Jones (2006) report contrary evidence however, appearing to show that symmetry detection is unaffected by inversion whereas preference for symmetrical faces is significantly reduced by inversion. The authors argue that mechanisms responsible for symmetry detection and for symmetry preference may too be dissociable. However, it is likely that participants' familiarity with the test images prior to symmetry detection may have affected their results and that their design was not sensitive enough to detect the effect of inversion on this task.

In a recent study by Chen and colleagues (2007) the close relationship between symmetry processing and face processing has been further established. Firstly, the authors found that regions of the brain responding to symmetry are located very close to the right occipital face area (OFA), which is heavily implicated in face processing (e.g. Rossion et al. 2003). In addition, researchers report that a large proportion of the BOLD response to faces in the right OFA can be attributed to the symmetry of faces (although symmetry alone does not evoke the same response). Taken together these results imply that brain areas processing symmetry information lend support to processing of face information in the OFA.

Though it is uncertain whether the processing of facial symmetry can be functionally dissociated from face recognition mechanisms in human cognition, facial symmetry has nonetheless provided a useful biometric for automatic face recognition systems. The use of facial asymmetry in face recognition systems is a relatively new development, with Liu *et al* (2002) providing the first demonstration that facial asymmetry can provide the basis for efficient identification. Early investigation has indicated that this biometric could be of use in situations requiring quick and accurate real-time identification across changes in expression (Liu et al. 2003; Mitra et al. 2007). A key question that will be addressed in this thesis is the extent to which facial asymmetry contributes to human face recognition and the extent to which this asymmetry is represented by cognitive mechanisms. Therefore the data reported here, in addition to resolving existing debate in psychological research, may inform the development of algorithms designed for automatic face recognition.

1.5 How are faces cognitively represented?

The process of face recognition must rely on some sophisticated process of abstraction. Take for instance the variation of surface features present in the images shown below (Figure 1.1). It is difficult when faced with this collection of photographs to articulate or determine just what aspect of the face is constant or what exactly its essence may be. Not only do the textural aspects of lighting, make-up and aging combine to produce vastly different patterns and absolute values of luminance, but the relations between the features of the face also vary between the images. Both the angle of the head relative to the camera and the camera lens itself contribute to variance in the two-dimensional positional relationships between component parts of the face. This degree of variation in the pattern of light produced by a given face poses a significant problem for computer systems designed to recognise faces, and changes in superficial image characteristics can severely disrupt performance on automatic recognition systems (Phillips et al. 2000, 2003).

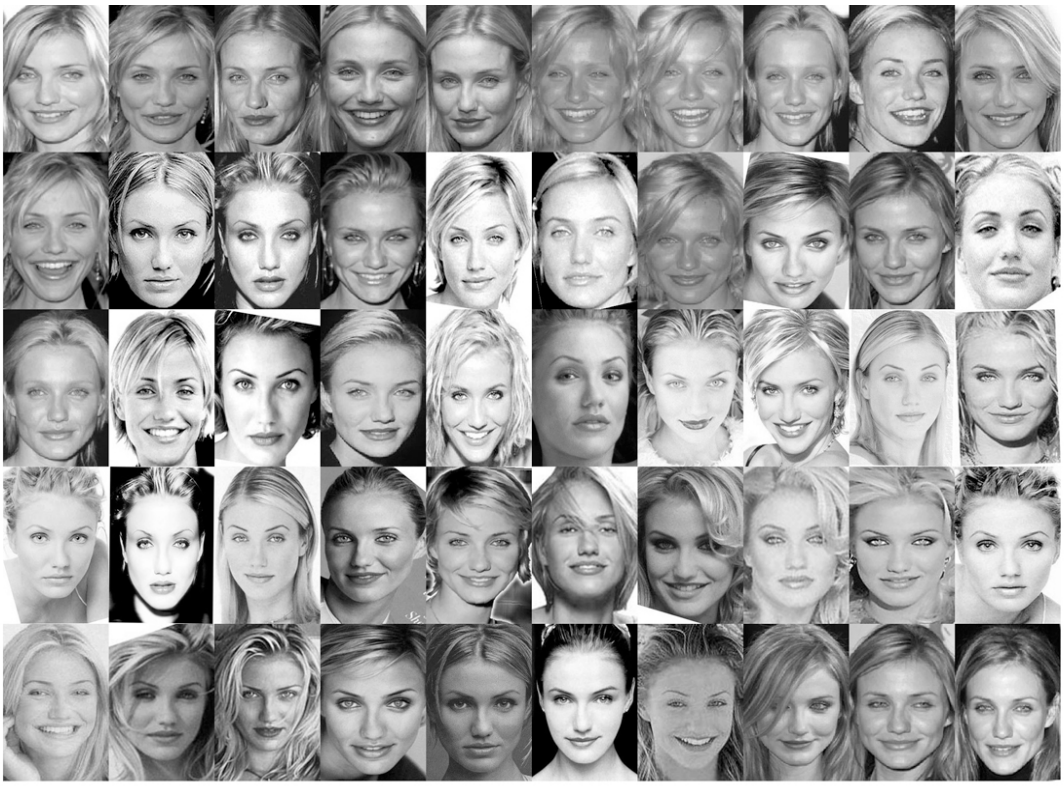

Figure 1.1: 50 photographs of the Hollywood actress Cameron Diaz.

The object constancy displayed by the face recognition system is necessarily robust, as the ability to recognise faces despite the significant variation in visual sensation is essential to successful social interaction. Memory representations of faces exhibit robustness in a greater sense than this however. For instance, we are able to remember a great number of faces and for a very long time, as was shown by Bahrick, Bahrick and Wittlinger (1975) who report 90% correct performance level when showing pictures of classmates to subjects for up to 35 years after graduation. Furthermore, we can recognise a familiar face even when the image is severely degraded (e.g. Sergent, 1986; Burton, Wilson, Cowan & Bruce, 1999), and recognition displays a stubborn resilience to profound image distortions such as stretching and blurring (Hole et al., 2002). When one considers the resilience and reliability of our face recognition abilities in this way, a picture begins to emerge of a cognitive system that is capable of embodying a huge number of abstract representations. The exact number of faces that can be recognised is difficult to determine, but it has been suggested that the average person can *"...probably identify several thousand faces"* (Ellis, 1981, p171). Each of these representations in turn is required to assimilate a large amount of variance (Burton et al, 2005).

Historically, debate regarding the nature of our mental representation of faces has centred on the question of whether the visual information is coded by segregating the separate features, or whether a configural mapping of the face results in a gestalt representation that enables 'holistic' processing (e.g. Tanaka & Farah, 1993). The term 'configural' is used to refer to a process that achieves perception of a given object by virtue of the spatial-relations between its component features. This is in contrast to 'featural' processing which represents an object as an index of individual features that can be matched to the perceived stimulus during the process of identification. The use of this terminology is often ambiguous however, especially with regards to what is meant by 'configural' information (see Bruce and Humphreys, 1994; Rhodes et al. 1993).

In an influential review of literature on this subject, Maurer *et al* (2002) distinguished three types of configural processing that are all implicated in the stages of processing underpinning face recognition. The first of these processes detects the presence of a face in our visual scene and this is accomplished by perception of what is referred to as ‘first order’ configural processing: The general form of the face is detected (i.e. two eyes above a nose, with a mouth directly below the nose) and from here the process of identification can proceed. Once the face has been detected a further process coheres the various features of the face to form a gestalt representation that can be represented as a function of the ‘second-order’ spatial relationships between the component features, and these spatial relationships are then processed in the final process.

The second-order configural information is implicated in the identification of a face, whereby the precise distance between the facial features provides an informative metric that the mind exploits for the task of recognition (see also Diamond & Carey, 1986; Leder & Bruce, 2000). As the task of discriminating between two faces must necessarily involve divining very small differences in the configural information (e.g. Farkas, 1994; Farkas & Cheung, 1981; Farkas & Munro, 1978), it is of little surprise that the perceptual mechanisms implicated in this task are finely tuned to the perception of these relations. Using experimental stimuli generated by the facial reconstruction tools Photofit and Indentikit (which are used by the police service when trying to generate a likeness to a criminal from eyewitness accounts), Haig (1984) demonstrated that adults are sensitive to extremely small changes in the distance between features. In fact, it would appear this sensitivity is constrained only by the limits of our visual acuity.

This fine sensitivity to the spatial arrangement of features is characteristic of face recognition, yet it has been reported that sensitivity is markedly impaired when a face is either turned upside down (e.g. Leder et al. 2001; Collishaw & Hole, 2000) or if its luminance is reversed (e.g. Kemp et al. 1990). These same manipulations have a far less drastic effect on recognition of objects and other non-face objects of expertise (Robbins & McKone, 2007), which has led researchers to argue that the processing of subtle second-order relations is specifically adapted for face recognition (McKone et al. 2007). Furthermore, the inversion effect (turning a face upside down) has little effect on tasks where subjects are required to make judgements based on the component features of the face (Friere et al 2000; Leder & Bruce, 2000; Leder et al. 2001), suggesting that processes involved in extracting featural and configural information are dissociable.

The distinction between featural and configural processing is far from straightforward however, and research on this topic has largely failed to demonstrate that such processes are in fact separable. Riesenhuber *et al* (2004) demonstrated that inversion impairs performance on a face matching procedure when faces differ only in terms of their features. This finding is inconsistent with previous research showing a null effect of inversion when participants are required to make judgements based on facial features (Friere et al 2000; Leder & Bruce, 2000; Le Grand et al, 2001; Leder et al., 2001). However, subsequent studies have also reported similar inversion effects for discrimination of faces altered in configurational and featural aspects (Malcolm et al. 2004; Yovel & Kanwisher, 2004a), and it has been shown that features are more easily matched when presented in the context of a whole face (Tanaka & Sengco, 1997; Tanaka & Farah, 1993). These findings would appear to implicate an interaction between processing of facial features and their global configuration, and fMRI studies attempting to show dissociation between these types of processing have produced ambiguous results (Yovel & Kanwisher, 2004b; Yovel & Duchaine, 2006; Maurer et al. 2007).

That processing of configural and featural information in faces has proved difficult to separate empirically is likely to stem from an inherent circularity in the definition of these apparently independent sources of information. For instance, although the eye region is generally considered to be a 'feature' (e.g. Riesenhuber et al. 2004, Leder et al. 2001), there is an abundance of 'configural' information within this area, such as the distance between the eyebrow and the eyelid. That said, it is difficult to see the value in making a psychological distinction between two types of processes if the two types of information that are apparently separated by these processes are neither perceptually nor physically dissociable. Just how inversion disrupts face-processing remains uncertain and it would appear unlikely that it selectively disrupts 'configural' processing, at least according to the current definition (see Maurer, 2002). It has been demonstrated that perception of vertical spacing is disrupted more by inversion than the perception of horizontal spacing (Goffaux et al. 2007), which suggests that inversion impedes face processing more selectively than had been previously thought.

An alternative hypothesis for why inversion disrupts face processing is that it prevents faces from being processed holistically. Farah *et al* (1995) provide strong evidence that the inversion effect is caused by selective disruption to holistic as opposed to configural processing in an experiment using dot pattern stimuli. Hypothesising that face recognition differs from other object processing because it relies on less 'part decomposition', they found that inversion worsened memory for random configurations of dots but did not affect performance when the dot patterns were sub-divided (i.e. grouped) by colour. This result would appear to suggest that inversion disrupts the processing of complex stimuli that are not segmented into 'parts'. In the same paper, Farah and colleagues also report a similar result obtained with face stimuli. They found that inversion does not affect the ability to recall a face broken into its component parts but it does impair performance on the same task when the face is presented in its normal arrangement.

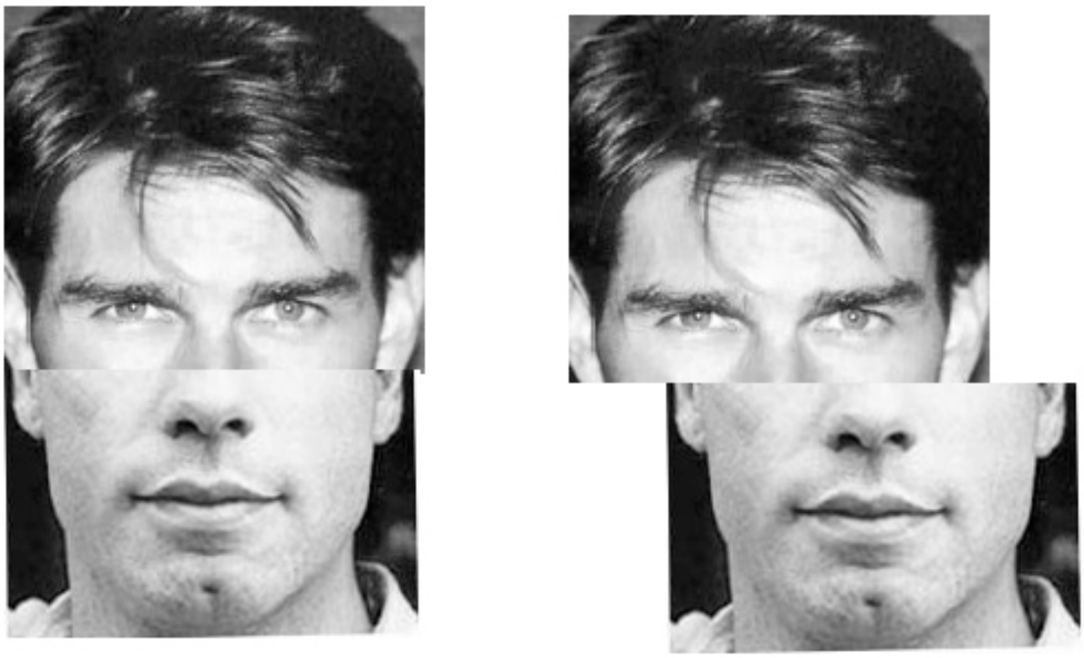

Figure 1.2: The composite face effect as reported by Young *et al* (1987). It is generally more difficult to identify the top (Tom Cruise) and bottom (John Travolta) halves in the image on the left than in the image on the right, where the face halves are misaligned. This effect demonstrates that face perception is driven by powerful tendency to combine elements into a unified perception.

Though there exists significant uncertainty regarding whether or not faces are processed 'configurally' (most likely due to difficulty in defining this term) the evidence that faces are processed 'holistically' is compelling (see Farah et al. 1998 for a review). In a classic demonstration of the inversion effect, Young *et al* (1987) measured participants accuracy at recognising images of familiar people showing either the top half or the bottom half of their face. When the top and bottom halves of two familiar faces were fused together performance on this task was very poor, however when the face halves were separated the ability to recognise the two identities was greatly improved (see figure 1.2). The implication of this result is that perceiving the face as a whole somehow disrupted the ability to perceive its parts. Importantly, this effect was not observed when stimuli were inverted; lending further support to argument that inversion disrupts the normal processing of faces.

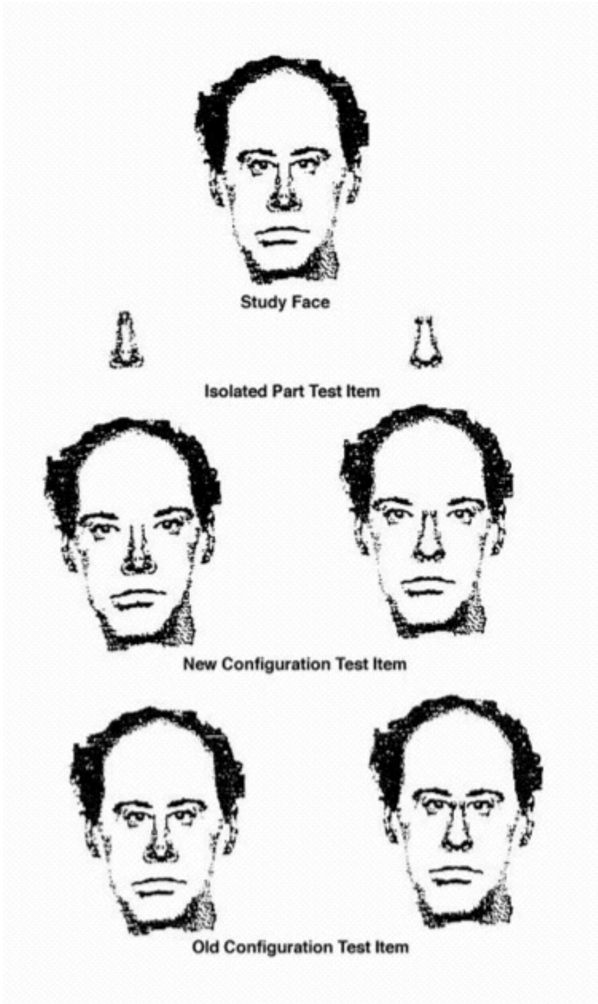

Figure 1.3: Stimulus conditions used by Tanaka & Sengenco (1997, Experiment 1). Performance was best when the feature appeared in the context of the original configuration (bottom row), and was better in the 'new configuration' relative to 'isolated' condition.

The tendency for visual information in faces to be combined in a holistic manner would appear powerful, and has been repeatedly demonstrated. Tanaka & Sengco (1997) found that individual parts of faces were recognized best when presented in the context of the original face configuration compared to when they were presented in the context of a new face configuration or in isolation (see Figure 1.3). As no such effect was found for their control stimuli, it is likely that processing of non-face objects is not as reliant on the perceptual whole (see also Farah et al. 1995; Tanaka & Farah, 1993). In addition, neurons in the temporal cortex of monkeys that show sensitivity to whole faces and to faces with masked features do not show any response to faces where features are presented in a scrambled configuration (see Desimone, 1991). As these neurons do not respond to the presence of facial features in an additive manner, it would suggest that they represent faces as a function of their global configuration. Furthermore, global processing of faces by rhesus monkeys has also been demonstrated behaviourally by Dahl and colleagues (2007), who report a standard human composite effect (Young et al. 1987) in this species.

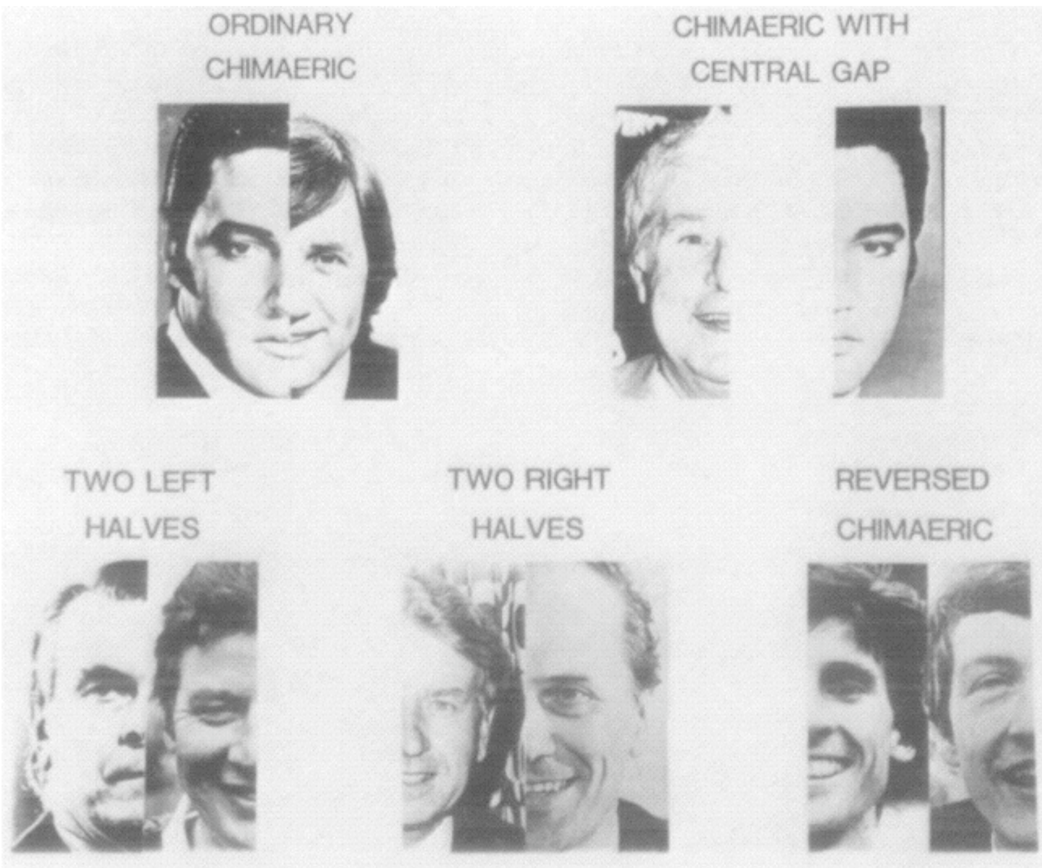

Figure 1.4: Stimulus conditions used by Young *et al* (1992). Patient B.Q. neglected the left side of the face stimuli only in the 'ordinary chimeric' condition.

The force of holistic perception in face perception is further evidenced by patient B.Q. who displayed symptoms of 'object-centred visual neglect' (Young et al. 1992). People with this condition neglect any information falling to one side of an attended object so that when viewing face images made by merging two face halves (bilaterally), only the face on the non-neglected side is perceived. However, this phenomenon does not occur when the half faces are presented with a gap between them suggesting that the patients' failure to recognise the left side of a face is restricted to instances where the face can be perceived as a whole (see figure 1.4). Furthermore, B.Q. neglected the left side of non-face objects significantly less than the left side of faces, which provides further evidence that faces are processed more holistically than other objects.

In addition to uncertainty regarding the manner in which faces are processed, there also exists significant debate as to whether these processes are exclusive to face perception. In a recent paper by Busey and Vanderkolk (2005) it was shown that fingerprint experts exhibit a similar inversion effect for faces and fingerprints on the electrophysiological N170 response. This data further supports the notion that the negative effect of inversion on configural processing is common to all 'objects of expertise' and replicates previous demonstrations of this phenomenon (e.g. Diamond and Carey, 1986; Gauthier & Tarr, 1997). Furthermore, brain-imaging studies have demonstrated that extensive training at discriminating between highly homogenous stimuli causes activation of neural networks implicated in face perception (e.g. Gauthier et al. 1999).

The issue of whether cognitive mechanisms involved in face perception are specific to this task remains highly contentious. In a recent review of behavioural data suggestive of similar processing of faces and other objects of expertise, Robbins & McKone (2007) contend that inversion effects for non-face objects are rarely demonstrated. In fact, in the twenty years since Diamond and Carey's (1986) initial result (that processing objects of expertise is disrupted by inversion), no research has replicated this effect. Similar effects of inversion on reaction times and physiological responses have been reported for face and non-face objects (e.g. Gauthier & Tarr, 1997; Gauthier et al. 1999), however inversion has never since been shown to reduce the accuracy with which objects of expertise are recognised. In addition, classical demonstrations of holistic processing in face perception such as the whole-part effect (Tanaka & Farah, 1993) and the composite effect (Young et al. 1987) have never been shown to occur in objects other than faces, regardless of the degree to which subjects were expert with the stimuli.

Robbins and McKone (2007) reproduced the conditions of Diamond and Carey's (1986) experiment and observed no inversion effect, no whole-part effect, and no composite effect for images of Labradors when testing Labrador experts. These results make it highly likely that Diamond and Carey's (1986) demonstration was artefactual, most probably resulting from the fact that their expert participants were overly familiarised with the test images (their stimuli were copied from a ubiquitous text that most dog lovers would be familiar with). In light of this research it would appear that face processing is indeed reliant on a form of perceptual processing that combines visual information in a qualitatively different fashion to the processing of non-face objects (see also McKone et al. 2006).

1.6 The effects of global geometric transformations on face processing

1.6.1 Stretching and Sheering

Though face processing is undoubtedly disrupted by inversion, it has been shown to be robust to certain affine transformations. Hole *et al* (2002) show that images of famous faces can be stretched to twice the original height without affecting participants' response behaviour, despite the profound alteration of the 'configural' mapping of the facial features (Figure 1.5; c). This manipulation grossly corrupts the true relationship between the facial features, and such stretching may be seen as an extension of the spatial distortion produced by a standard camera lens (where any 'stretch' is typically slight). Given that the shape information in the face is altered by this transformation (but not by inversion), it is surprising that it does not affect our ability to recognise the face, or the speed at which we do so.

Hole *et al* (2002) did find that 'shearing' an image (Figure 1.5; d) of a familiar person disrupted recognition performance however, as did stretching the image to twice its width, but these effects were slight and could not be considered catastrophic to recognition performance. Furthermore, where linear transformations were applied to only one half of the face, recognition performance suffered. Again however, the effect of this transformation on performance was not as great as might be expected given the nature of the distortion with performance dropping from a baseline of 93% to 79% accuracy. It would appear from this data that it is actually quite difficult to disrupt face recognition by distorting positional information, and recognition would appear robust even to severe levels of geometric distortion (see Figure 1.6).

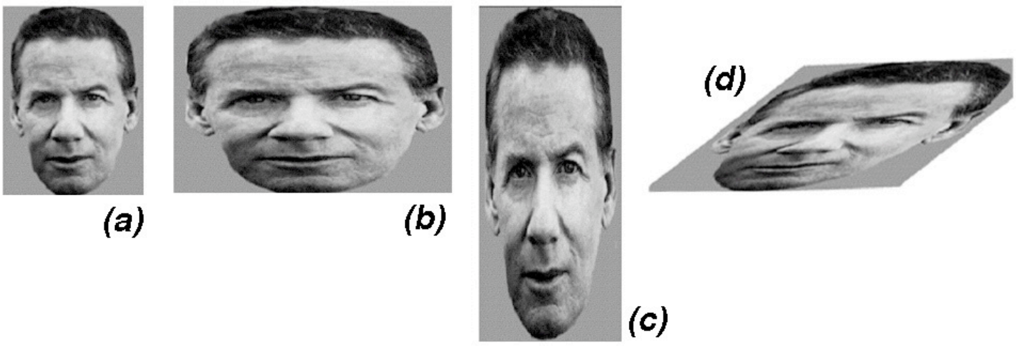

Figure 1.5: A veridical representation of a familiar face (a) can be subjected to linear distortion (b, c) without effecting recognition behaviour. Illustration reproduced from Hole *et al* (2002).

Figure 1.6: The reader is invited to identify the two images, which have been subjected to a severe degree of non-global distortion.

As discussed in the previous section, the distinction between configural and featural processing is somewhat arbitrary in the context of human face processing, and the results reported by Hole *et al* (2002) complicate this issue further. Not only do the geometric distortions executed on their test stimuli grossly alter the inter-feature distance measurements, but they also drastically change the appearance of the features themselves. Furthermore, and perhaps most critically, these distortions also play havoc with the distance *ratios* across x and y coordinates. Such data would appear to preclude any explanation of configural processing that posits a simple x-y coordinate matching system as its basis (e.g. Cooper & Brooks, 2004). In addition, these findings highlight the need for a revised explanation of face recognition: one that relies not on cognitive representations of simple distance metrics at its foundation, but on a more sophisticated and flexible representation tolerant to systematic distortions in configural relations.

Discussing their findings, Hole *et al* (2002) offer two explanations for the surprisingly slight effect of their experimental manipulations. The first is that the 'reverse transformation' was applied to the stretched stimuli prior to recognition, and that this normalisation was easily applied to images stretched in only one dimension. The existence of uni-dimensional normalisation in face processing is implicated further by reported dissociations between configural information in horizontal and vertical axes. Using factor analysis, Fellous (1997) demonstrated that the correlation between measurements of facial configuration is explained by groups of dimensions belonging to either vertical or horizontal measurements (but not to both) and that horizontal and vertical facial measurements are therefore statistically independent. Furthermore, there is more between-face variance in the horizontal dimension relative to the vertical dimension of faces (Ferrario et al. 1997). In addition, it has recently been suggested that 'iso-dimensional' normalisation may have developed to tolerate rotations of the head in depth around the x and y axes, which causes a similar compression of configural information to geometric stretches (Sinha et al. 2006).

A second proposed explanation for Hole's findings is that the cognitive representation of the perceived face is transformed to match the incoming stimulus. This explanation implicates the existence of a mutable representation that is capable of altering its state in accordance with the appearance of the familiar stimulus, and as such predicts that stretching should disrupt unfamiliar face matching. As yet however there has not been a successful demonstration of disrupted performance on an unfamiliar matching task. A recent study reported null effects using this paradigm (Hole, unpublished data), suggesting that the perceptual invariance is a result of some automatic process of normalisation that is not contingent upon an existing representation of the perceived individual.

1.6.2 Inversion and Rotation

As has already been discussed, reversing an image of a face across the horizontal axis severely impairs normal face processing (e.g. Yin, 1969), and this effect is likely to be caused by disruption to normal holistic processing (e.g. Farah et al. 1998). Since Yin (1969) found that faces are more difficult to recognise upside down than they are the right way up, the 'inversion effect' has been repeatedly demonstrated and has become perhaps the single most important finding in the field of face perception. The effect of inversion is extremely robust and disrupts many aspects of normal face perception including emotion processing (Calder et al. 2000), perception of attractiveness (Little & Jones, 2006) and gender discrimination (Bruyer et al. 1993; Stevenage & Osborne, 2006). Historically, it has been considered that inversion leads to a qualitative shift in the way faces are processed (e.g. Valentine, 1988), however recent reports have cast doubt over this assertion.

It has been argued that inversion causes disruption to face processing because of difficulty in mentally rotating inverted stimuli to their canonical (upright) orientation (e.g. Rock, 1973) and there exists a significant body of data supporting this theory. Valentine and Bruce (1988) were the first to show that face recognition performance varies as a function of the angle of rotation, and since then there have been further demonstrations of this monotonic relationship (Bruyer et al. 1993; Collishaw & Hole, 2002; Stevenage & Osborne, 2006). By demonstrating the linear relationship between angle of rotation and recognition using blurred faces, Collingshaw & Hole (2002) demonstrated that configural processing is increasingly disrupted as the image of the face is oriented away from upright. This finding is evidence against any 'processing shift' caused by inversion but instead points towards a difficulty in retaining the ‘facial gestalt’ during the process of mental rotation. Therefore, it would appear that the ‘inversion effect’ may not be the product of inversion *per se*, but instead is the cumulative effect of angular distance from vertical orientation. One may even argue, in light of this data, that the ‘inversion effect’ would be more accurately named the ‘rotation effect’.

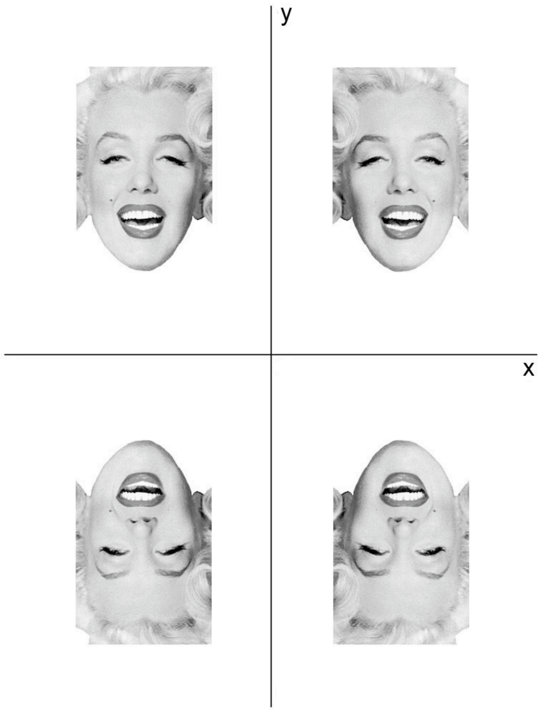

Figure 1.7: The original image (in the upper right quadrant) has been subjected to reflection across the x and y axes. N.B. reflection across the x-axis (classical inversion) comprises of lateral and vertical reversal, as evidenced by the location of Marilyn Monroe’s beauty spot.

It should be noted however that rotation transformation is different to inversion (i.e. vertical reflection), and it produces different stimuli (see Phillips & Rawles, 1997). Whereas the classical inversion manipulation reflects information across an axis perpendicular to the axis of symmetry, rotation instead rotates the face through 180 degrees. Thus, the resulting stimuli from these manipulations are mirror reversals of one another (see figure 1.7), and it has been demonstrated that the effect of inversion is confounded by the effect of lateral reversal in a memory task (McKelvie, 1987). If inversion is indeed confounded by mirror reversal, then this would pose a problem for those researchers that argue the disproportionate disruption caused by inversion on face processing relative to processing of non-face objects is of great theoretical importance (e.g. McKone et al. 2006). Object recognition is invariant to mirror orientation (e.g. Biedermann & Cooper, 1991), so if face recognition is negatively affected by mirror reversal, then this may explain why inversion impedes face recognition more than it does object recognition.

1.6.3 Mirror Reflection

The effect of lateral reversal on face recognition has been subject to a limited amount of research relative to the extensive use of the classic inversion manipulation. This is due to the perceived importance of the theoretical distinctions associated with the 'inversion effect', which is considered to selectively disrupt mechanisms intrinsic to face recognition (e.g. McKone et al. 2006). The present thesis will look at the effect of lateral reversal in detail and will try to resolve the question of whether or not asymmetry in the configuration of facial landmarks is explicitly coded in memory. This will inform current debate in face recognition research concerned with whether or not face recognition and symmetry detection are functionally dissociable (e.g. Rhodes et al. 2005), and also elucidate the contribution of facial configuration to face recognition.

It is considered that whereas reversal across the horizontal axis alters the perception of both first and second order configural information (see Maurer et al. 2002), lateral reversal preserves the former whilst altering the latter. Therefore, it is expected that disruption of second order configuration will disrupt the task of recognition. The effect of mirror reversal on face recognition was first investigated by Mita, Dermer and Knight (1977). They sought to replicate the Mere Exposure Effect (Zajonc, 1968) using an affective 'likeability' decision to faces in their original and reversed orientations. Subjects were shown an image of either themselves or of close friends in both original and reversed lateral orientation and had to decide which image they 'liked' better. Their data showed that people are more likely to prefer images of themselves in the mirror orientation (which is the orientation most often available to our perception), and that the opposite pattern was true when making likeness decisions to close friends.

Brédart (2003) replicated this finding using a more explicit test of orientation memory. He asked participants to indicate which of two images (original or reversed) were shown in the orientation most familiar to them for both images of co-workers and of themselves. Twenty-eight out of thirty-two participants correctly chose the veridical representation of their co-workers and twenty-four out of thirty-two correctly chose the mirror reversal of their own face. This data is in agreement with a previous study by Rhodes (1986) showing that subjects are more likely to choose the veridical orientation when asked to indicate which image is the best 'likeness' of the person.

Given that the cognitive system appears to be selectively responsive to faces in their most commonly perceived lateral orientation, one might expect to find a detectable effect of mirror reversal on the recognition of familiar faces. As yet however, there is a very limited amount of research supporting this assertion. McKelvie (1983) tested participants' memory for unfamiliar faces and found that, if the images were reversed prior to test, memory performance was disrupted. Though McKelvie replicated these findings in 5 separate experiments, all of these experiments used the same images at learning as in test. This somewhat constrains the scope of the conclusions, as it has previously been demonstrated that memory for images of faces may be achievable without using mechanisms involved in human face recognition (Dyer et al. 2002).

A more recent study by Brooks *et al* (2002) used a repetition-priming paradigm to determine whether or not our memory representations for faces are sensitive to the lateral orientation of the perceived stimulus. They found no difference in response latency between original and reversed conditions. This contradicts McKelvie (1983) and suggests that the mental representations underpinning face recognition are coded in a manner that transcends the left-right orientation of the external world. This is an appealing possibility, as the same invariance to lateral orientation has previously been demonstrated in object recognition (Biederman & Cooper, 1991). Furthermore, this data is consistent with models of object recognition proposing a mode of representation that is positionally invariant (e.g. Hummel and Biederman, 1992).

1.7 Thesis Overview

The empirical approach employed in this thesis is characterised by the stimulus manipulation of mirror reversal. The effect of this geometric distortion on a number of face processing tasks will be investigated in an attempt to improve understanding of how faces are cognitively represented. Lateral orientation does not appear to be specified in memory representations of other objects (e.g. Hummel and Biederman, 1992), but faces may be exceptional in this respect as asymmetry in the face provides information enabling powerful identity discrimination (e.g. Burke & Healey, 1993), and facial asymmetry conveys important social information (e.g. Rhodes, 1999).

The first experimental chapter will assess the contribution of bilateral information in an unfamiliar face-matching task. Following that, chapter 3 will examine the extent to which the left-right orientation of facial asymmetry is represented in memory by quantifying the effect of mirror reversal on various face recognition tasks. Chapters 4 and 5 then employ a two alternative forced-choice procedure (2AFC) where an image of a face is presented alongside its mirror-reversal and participants must decide which image is presented in the 'real-world' orientation. This procedure is considered to be a sensitive test of memory for lateral orientation of familiar objects (e.g. Kelly et al. 2001) and will be used to assess whether memory for the lateral orientation of familiar faces can be detected (chapter 4). A slightly different 2AFC procedure is used in chapter 5 where participants are asked to make a 'likeability' decision (see Zajonc, 1968) to simultaneous presentation of images of *unfamiliar* faces in normal and reversed orientations. It has been demonstrated that faces are generally asymmetric (e.g. Farkas & Cheung, 1981), and this may produce cognitive sensitivity to the mirror reversal of unfamiliar faces. In the final experimental chapter, two experiments are designed that demonstrate facial configuration is *processed* asymmetrically. In addition to asymmetrical morphology, perceptual asymmetries may contribute to the asymmetry of cognitive representations, and the interaction between these two factors is discussed in chapter 6.

Chapter 2

Facial asymmetry and unfamiliar face matching

Introduction

The typical human face displays mild structural asymmetry along the vertical meridian, with average distances differing between right and left measurements by around 3 percent (Farkas & Cheung, 1981). In a more recent study, average measurements from the facial landmarks to centre points were found to vary from 4% to 12% depending on the landmark measured (Ferrario et al. 2001). Even in studies using samples with highly symmetrical facial features (e.g. professional models, beauty contest winners), researchers have failed to find an instance where a subject did not demonstrate asymmetry in one or more of the dimensions measured (Peck et al., 1991). Presently, this naturally occurring asymmetry is of theoretical interest primarily to evolutionary biologists (e.g. Parsons, 1990; Polak, 2003), and to psychologists studying the perception of beauty (e.g. Thornhill & Gangestad, 1994; Gangestad & Thornhill, 1997; Simmons et al, 2004; for a recent review see Rhodes, 2006). However it is surprising that research concerning the cognitive process of face recognition has rarely focused on this source of variation.

It has been argued that cognitive mechanisms specialised to compute fine-grain configural information have evolved specifically to process faces (Le Grand et al 2004, McKone, Kanwisher & Duchaine 2007, Tanaka & Farah 1993). Furthermore, it would appear that this 'configural' information is integral to both our memory and perception of faces, with processes requiring sensitivity to the fine-grain spatial relationships between facial features having been heavily implicated in tasks involving recognition of familiar faces (Sergent 1986, Hayes 1988, Collishaw & Hole 2000, McKone, Martinin & Nakayama 2001) and in face discrimination tasks with novel stimuli (Freire et al. 2000, Le Grand et al 2001, Leder et al 2001, Leder & Bruce 2000). Given that humans show a high level of sensitivity to deviations from perfect symmetry in faces, and that the degree to which this source of asymmetry is present in a given face is thought to vary moderately between human individuals (Simmons et al. 2004), it is likely that the holistic property of asymmetry individuates facial identity. As such, a central hypothesis to be tested in this thesis is that the quality (and quantity) of this asymmetry contributes to the visual information that makes a given face unique.

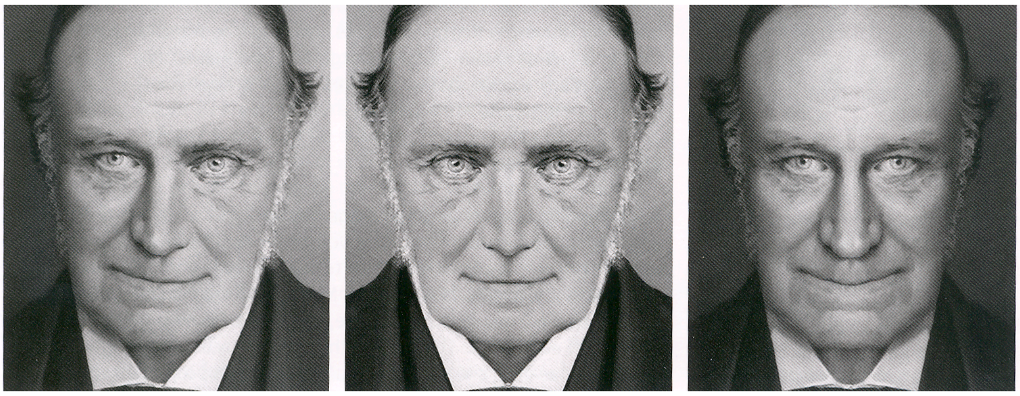

Figure 2.1: A portrait of John Innes (left) has been used to create a right (middle) and left (right) 'chimera', emphasising the differences between the left and right sides of a face (Taken from Roberts, 2006).

One can illustrate the asymmetrical nature of a human face by creating what are known as 'chimeric faces' (see figure 2.1). These images are generated by firstly generating a mirror reversed copy of a straight-on portrait, then marrying the left half of one image with the right half of the other (reversed) image, and then by fusing the remaining hemifaces in a similar manner (see Wolff, 1933). The resulting images (the middle and right pictures of figure 2.1) are strikingly distinct: So much so that they typically appear to depict two separate identities altogether. Chimeric faces have been used to investigate the asymmetry of perceptual systems responsible for face recognition (Burt & Perrett, 1997; Butler et al. 2005), and these findings will be discussed later in the thesis (chapter 6). For the moment however, they are relevant only in as far as they highlight an important perceptual consequence of facial asymmetry.

Here I design two experiments to assess the extent to which the configural information that is distributed across the two hemispheres of the face is implicated in successful performance on an unfamiliar face-matching task. Previous research has shown that information specifying relations between the facial features is important when confronted with the task of individuating previously unfamiliar faces. Using an unfamiliar face discrimination task, Freire *et al* (2000) found an inversion effect in the order of 30% when faces differed primarily in configural information, but no effect of inversion whatsoever when discriminating faces that differed mainly in featural information. Though appearing to implicate configural processing in unfamiliar face matching, this conclusion has been challenged by more recent evidence. Megreya and Burton (2006) show that performance on upright unfamiliar face matching is highly correlated with performance on the same task when stimuli are inverted; a result which is not consistent with the view that these two tasks rely on independent processing mechanisms.

A further sensitivity to bilaterally distributed configural information in tasks involving unfamiliar faces was demonstrated by Leder *et al* (2001). Participants were shown two images of unfamiliar faces side-by-side and asked to decide in which face the interocular distance was greatest. They found that subjects performed this task better when the faces were presented in upright orientation relative to inverted orientation, and that this effect occurred even when the eyes were presented in isolation. Again this illustrates that some degree of configural processing occurs during discrimination tasks involving unfamiliar faces, and also that this configural sensitivity extends to relationships between the two facial hemispheres. In the two experiments reported in this chapter I investigate more closely the importance of this bilaterally distributed configural information in an unfamiliar face discrimination task.

It is subjectively apparent when looking at figure 2.1 that the two chimeric images appear to be portraits of two different people (if one ignores the hairstyle and clothing), an effect that may be predicted given the fact that the two hemi-faces are morphologically distinct. This observation prompts a testable hypothesis: When shown two opposing hemi-faces (one left, one right) people will be unable to reliably decide if the hemi faces belong to the same person, or to two different people. So the first experiment will investigate the extent to which the visual information contained in opposing hemifaces is perceptually distinct. That is whether simultaneous presentation of the two sides of the face (each from a different image of the person) will result in participants perceiving two separate identities. If it transpires that this is the case, then it can be concluded that the coherence afforded by the facial gestalt is not a result of any intrinsic similarity between the two sides of the face, but rather this holistic perception is constrained by familiarity with the face in question.

Previous research has shown that when target and 'line-up' array images are captured using different cameras, successfully picking out a target image from an array of ten images can prove difficult, with accuracy typically in the region of 70% (Bruce et al. 1999). Further studies have shown that this task remains difficult even when the array is reduced to just a target image alongside one of the array images, with subjects being forced to decide if the images are of the same person, or of two different people (Megreya & Burton, 2006). In this situation mean accuracy is around 80%, which given that chance performance in this task is 50% (relative to the 10% level in the 1-in-10 task) represents a meagre improvement.

The difficulty of the unfamiliar face-matching task is considered to stem in part from the distorting effect that camera lenses have on their subject (Burton et al. 2005). Lenses of different focal length will produce qualitatively different projections of the subject, and this distortion can impede our ability to identify two faces as belonging to the same person. We encounter such distortions in our everyday life through our experience with both static and moving images, and our ability to recognise faces in spite of these various perversions is testimony to the sophistication of our minds face recognition systems. Modern day wisdom has it that television cameras 'add two pounds' to the weight of those who fall victim to their distortion, yet this degree of misrepresentation does not disrupt our ability to recognise this person. Furthermore, it has been demonstrated that face recognition is robust even to fairly profound degrees of geometric alteration. Stretching an image of a familiar face to twice its original aspect ratio in either the horizontal or vertical dimension leads to no detectable effect on recognition (Hole et al. 2002).

Though robust to many distortions, one can disrupt both familiar (e.g. Hole et al. 2002) and unfamiliar (Freire et al. 2000, Megreya & Burton 2006) face processing by simply turning an image upside down. This effect has been taken to reflect a disabling of mechanisms responsible for configural processing, thus highlighting the importance of this information in face processing tasks. Previous demonstrations of the inversion effect in the context of unfamiliar face processing (Freire et al. 2000; Megreya & Burton 2006) have shown performance to be extremely poor. In a second experiment I investigated the effect of an analogous manipulation by mirror reversing the target image. In this case the bilaterally distributed configural information (which is unique to the individual) has been reversed, but the facial schema (common to all faces) has not been violated, as in the case of inversion. This said it is not anticipated that this manipulation will impede normal performance on a face matching procedure as profoundly as has been reported in studies employing an inversion manipulation. Lateral reversal will not disrupt processing of configuration (as first order relations are unaffected), but it will alter the information pertaining to the identity of a face (as second order relations are reversed). Therefore it is anticipated that performance on the face-matching task will be disrupted by lateral reversal, because the two images will be more likely to be perceived as two different people.

In the final experiment of this chapter a somewhat separate issue regarding unfamiliar face matching is addressed. Here I investigate whether performance on a task where participants had to decide whether one image in a pair was a mirror reversal of the other (or whether they were presented in identical orientation) correlated with unfamiliar face matching performance. The 'mirror discrimination' task used in this experiment forces participants to make image-level comparisons between two images presented adjacent to one another. It is hypothesised that this strategy is similar to the 'image matching' strategy used by participants when performing unfamiliar face matching decisions (Megreya & Burton, 2006), and that performance on these two tasks will therefore be highly correlated. Furthermore, it has been argued that mirror discrimination ability may be related to our ability to make within-category discriminations (Rentschler & Jüttner, 2007), and this argument would also predict performance on these two tasks to be correlated.

Experiment 1: Matching Unfamiliar Hemifaces

Introduction

Unfamiliar face matching is thought to be characterised by a matching strategy (Mergeya & Burton 2006) that attributes particular reliance on the external features of a face (Bonner et al. 2003). Others have argued that configural information is implicated in unfamiliar face matching (Freire et al., 2000; Leder et al., 2001).

Regardless of the process used to achieve normal performance however, it is anticipated that the task of matching given from two opposing sides of the face will be extremely error-prone. The asymmetry of a face not only ensures that there will be no direct mapping between the features of the two images in this condition, but also that relationships between the features of face-halves of 'same' pairings will be incongruent. It is nevertheless predicted that performance in this condition will be superior to chance, as the visual characteristics of external cues (i.e. hairstyle), are somewhat transferable across hemifields. In addition, it can be reasonably assumed that although the two hemi-faces may be morphologically distinct they nevertheless are more similar in form than the hemifaces of two separate identities.

It was further hypothesised that when presented with corresponding hemifaces (i.e. l-l or r-r pairings) subjects' performance will be much improved relative to the opposite hemiface condition, and may even be equivalent to the whole-face control condition. In the l-l and r-r conditions the facial features and configuration are correspondent between the two images so any detrimental effect could be explained by a reduction in the visual information presented to participants.

Method

Stimuli and Materials

One hundred and sixty face-image pairs were selected from a database used in previous studies (see Bruce et al. 1999 for a full description). These image pairs were created from images of eighty trainee policemen ranging between 18 and 35 years of age. For each of the eighty identities I created a matching pair, where two images of the identity (taken using different cameras) were presented side-by-side, and a mismatched pair, where the identity was presented alongside an image of a similar identity. Similar identities had been previously identified using a card-sorting procedure outlined by Bruce *et al* (1999).

The experimental manipulation was implemented by removing pixel information either to the right of the axis of symmetry in each image (left hemiface condition), or to the left of the axis of symmetry (right hemiface condition), or from the opposing sides in each image (opposite hemiface condition). All the face-pairs were subjected to these manipulations giving a complete set of 160 original pairs (80 same/80 different), 160 left hemiface pairs, 160 right hemiface, and 160 'opposite' hemiface pairs from which experimental stimuli could be drawn. Figure 2.2 gives examples of the stimuli used.

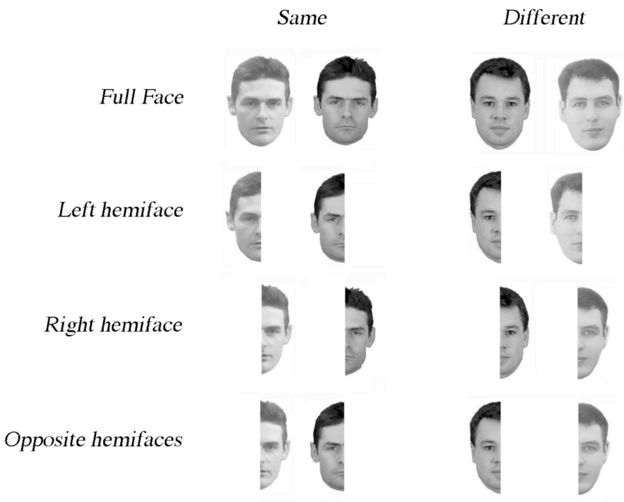

Figure 2.2: Example stimuli for experimental conditions. N.B. For "opposite hemiface" condition 20 trials (10 same/ 10 different) were with right hemiface presented right of centre (see 'different' in figure) and 20 trials were with right hemiface left of centre (see 'same').

Design and Procedure

24 subjects from the student population at the University of Glasgow consented to participate in the experiment. There were 9 males and 15 females (mean age= 22.1) and all received cash payment as reward for their participation.

Subjects were sat at a Macintosh workstation running the experimental software PsyScope (Cohen et al. 1993). Face pairs appeared one at a time (preceded by a fixation cross for 750ms) and remained on the screen until subjects made their response. Subjects were instructed to respond as to whether the two images on the screen were of the same person or of two different people. There were 40 (20 same and 20 different) trials per condition and counterbalancing ensured that, across subjects, each face-pair appeared in each condition an equal amount of times. In the 'opposite hemiface' condition half the image pairs (10 same and 10 different) appeared with the right hemiface on the right hand side of the screen and half with the right hemiface on the left.

The experiment lasted approximately ten minutes and there were no breaks. After completion of the task subjects were debriefed as to the aims of the experiment and the experimental manipulations used.

Results

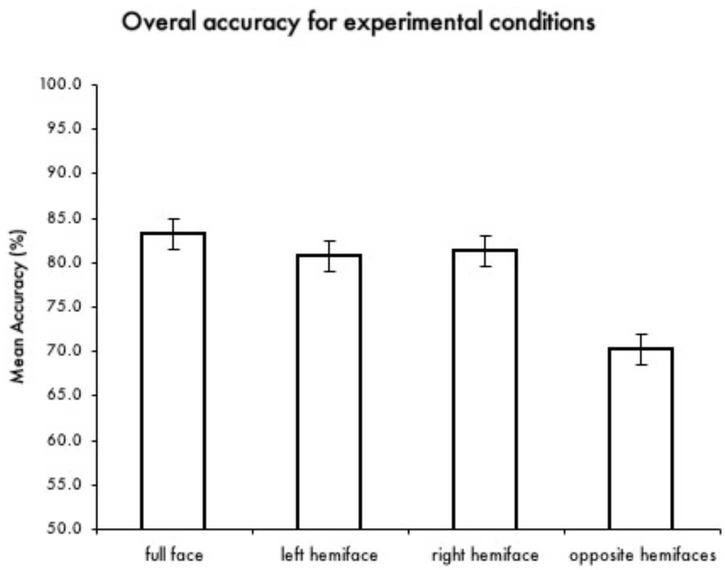

Figure 2.3: Mean accuracy in face matching task (error bars denote 95% confidence intervals which were calculated using the method described by Loftus & Masson, 1994).

A one-way within subjects ANOVA showed the factor of experimental condition to be statistically reliable [$F(3,23) = 11.46$; $p<0.01$]. Post-hoc Tukey's HSD tests confirmed that performance on the opposite hemiface condition was significantly poorer than in the full face condition ($p<0.001$), and was also poorer than in the remaining two experimental conditions where left hemifaces ($p<0.001$) and right hemifaces ($p<0.001$) were presented. There were no other significant differences in accuracy data.

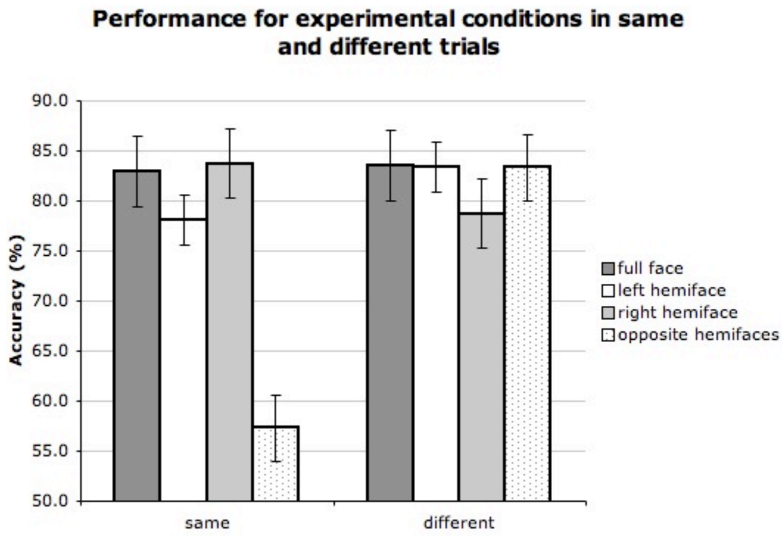

Figure 2.4: Mean accuracy for same and different trials (error bars denote standard error)

A two-way within-subjects ANOVA was performed on the data with "trial type" (same/ different) and experimental condition (full face/left hemiface/right hemiface/opposite hemifaces) as factors. This confirmed the main effect of experimental condition to be reliable [$F(3,23)=11.46$; $p<0.01$]. The interaction between factors can also be reported as reliable [$F(3,23)=13.47$; $p<0.01$], however there was no significant effect of trial type. Analysis of simple main effects showed the effect of condition only to be reliable for trials in which the images were of the same person [$F(3,23)=26.11$; $p<0.01$]. Furthermore the effect of trial type was only reliable in the opposite hemiface condition [$F(3,23)=10.79$; $p<0.01$].

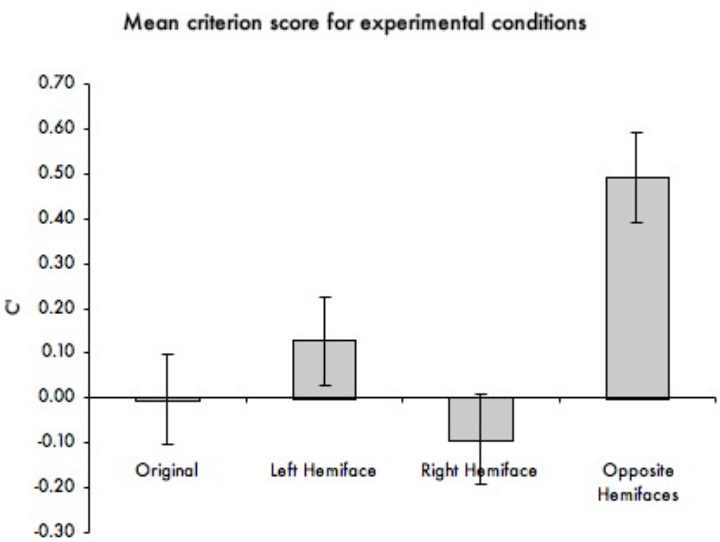

Figure 2.5: Mean response bias for experimental conditions (error bars denote standard error)

The interaction between stimulus congruency (same identity/ different identity) and experimental condition indicated a difference in response bias between experimental conditions. To investigate this more closely criterion scores (C') and sensitivity (d') scores were calculated for each of the subjects and analysed using two one-way ANOVAs. For criterion values the main effect of experimental condition was found to be reliable [$F(3,23)=11.41$; $p<0.01$]. Post-hoc Tukey's HSD tests revealed a significantly different response bias in the opposite hemiface condition relative to full face ($p<0.001$), right hemiface ($p<0.001$) and left hemi face conditions ($p<0.01$). The significantly larger (positive) criterion value in the opposite hemiface condition is indicative of a bias towards 'different' responses in this condition. Although the difference between matching left hemifaces and right hemifaces indicates that different response strategies are used in these two conditions, this difference was not found to be significant ($q=2.92$; $p>0.05$).

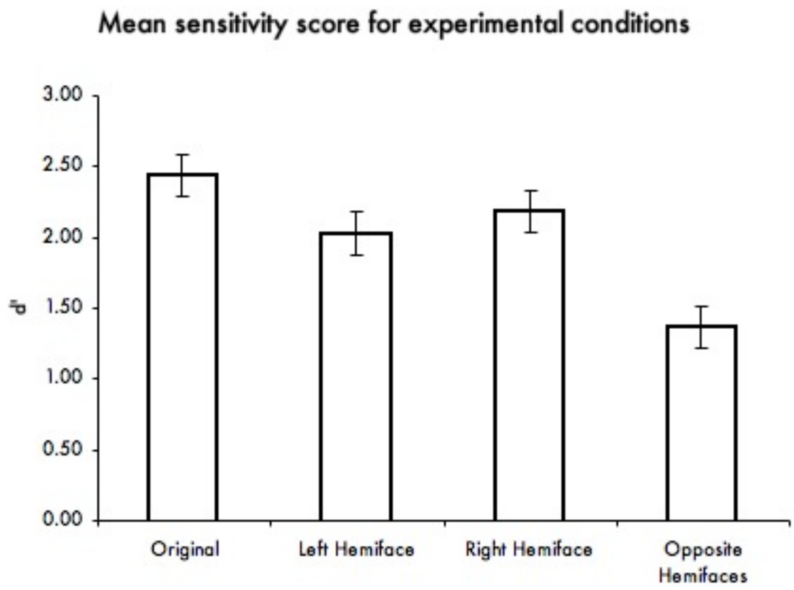

Figure 2.6: NB Error bars denote standard error.

A reliable main effect of experimental condition was also detected in the measure of response sensitivity (d') [$F(3,23)=10.45$; $p<0.01$]. Post-hoc Tukey's HSD tests showed that sensitivity in the opposite hemiface condition was significantly less than in the three other conditions (fullface, $p<0.001$; left hemifaces, $p<0.01$; right hemifaces, $p<0.001$]).

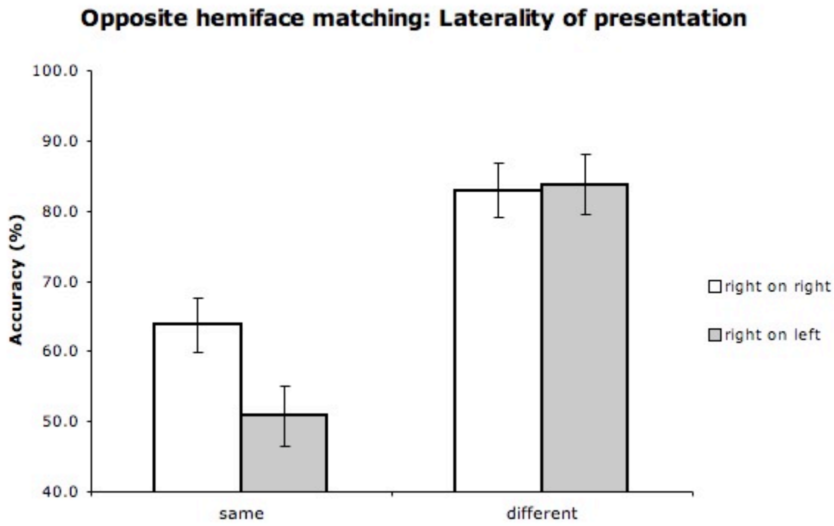

Figure 2.7: Mean accuracy data for opposite hemiface condition (error bars denote standard error)

A two-way within-subjects ANOVA was performed on data from the 'opposite hemifaces' condition with "trial type" (same/ different) and experimental condition (right hemiface to right of centre/ right hemiface to left of centre) as factors. The main effect of experimental condition was not found to be reliable in the context of this analysis [$F(1,23)=2.291$; $p>0.1$]. There was however a reliable main effect of trial type [$F(1,23)= 23.109$; $p<0.001$]. Although the interaction between factors was not statistically significant there was a trend apparent [$F(1,23)= 3.507$; $p=0.073$], with the simple main effect of condition only being significant in trials where the images were of the same person [$F(1,23)= 6.675$; $p<0.05$]. This result would appear to indicate that participants found it easier to cohere the two hemifaces into a whole face when they were placed in a natural spatial orientation (i.e. left on left and right on right), but that this compositional factor did not influence 'different' judgements. It is considered that this result is reflective of the greater ease with which subjects could perceptually 'combine' the halves into a unitary representation in this condition.

Discussion

There was no significant advantage for matching whole face relative to corresponding hemifaces in the unfamiliar face discrimination task. This finding indicates that bilaterally distributed configural information may not be important when discriminating between unfamiliar faces. This bilaterally distributed information is not exclusively configural however. Because only half of the nose and mouth is visible when occluding half the face, featural information is equivalently, if not additionally obscured by the manipulation. Indeed, it is surprising that this manipulation does not lead to a significant effect on behaviour and this shows that unfamiliar face matching can be performed just as accurately based on half of the available information. That normal discrimination occurs in the context of this constraint suggests that this task may rely on matching of small regions of visual information, as has been previously proposed by Megreya and Burton (2006).

The principal hypothesis was that in an unfamiliar matching task, showing participants opposite sides of a persons face would reduce performance relative to full-face and correspondent half-face presentations. This was found to be the case with performance in the opposite hemiface condition being roughly 12% worse than in the other three conditions. Though matching performance was far greater than chance (50%) when opposing hemifaces were presented, this decrease in performance is nonetheless large and is reflective of both a decreased sensitivity and an increased tendency to make ‘different’ responses. The reduced sensitivity and elevated criterion characterising response behaviour in the opposite hemiface condition shows that the small asymmetries present in the human face can nonetheless give rise to large changes in behaviour. Participants were commonly mistaking two images of the same person as being two different people, presumably because the visual information in one side of a face does not predict the visual information in the other side of the same face with sufficient accuracy to elicit a ‘same’ response.

Though performance in the opposite hemiface condition was reduced relative to the other experimental conditions, and despite the significant alterations of response behaviour induced by this manipulation, participants were still fairly accurate in deciding whether or not the two images were of the same person. With an average accuracy in this condition of 70% (sd=9.8), the effect of the manipulation could hardly be considered catastrophic. A question therefore arises: what transferable information is there in the left side of a face that can be matched with information on the right half of a face?

It is unclear what information subjects were basing their responses on in the opposite hemiface condition. Prior to the experiment I had considered that the external features of the faces would enable successful performance on the matching task. This is because aspects of their hairstyle and face-shape may be more easily transposed onto the opposite hemiface than an analogous 'reflection' of the internal features, which are more visually complex by nature. Even if this assumed to be the case however, it is not clear whether this would affect task performance positively, or whether (due to the relative homogeneity of the stimulus set) it would impede performance. As all the images were of young male police officers the hairstyles tended to be similar in length and style, this cue may have been misleading. However, if this was the case, then one would expect response behaviour to be biased in the opposite direction to that which I report here (i.e. an increased tendency to make 'same' responses).

Therefore it is more likely that the above chance performance is due to the similarity of the facial features across hemifields, or of the face shape, or of the configuration between the facial features. That the two sides of a face are morphologically distinct does not perhaps give ample reason to believe that they cannot be recognised as two sides of the same face. After all, there is no direct manner in which the two images shown in their entirety relate to each other either. If their 'visual templates' did overlap in an exact manner then performance would be higher in this condition than the 84% reported here.

The difficulty of this task is by virtue of the fact that the images are taken from different cameras (both distorting the aspect ratio in a qualitatively different manner) and that they vary to a small extent in the angle of the face with respect to the camera. This being the case, it may be that performance on this task involves a process that computes the similarity of the two images (or regions within the image), and that a 'same' response is dependent on this perceived similarity exceeding a certain threshold (see Burton et al. 1990). Thought of in this way, it is possible to see how two face-halves that are divergent in their morphology may yet be similar enough to elicit a 'same' response with sufficient frequency to produce a performance level of the order reported for the opposite hemiface condition.

This explanation has been previously proposed by Troje & Bülthoff (1998), who argue that viewpoint generalisation across mirror symmetric head rotation is achieved by mirror reversing a stored representation and comparing this to a test image. Because a high level of object constancy is displayed between opposing three-quarter views of faces (Troje & Bülthoff, 1998; Hill, Schyns & Akamatsu, 1997), it is considered that a mirror reversal of presented stimuli is readily produced by cognitive systems, and that this representation approximates well to the novel three-quarter view. In light of the results reported here, it would appear that this object constancy is not as reliable when mirror approximations are made on the basis of one face-half. The fact that photographs of real people taken with real cameras were used as stimuli here (as compared with the computer generated stimuli used by Troje et al.) may also have contributed to the difficulty of this task.

Experiment 2: The Effect of Mirror Reversal on Unfamiliar Face Matching

Introduction

In experiment 1 it has been shown that one can induce the perception of two separate people by showing participants opposing halves of an unfamiliar persons face. This is by virtue of the fact that the two halves of any given face are morphologically and perceptually distinct. Given that facial asymmetry has been shown to be profound enough to elicit such effects, an experiment was designed to test whether or not the same effect could be induced (i.e. a greater proportion of 'different' responses) by mirror reversing one of the images in the face-pair. Successful performance in this task cannot be achieved by matching the configurational information, as this is essentially asymmetric, and has been reversed. Furthermore, a feature matching strategy would also lead to increased error, as facial features are also typically asymmetric (e.g. Farkas, 1994).

In the same way that the classic inversion manipulation alters the configurational template common to all faces (i.e. two eyes above a nose, which in turn is above a mouth etc.), lateral reversal alters the configurational template describing that individual face. Therefore it is predicted that the mirror-reversal of one image in a matched pair will result in an increase in ‘different’ responses to that pair.

Method

Stimuli and Materials

A randomly selected subset of sixty identities were selected from the eighty used in the previous experiment. With 60 same pairs and 60 different pairs, there were 120 face pairs used in total. Images were presented in greyscale on a 1024 X 860 display with typical dimensions of 500 by 350 pixels, with each image in the pair being separated (measuring from the centre) by 250 pixels (see figure 2.8).

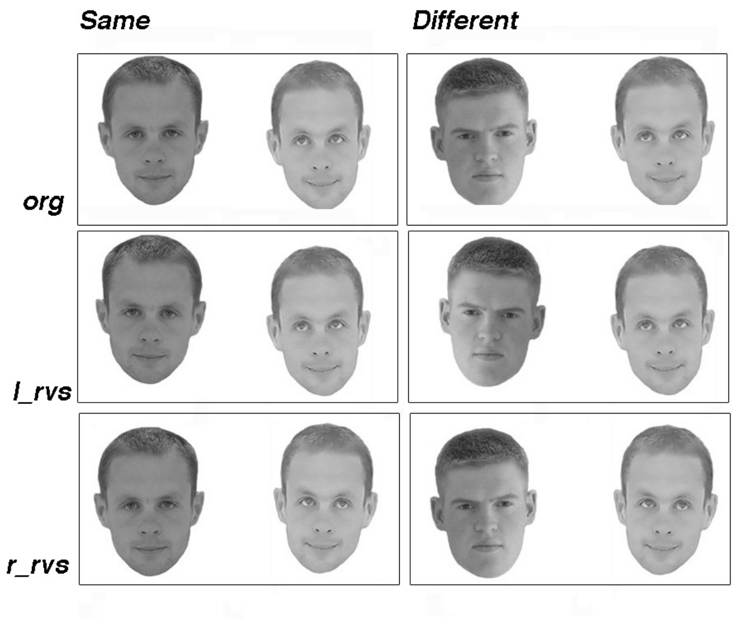

Figure 2.8: Example stimuli for experimental conditions

The experimental manipulation was implemented by mirror reversing either the left image or the right image. All the face-pairs were subjected to these manipulations giving a complete set of 120 original pairs, 120 left-reversed pairs and 120 right-reversed pairs from which experimental stimuli could be drawn.

Design and Procedure

Thirty-six subjects from the student population at the University of Glasgow consented to participate in the experiment. There were 17 males and 19 females (mean age= 21.4) and they were all paid in cash.

Subjects were sat at a Macintosh workstation running the experimental software Psyscope. Face pairs appeared one at a time (preceded by a fixation cross for 750ms) and remained on the screen until subjects made their response. Subjects were instructed to respond as to whether the two images on the screen were of the same person or of two different people. There were 40 trials per condition (20 same/20 different) and counterbalancing ensured that, across subjects, each face-pair appeared in each condition an equal amount of times.

The experiment lasted approximately seven minutes and there were no breaks. After completion of the task subjects were debriefed as to the aims of the experiment and the experimental manipulations used.

Results

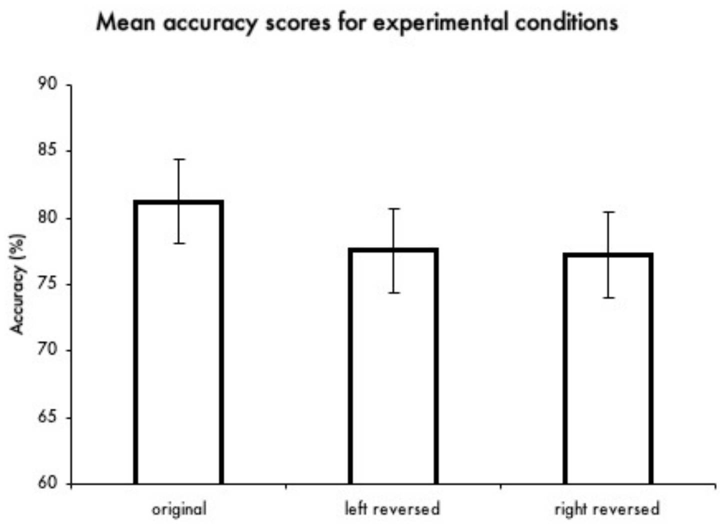

Figure 2.9: NB Error bars denote 95% confidence intervals, which were calculated using the method described by Loftus & Masson (1994).

Though the suggestion from the accuracy data charted above is that mirror reversing one image in a face pair makes the discrimination task more difficult, the main effect of trial type does not reach statistical significance [$F(2,35)=2.03$; $p=0.14$]. In addition, there was no effect of reversal on response criterion [$F(2,35)=0.66$; $p=0.52$] or on sensitivity [$F(2,35)=1.02$; $p=0.37$]. Overall response latency (i.e. "same" and "different" responses combined) was statistically equivalent ($F<1$) for original (M=1758ms, SD=936), left reversed (M=1702ms, SD=762) and right reversed face pairs (M=1707ms, SD=694).

Discussion

It was hypothesised that the configural and featural asymmetry that is unique to an individuals face, if reversed, would lead to the perception of a separate identity. This hypothesis was not supported by the data reported here. Reversal of one image produced no detectable effect on the accuracy or response behaviour of participants on this task, which leads us to believe that the face-matching task is no more difficult when one of the images is reversed. This finding can be explained in one of two ways. Either the asymmetries present in the faces were not sufficient to significantly alter the appearance of the face after lateral reversal, or subjects were able to represent the images in both their veridical and mirror reversed orientation, therefore enabling a discrimination judgement to be made independently of lateral orientation.

It is unlikely that the lack of an effect results from any lack of sensitivity to the differences in relationships between features or even in the direction of the asymmetry present in the features themselves. If this were the case then a difference between the corresponding hemiface matching and opposite hemiface matching tasks in the previous experiment would not have been observed. If the half-faces were perceptually equivalent there would have been no decrement in performance in the opposite hemiface condition of experiment 1. That the two hemispheres of a face are perceptually distinct has already been established.

Another explanation is that the face-matching task can be achieved by somehow performing the opposite transformation on one of the images prior to a match decision. If this is the case however, the reverse transformation process is automatic and does not involve conscious deliberation. The response latencies across experimental conditions were equivalent so there is no evidence that any extra processing is required when one of the images has been reversed. The notion that two mirror representations are available to perceptual systems is an appealing one, and it has been proposed that the visual system is structurally disposed to represent both the veridical and laterally reversed representations of a perceived object (Corballis & Beale, 1976; Noble, 1966, 1968). In addition, it has been argued that viewpoint generalisation in face processing may be achieved by mirror reversing previously encountered views (Troje & Bülthoff, 1998).

There is an additional, more prosaic explanation of our findings in this experiment. It could be that performance was unaffected due to the salience of external features in this task, and given the previous literature indicating that this information is important in tasks involving unfamiliar faces (Bonner et al. 2000), this remains a strong possibility. If a match decision was made based on the hairstyle (or hairline) of the person for example, it may be that this information is more readily reversible than the internal configurational information. As was discussed in experiment one however, the hairstyles of the policemen pictured in our stimuli are highly homogeneous, and this cue would commonly mislead participants if used as a discrimination criteria.

Experiment 3: Discriminating between mirror images of unfamiliar faces

Introduction

It has been previously contended that the task of matching two previously unfamiliar faces is achieved through a rather unsophisticated "image matching" strategy (Bruce et al. 1999; Hancock et al. 2000; Megreya & Burton, 2006). In their 2006 paper, Megreya and Burton demonstrated that the ability to match two unfamiliar faces was highly correlated with performance on the same task when the test images were turned upside down. As 'inverting' face stimuli is thought to disable processes that have adapted to allow facial recognition, this is taken as evidence that unfamiliar faces are processed in a qualitatively different way to familiar faces. Furthermore, significant positive correlations between unfamiliar face matching and object matching tasks reported in this paper suggest that the process involved in matching unfamiliar faces is largely equivalent to an 'image matching' strategy. Importantly however, performance on the inverted task was markedly inferior to performance on upright face matching. This suggests that the process involved is not exclusively reliant on 'image matching', or at least that this process is in some way facilitated by existing face processing systems.

The present experiment was designed to test whether performance on a 'mirror discrimination' task would correlate with unfamiliar face matching ability. This task involves deciding whether or not one of two otherwise identical images has been mirror reversed and as such requires participants to make fine-grain image-level comparisons between two adjacent images. It is considered therefore that mirror discrimination provides a good analogy for the unfamiliar face-matching task. The definitive feature of human performance on this task is its reliance on image matching strategies (Megreya & Burton, 2006) and these same strategies are important for mirror discrimination. As faces are broadly symmetrical, lateral orientation of the face stimuli used in this experiment can only be disambiguated by subtle image level characteristics. It is therefore hypothesised that performance on these two task should be highly correlated.

Mirror Discrimination

Though the task of 'mirror discriminating' has been chosen for use in the present experiment because of its reliance on discerning differences between two very similar images, this same task has been used to address a quite separate issue in previous research. It has been noted that during normal development the visual system can often confuse stimuli which are the mirror rotations of one another, and this phenomenon has been most commonly reported during the acquisition of reading and writing skills, where characters such as 'p' and 'q' and syllables such as 'on' and 'no' are confused (Mach, 1914).

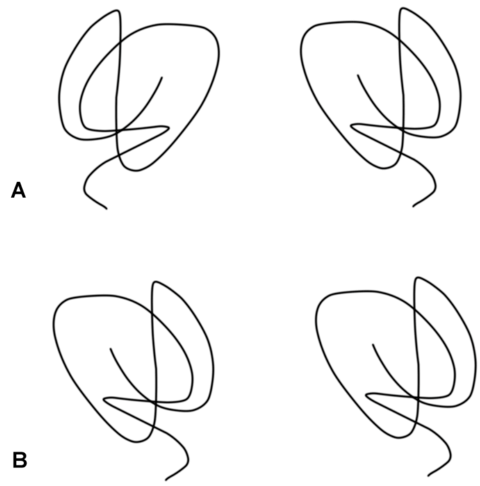

Figure 2.10: Example of mirror discrimination stimuli similar to those used by Warrington & Davidoff (2000). Both 'different' (A) and 'same' (B) pairs are shown

The mirror discrimination procedure used here is borrowed from existing literature investigating a neuropsychological phenomenon. Whilst difficulty in distinguishing mirror reversals of letters and objects is mostly overcome during the course of cognitive development, in particular cases of brain disease or damage individuals have been shown to have great difficult in perceiving any difference between two images that are mirror reversals of one another (Tarr & Bülthoff, 1995; Tumbull et al, 1995; Davidoff & Warrington, 1999; Warrington & Davidoff, 2000). Examples of the stimuli used in these experiments are illustrated in figure 2.10.

Interestingly, in each of the reports cited above the deficits apparent in participants' ability to discriminate between mirror images occur in the context of a preserved capacity for object recognition. Furthermore, in one of these studies it is reported that mirror discrimination performance is greatly improved when the subject is unable to recognise the object (Warrington & Davidoff, 2000). This finding, in combination with previous demonstrations of dissociation between mirror discrimination and object recognition in neuropsychological (Davidoff & Warrington, 2001; Priftis et al. 2003; Turnbull & McCarthy, 1996) and behavioural data (Jolicoeur, 1989) provides strong evidence that the demands of the mirror discrimination task can be dissociated from those of recognition. This has previously been demonstrated to be the case when matching unfamiliar faces (i.e. matching uses different processes than recognition, Megreya & Burton, 2006) and therefore provides further reason to expect that performance on these two tasks will be correlated.

Unfamiliar Face Matching

The unfamiliar face-matching task used in this experiment has been described in previous experiments, however the stimuli used here are drawn from a different database.

Method

Mirror Discrimination Task

Stimuli and Materials

Forty images of unfamiliar faces (20 male and 20 female) were taken from an in-house database consisting of high-resolution images taken under diffuse lighting conditions. These images, all of which were taken with the subjects looking straight at the camera, were used to create two types of image pair per identity. Either the image was placed in horizontal alignment alongside an exact duplicate of the image, or it was placed alongside a mirror reversed duplicate of the image (see figure 2.11). These image pair arrays measured 600 pixels by 400 pixels and each face was positioned with their eyes level with the vertical midline and offset by 200 pixels to the left or right of the horizontal midline. For each identity an 'identical' image and a 'reversed' image pair were created giving a total of 40 'identical' and 40 'reversed' image pairs. Though the stimuli in this task were taken from the same database as was used to create the face pairs in the unfamiliar matching task, it was ensured that no identities were used in both tasks.

Figure 2.11: Example trial stimuli from the Mirror Discrimination task (‘different’ pair).

Design and Procedure

Forty-six people were recruited by way of an advert posted on the World Wide Web. Participants ranged from 18 to 66 years (mean age=32.2) and there were 22 males in the sample. All participants had either normal or corrected vision.

Participants were sat in front of a Macintosh workstation running the experimental software Psyscope and were given the task instructions. Subjects were told that they would be shown a series of pairs of identical images and that in some cases these images would be in the same left-right orientation and that in others the two images would be mirror reflections of one another. Their task was to decide whether the images were identical or whether one had been reversed and they made their response by pressing one of two vertically adjacent keys on the keyboard. Face pairs were presented in a random order and responses and response times were recorded after each trial.

Unfamiliar Face Matching Task

Stimuli and Materials

Face pairs were taken from the same in-house database used to create the experimental stimuli in the mirror discrimination task. The stimulus materials were taken from an existing test of unfamiliar face matching ability in which high-resolution images of individuals are either presented alongside Digital Video (DV) captures of the same person (matched pairs) or alongside a DV capture of a different person similar in appearance. The DV and high-resolution images were taken under the same diffuse lighting conditions and the two images were taken ten minutes apart. Similar pairs were created on the basis of similarity data collected by way of a procedure outlined by Bruce *et al* (1999), whereby participants sort the database into piles based on the perceived similarity of the faces. The faces that were most commonly sorted together were paired to make "different" pairs. Same pairs were selected so that for each different pair there was a corresponding same pair (i.e. each target identity appeared in a 'same' and a 'different' trial). In all there were 168 stimulus pairs (84 same/ 84 different).

Figure 2.12: Example trial stimuli from the Unfamiliar Face Matching task ('different' pair).

Design and Procedure

The same forty-six subjects that participated in the mirror discrimination task also completed the unfamiliar face-matching task. Participants were presented with all 168 face-pairs in a randomly ordered series displayed on a Macintosh 15” monitor at a resolution of 1152 X 864 pixels. They were asked to respond as to whether the two images were of the same person or of two different people and the face-pair remained on the screen until the participant made their response. Although subjects were told that accuracy was more important than speed, they were further encouraged not to take too long in deliberating over each decision. The task typically took 10 minutes to complete.

Results

Performance on the mirror discrimination task [mean= 92.8% correct; sd=7.0] was superior to performance on the unfamiliar face matching task [mean= 91.3%; sd=8.0], but this difference was not found to be statistically reliable [t(46)= 0.63; p>0.05].

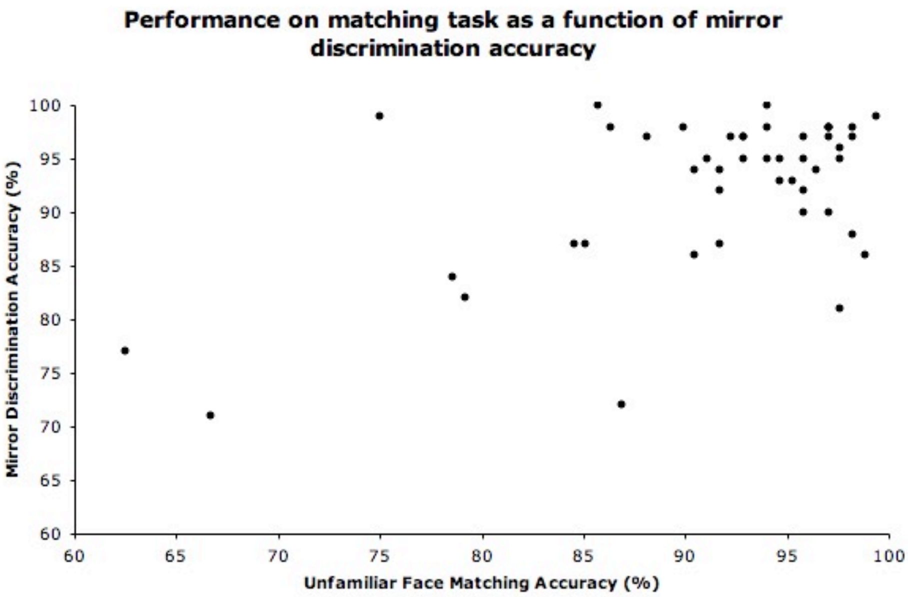

Figure 2.13: Scatter plot of overall accuracy on unfamiliar face matching and mirror discrimination tasks.

Overall performance accuracy in the mirror discrimination task correlated with performance in the unfamiliar face-matching task, and this result was highly significant using a Pearsons-r test of correlation [$r=0.544$; $p<0.001$]. However, it was considered that the outlying participants represented in the scatter plot above contributed to the significance of this correlation. When the analysis was repeated using a Spearmans Ranked Correlation Coefficient the analysis was non-significant [rho= 0.242; $p=0.105$]. It is likely that the high performance on these two tasks has obscured any pattern of correlation that may have existed between performance on these two tasks.

Differences between performance on these two tasks were detected when performance was treated separately according to trial type (see Table 1). Whereas performance on 'same' trials was highly correlated with performance on 'different' trials in the mirror discrimination task [$r=0.743$; $p<0.001$], the correlation between same and different trial performance was not as strong for the unfamiliar face matching task [$r=0.496$; $p<0.001$], and the difference between these correlation coefficients was found to be statistically reliable [$t(43)=3.63$; $p<0.001$]. Correlation coefficients computed using the Spearman-rho test also showed performance on 'same' trials to be correlated with performance on 'different' trials for both unfamiliar face matching [rho= 0.341; $p<0.05$] and mirror discrimination [rho=0.626; $p<0.01$]. These results show that 'hit' and 'false alarm' responses did not have such a strong negative association in the unfamiliar face-matching task as they did in the mirror discrimination task. This difference may provide an important dissociation between performance in these two tasks.

	UFM match	*UFM mismatch*	*MD match*	*MD mismatch*
UFM match	-	0.496***	0.447***	0.579***
UFM mismatch	-	-	0.285	0.526**
MD match	-	-	-	0.743***
MD mismatch	-	-	-	-

Table 2.1: Pearson-r correlation coefficients for match/mismatch trials in the two cognitive tests (UFM= Unfamiliar Face Matching, MD= Mirror Discrimination; * = $p<0.05$, ** = $p < 0.01$, *** = $p < 0.001$)

Discussion

Unfortunately this experiment did not provide a valid test of the hypothesis that unfamiliar face matching relies on an image matching strategy (e.g. Megreya & Burton, 2006). As the mirror discrimination task requires purely image-level comparison of two adjacently presented face images, it was considered to be an example of a visual-matching task involving faces to which our existing face recognition systems can not lend any beneficial support, and correlation between this task and unfamiliar face matching would have provided support for the hypothesis proposed by Megreya and Burton (2006). However, as performance on both these tasks was generally very high in this experiment, the analysis reported did not provide reliable evidence for the hypothesised correlation. The weak (non-significant) correlation detected using the ranked correlation test may suggest that these tasks are related, and that increasing their difficulty might produce a statistically significant result.

Although there is a suggestion in this data that performance on the two tasks was broadly associated, a reliable dissociation was detected when performance was separated into accuracy scores for 'same' and 'different' responses. The negative correlation between hit and false alarm rates on the mirror discrimination task is very high, however the correlation between these two response behaviours is not as high in the unfamiliar face-matching task. A similar finding has previously been reported for unfamiliar face matching performance, with researchers reporting there to be no significant correlation between hits and FA's whatsoever (Megreya & Burton, 2007). Therefore, although the dissociation between performance on 'same' and 'different' trials was not as profound in this case, the correlation between hits and FA's was weak in relation to the r-value for this same relationship in the mirror discrimination task.

The difference between performance behaviour on these two tasks is likely to stem from an increased reliance on face recognition systems when matching unfamiliar faces relative to when participants were performing the mirror discrimination task. It has previously been demonstrated that unfamiliar face matching performance reliably predicts participants performance both on the same task inverted (i.e. the stimuli turned upside down) as well as on other object-matching tasks, but that it does not predict performance in tasks involving familiar faces (Megreya & Burton, 2006). Whilst this may reflect an increased reliance on image-based strategies during unfamiliar face matching, it is not certain from this data that mechanisms employed for recognising familiar faces are disengaged entirely during unfamiliar face processing.

It may be that as a purely image-matching strategy does not provide a reliable technique for deciding if two different images are of the same person (Megreya & Burton, 2006), other processes are additionally employed to enable successful performance. The contribution of additional processes to behaviour during unfamiliar face matching most probably account for the differences in the extent to which false positive response behaviour predicts hit response behaviour in the two tasks reported here. Furthermore, it is likely that these additional processes emanate from the engagement of existing face recognition mechanisms when performing the unfamiliar face-matching task.

Chapter Summary

In experiment 3 I tested participants' ability to discriminate between mirror images of faces, a task that involves close examination of image level stimulus properties. This ability was not found to correlate with unfamiliar face matching, however the experiment did not provide a valid test of the hypothesis that performance on this task is reliant upon a strategy most akin to mage-matching (Megreya & Buron, 2006). The high level of performance on these two tasks made it difficult to draw conclusions regarding the relationship between them and therefore this experiment contributes little to our understanding of the processes involved in unfamiliar face matching. The poor performance for matching opposite hemifaces reported in experiment one does however lend support to the notion that unfamiliar face matching is commonly achieved by matching surface information in two images. An image-matching strategy would be a particularly unreliable way in which to perform this task, as the visual information in one hemiface relates in no exact manner to its opposite half.

Thus, the image-matching hypothesis would predict the poor performance reported here.

The findings of experiments one and two suggest that a given face is perceptually equivalent to its mirror reversal (experiment two), even though the two sides of the face are perceptually distinct (experiment one). This slightly paradoxical finding suggests that when presented in full, the information present in a face can be readily reversed (in lateral orientation), but that the visual information in one half of a face does not reliably predict the visual information in the other half (or at least we find it difficult to project this information across the vertical meridian).

An intriguing question is whether the effect reported in experiment one is a result of the incongruence between the face halves, or whether it is due to a cognitive difficulty in reflecting the information across the vertical meridian. It may be that the information each side of this meridian in a given face is sufficiently similar to elicit a 'same' response in most cases, but that an inaccuracy in reflecting the image leads to a reduced correspondence between the face-halves and therefore to an increase in 'different' responses. However, if this is the case then it is puzzling that no effect of reversal was detected in experiment two. Presumably a cognitive 'mirror reversal' should be equally achievable for whole faces and hemi faces, unless this process relies on a whole percept being available to the senses. However, the fact that performance on the opposite hemiface condition was superior when half-faces were presented in their natural configuration (i.e. left hemiface on left and right on right) may be evidence that mental reversal is more difficult in the hemiface condition. This effect was not statistically reliable however, and therefore this argument would need to be subject to further empirical testing before being accepted.

The most surprising result in the present chapter is that there was no effect of mirror reversal on the unfamiliar face matching procedure. Given that we discriminate between individual faces based on very small differences in facial configuration, and that we are extremely sensitive to deviations from symmetry in faces (Simmons et al. 2004), one would expect reversal of this information to produce a perceptually distinct face. Furthermore, if unfamiliar face matching is reliant on an 'image-matching' strategy as is suggested by previous research (e..g Megreya & Burton,

2006) as well as the data reported in experiment three, it would be expected that the effect of mirror reversing an image would be greater as this manipulation pronounces surface differences. However, in experiment two subjects are behaving as though an image of an unfamiliar face is perceptually equivalent to the mirror representation of that face. Our data suggests that a normalisation process, which involves mirror reversing one of the images, occurs prior to a match/mismatch decision being made and that this pre-processing is relatively automatic. In order for successful matching performance to be maintained despite incongruous lateral orientation however, the linear reversal occurring in our perceptual systems must be precise.

Such precision does not characterise performance on tasks where images have been reversed relative to the x-axis. The ‘inversion effect’ is traditionally thought to occur due to the disruption this manipulation causes towards the processes responsible for configural processing. Whereas traditionally this reliance on an upright orientation has been thought to be an artefact of visual expertise, more recent research has instead proposed that it reflects a difficulty in normalising the incoming visual input (Collishaw & Hole, 2002). Data showing a linear relationship between the angle of rotation from upright and performance in familiar face recognition tasks is taken as evidence to support this hypothesis (Collishaw & Hole, 2002; Parr & Heintz, 2006; Stevenage & Oborne, 2006; Valentine & Bruce, 1988). The implication of these studies is that during the process of normalisation (i.e. mentally rotating the image to an upright orientation) some veridical aspects of the image are altered, leading to reduced recognition and matching performance.

The results of experiment two indicate that the normalisation of mirror-reversed images is in no way as cumbersome as the rotation process described above. As there is no reliable effect of mirror reversal on matching accuracy or response latency, any processing that represents an image of a face in its reverse lateral orientation would appear to do so in a precise and automatic manner.

Chapter 3

Is face recognition impaired by mirror reversal?

Introduction

We encounter images that have been reversed in left-right orientation during our everyday lives. When we look at photographs of ourselves we are seeing them in the reverse orientation relative to our normal experience, as we would typically inspect our own face whilst looking into a mirror. In addition, we encounter images of other people in a reversed orientation. Take for example the situation shown in Figure 3.1 where an image of a famous celebrity has been sourced from the World Wide Web using the Google Images search engine. On the Google frame at the top a thumbnail of the original image is shown, and below that the image is shown in its published form. It is apparent that the original image has been mirror reversed so as to conform to the composition of the KINERASE website. This technique is commonplace in the world of publishing, and academic journals are not immune from such practise (see Figure 3.2).

Figure 3.1: Google Image search result for "Courtney Cox" showing lateral image reversal

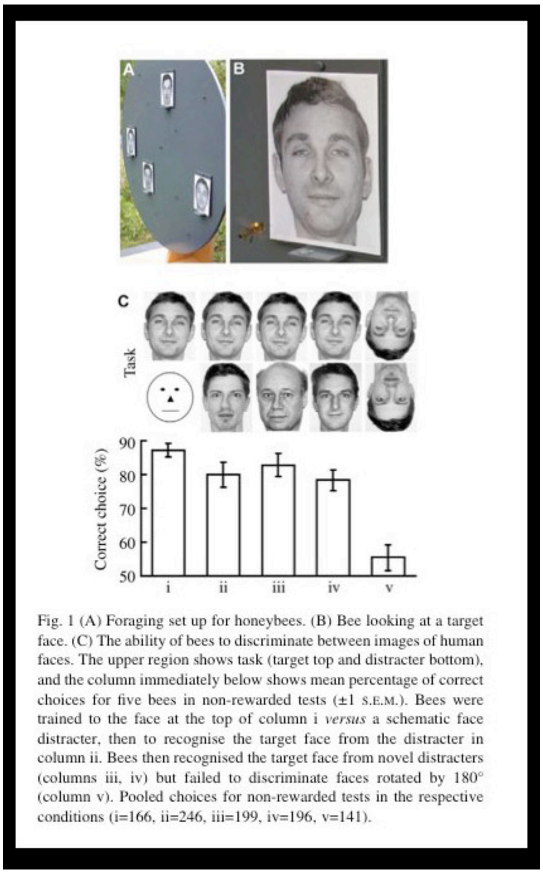

Figure 3.2: A figure taken from the journal *The Journal of Experimental Biology* (Dyer, Neumeyer & Chittka, 2005) where the original image (C, top row) has been subjected to reversal prior to publication (B).

Despite the prevalence of mirror images in our everyday lives, we do not appear to be affected by these transformations. In fact, we are generally unaware of the fact that this occurs, and would not appear to have trouble in recognising people despite mirror-transformation. In addition, we have no trouble in recognising images of ourselves in the opposite orientation to which we would normally experience our own face (though this perception can feel peculiar to subjects, see Brédart 2003). The results of the previous chapter suggest that a form of object constancy that allows us to perceive a mirror image as equivalent to the face in its veridical orientation mediates our perception of faces, and this result could have some important implications for the study of face recognition.

Historically, the importance of configural processing in familiar face recognition has been emphasised (e.g. Haig, 1984; Friere et al. 2000; Leder & Bruce, 1998). For this reason it is surprising that the mirror reversal manipulation in experiment 2 did not significantly alter participants' response behaviour. After all, the direction of the metric relations between the component features has been reversed in one of the test stimuli. It has previously been argued that unfamiliar face matching is made more difficult when the two images presented are taken using different cameras (Bruce et al. 1999). This effect is by virtue of the small distortions in aspect ratio that the different lenses produce. If unfamiliar face matching is disrupted by these small perturbations then one might reasonably expect to detect a greater effect of mirror reversal, which alters the configural relationships more severely. It should be noted however that whilst the relationships are altered by reversal, none of the original shape information has been lost subsequent to this transformation: it is only the orientation that has changed (see Rock, 1973).

The process of face recognition is thought to rely on different processes than basic-level object recognition (e.g. Duchaine & Nakayama, 2006; McKone, Kanwisher & Duchaine, 2007), and the 'second-order' relations between facial features are more important in face recognition than in object recognition (Robbins & McKone, 2007). Given that reversal changes the second-order relations between the facial features, one might expect that this would disrupt the recognition of a face. However, this assumes that our memory representations for faces are coded in a manner that preserves both the subtle asymmetries and the direction of these asymmetries (i.e. that the representation specifies the veridical left-right orientation). The experiments reported here each attempt to detect an effect of mirror reversal on face recognition using paradigms previously used to investigate familiar face processing. If differences in behavioural response between veridical and reversed orientations are detected, then it can be concluded that memories for faces retain the left-right orientation of their external analogues.

There have been a number of previous reports of cognitive sensitivity to the mirror reversal of familiar faces (e.g. Mita et al. 1977; Rhodes, 1986; Brédart, 2003). As yet however, it has not been convincingly demonstrated that this stimulus manipulation disrupts face processing. McKelvie (1983) found that memory for images of faces was worse when test images were presented in the reverse left-right orientation relative to when images were presented in their original orientation. Although this result demonstrates that the orientation of *image memory* is disrupted by mirror reversal it does not necessarily follow that stored representations of the faces necessarily specify lateral orientation. Certainly, McKelvie's result implies that this effect (previously reported with images of common objects by Madigan & Rouse, 1974) is robust enough to occur when using highly symmetric stimuli, however it does not necessarily reflect disrupted face processing. Dyer and colleagues (2005) have demonstrated that bees are capable of remembering images of faces. This result does not imply that a nervous system as small as this is capable of recognising faces *per se.* More likely is that memory for specific images of faces can be dissociated from memory for faces. Indeed, there exists evidence to suggest that face recognition proceeds normally when presented with a mirror reversed image of a familiar face. Brooks *et al* (2002) demonstrated using a two-alternative forced-choice procedure (2AFC) that mirror reversal does not significantly increase response latency to name verification responses and that identity priming is invariant to the left-right orientation of familiar faces.

Experiment 4: The effect of Mirror Reversal on Face Recognition

Introduction

Brooks *et al* (2002) reported that laterally reversing an image of a face at the test phase of a priming procedure resulted in statistically equivalent response latencies to those corresponding to face images in their original orientation. Taking this to confirm their hypothesis that the mental representation of faces is invariant to mirror reflection, they sought no further confirmation of their finding. In the present experiment a 2AFC name-verification paradigm was used to verify this conclusion. This methodology has been used to inspect the nature of mental representations in previous research (e.g. Burton et al., 2005).

This same task was used in the priming phase and test phase by Brooks *et al* (2002), yet they failed to find a difference between the original and reversed conditions in either test. Given that previous studies have shown that people are sensitive to mirror reversal of faces (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003), and that familiar face recognition relies on processing subtle relations between features that are altered by this manipulation (see Maurer et al, 2002), a replication of this result was considered necessary. If there is a difference in response behaviour between presentations of familiar faces in their original and mirror orientation then the null hypothesis can be rejected (i.e. that the mental representation of faces is invariant to mirror reflection).

Method

Stimuli and Materials

Twenty-five famous faces were sourced from the Internet. Images were chosen that contained writing in either the foreground or the background so as to ensure that pictures had not been mirror-reversed prior to publication on the world-wide-web. Face images were rotated so as to bring the pupils into alignment with the horizontal plane, cropped around the outline of the head, converted to greyscale and resized to 380 pixels X 570 pixels. For each image I flipped the image horizontally using Adobe Photoshop to create a mirror-reversed copy. The face images measured 8cm by 11cm, pertaining a visual angle of 7.6 by 10.5 degrees (at 60cm), and were centred 13cm apart (visual angle= 6.2 degrees).

Method and Procedure

26 subjects from the undergraduate and postgraduate populations at the University of Glasgow agreed to participate in the study. The sample consisted of 17 female and 11 male subjects ranging from 17 to 29 years in age (mean 22). Subjects were sat in front of a Macintosh workstation running the software Psyscope and the experimenter explained their task to them. Subjects were told that in each trial they would be shown a famous name followed by a famous face and that they were to respond using the keyboard in front of them as quickly and as accurately as possible.

Each trial consisted of a fixation cross in the centre of the screen for 1000ms, followed by a famous name for 1500ms and then by the presentation of a face for 200ms. Subjects' responses and response latencies were recorded. Each of the 25 stimuli appeared twice in each of the two conditions (one positive/matched and one negative/mismatched trial per condition) giving a total of 100 trials. Order of stimulus presentation was fully randomised. The experiment was approximately 7 minutes in duration.

Results

Two subjects were extracted prior to analysis on account of their performance falling below two standard deviations of the group mean.

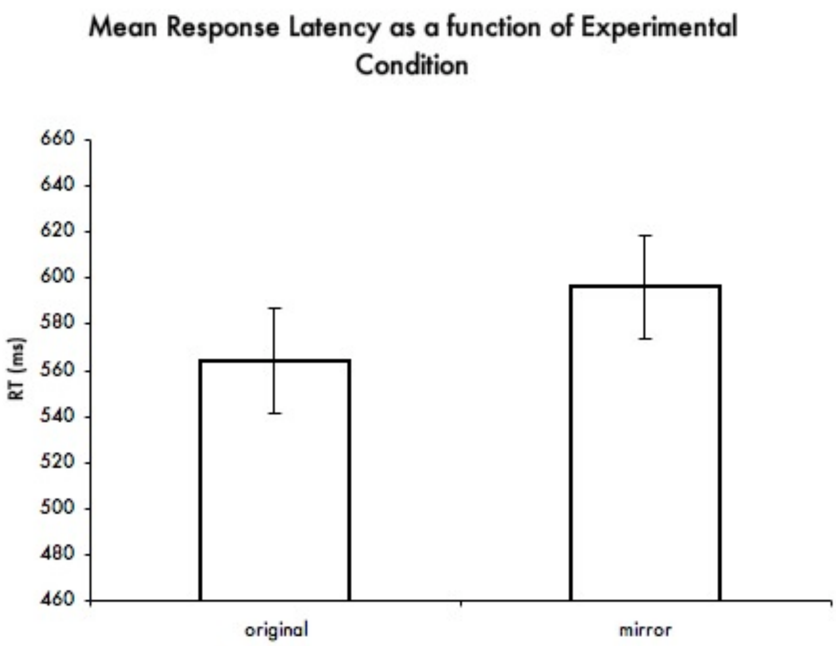

Figure 3.3: NB Error bars denote 95% confidence intervals.

A related sample t-test shows the difference between mean response latencies (for positive trials) for the two experimental conditions [original= 564ms (sd= 70), reversed= 596ms (sd= 79)] to be reliable [$t(1,23)=2.896$; $p<0.01$]. There were no significant differences between error rates for the two experimental conditions [$t(1,23)= 0.496$; $p=0.625$].

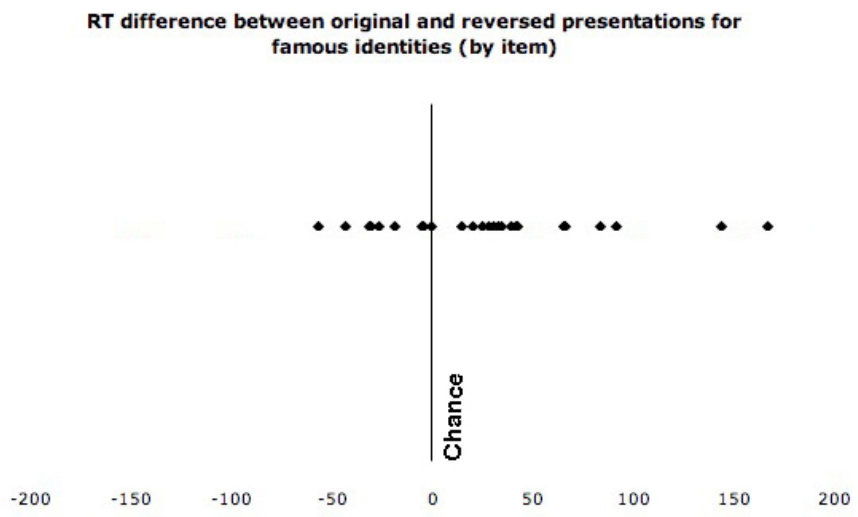

Figure 3.4: Difference scores were calculated by subtracting average response latency to the veridically oriented version of the test stimulus from average RT to the mirror reversed copy.

An additional by-item analysis was executed on the data by firstly calculating the difference in average response latency between mirror-reversed images and non-reversed images (positive values signal faster response to the veridical orientation). These differences were then compared to difference scores as predicted by chance performance (i.e. with a mean of zero) and were found to differ significantly from chance [$t(24)=2.57$; $p<0.05$].

Discussion

The results reported above are in direct contradiction to Brooks *et al* (2002), as they suggest that mental representations of familiar faces contain information relating to their veridical left-right orientation. This is most likely a result of the different methodology used to investigate the question of reflectional invariance in the present study. The name verification paradigm used in their study was slightly different to the one used here as the name preceded the face in the present experiment, whereas subjects in their study responded to the presentation of a name that had been preceded by a face. Nevertheless it is difficult to explain why they found no difference between identical image and reflectional change conditions. It may be that recalling the visual representation from memory is more effectively achieved using the method reported in the present experiment. Regardless of these slight methodological concerns however, the effect of reversal is clear and demonstrates cognitive sensitivity to mirror reversal.

Further investigation may also seek to investigate the process underlying the delay in reaction time elicited by mirror reversal in the present study. It may be that the mind is in some way applying the reverse transformation to the reversed image prior to recognition: this possibility was previously suggested by Hole *et al* (2002) as a way of explaining subjects invariant performance to globally transformed face images. However it could also be that the increased response latency is due to a process of transformation that aligns our stored representation to the incoming (reversed) stimulus. Given that experiment 2 in the present thesis showed no effect of reversal on an unfamiliar face-matching task, it is likely that the normalisation is not dependent on the existence of a stored representation, but rather this ‘mirror’ transformation is made independently of such constructs.

A further explanation is that this effect is due to the novel nature of the reversed image, which captures attention mechanisms, thus delaying response. It is therefore not certain that the increased response latency is reflective of additional cognitive processing. Regardless of the explanation chosen however, it is clear that the cognitive system is responding selectively to mirror reversed images and therefore that lateral orientation is specified at some level of cognitive representation. This result is therefore in keeping with previous studies demonstrating a sensitivity to mirror reversal (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003).

Experiment 5: The effect of Mirror Reversal on Recognition of Half-Faces

Introduction

An experiment was designed to extend the findings reported in experiment one. I sought to determine whether the same decrement in reaction time, found to be a result of mirror reversal, would occur when only half a face was shown (i.e. when bilaterally distributed configural information was not displayed).

Method

Stimuli and Materials

Forty famous faces were sourced from the Internet. Images were chosen that contained writing in either the foreground or the background so as to ensure that pictures had not been mirror-reversed prior to publication on the world-wide-web. The images used in this experiment had not been used in any previous experiments, though some of the identities had been. Face images were rotated so as to bring the pupils into alignment with the horizontal plane, cropped around the outline of the head, converted to greyscale and resized to 380 pixels X 570 pixels. For each image a mirror-reversed copy was created by flipping the image horizontally using Adobe Photoshop.

For each subset of images (original/reversed) a further two subsets containing forty images each were created. One subset was created by erasing the pixel information to the left of the axis of symmetry (right half-face), and another by erasing the pixel information to the right of the axis of symmetry (left half-face). The full-face images measured 8cm by 11cm, pertaining a visual angle of 7.6 by 10.5 degrees (at 60cm), and were centred 13cm apart (visual angle= 6.2 degrees). Example stimuli are shown in figure 3.5.

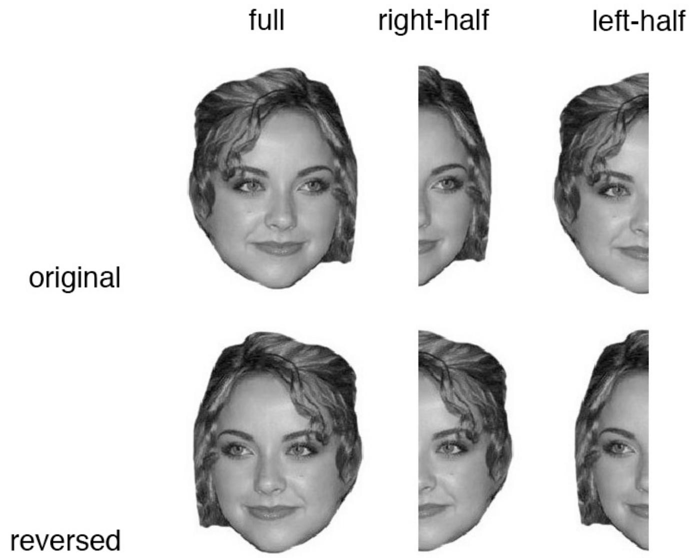

Figure 3.5: Example stimuli from experiment 5.

Method and Procedure

38 subjects from the undergraduate and postgraduate populations at the University of Glasgow agreed to participate in the study. The sample consisted of 21 female and 17 male subjects ranging from 18 to 29 years in age (mean 21). Subjects were sat in front of a Macintosh workstation running the software Psyscope and the experimenter explained their task to them. Subjects were told that in each trial they would be shown a famous name followed by a famous face and that they were to respond using the keyboard in front of them as quickly and as accurately as possible. They were told that sometimes only half a face would appear but they were to not let this distract them, but instead they must concentrate on making their response as early as possible.

Each trial consisted of a fixation cross in the centre of the screen for 1000ms, followed by a famous name for 1500ms and then by the presentation of a face (or half-face) for 200ms. Subjects' responses and response latencies were recorded. Each of the 40 stimuli appeared twice in each of the six conditions (one positive and one negative trial per condition) giving a total of 480 trials. The experiment was approximately 20 minutes in duration.

Results

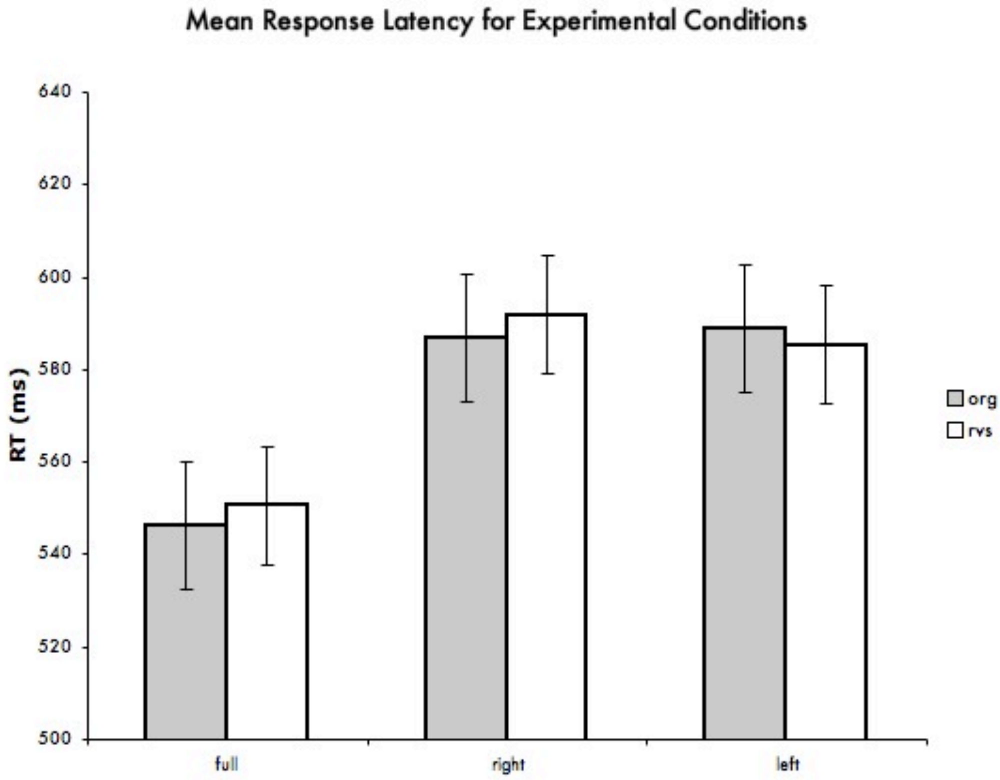

Figure 3.6: Mean RT's for positive trials (correct response). NB Error bars denote standard error.

	FULL FACE	LEFT HALF	RIGHT HALF
ORIGINAL	87.7 (SD=14.7)	86.6 (SD=14.7)	85.5 (SD=14.5)
REVERSED	87.7 (SD=15.0)	86.0 (SD=15.1)	86.5 (SD=14.8)

Table 3.1: Mean accuracy data (percentage correct) for name verification task

Data from this experiment was subjected to a 2 (original/reversed) X 3 (full/right half/left half) within-subjects ANOVA. There was found to be a significant effect of presentation type [F(2,37)=14.462; p<0.01] with full-face presentation resulting in significantly faster reaction times than both half-face conditions. There was no significant main effect of orientation and no interaction. Furthermore, there were no significant effects of the experimental manipulations on participants' response accuracy (see Table 1).

Discussion

The only significant effect to be reported in this experiment was the significantly faster reaction times to whole face relative to hemi-face presentations. That participants recognised faces more quickly when the full face was shown is hardly surprising given that half of the visual information is unavailable in the hemiface conditions. What is more interesting however is that there was no effect of deleting half the available visual information on participants' response accuracy. Given that the configural information that is considered important in face recognition is distributed bilaterally, one might reasonably expect that obscuring one half of the face might impair participants' ability to recognise the person. This does not appear to be the case however. It would appear that participants are able to recognise faces on the basis of a reduced amount of information, and this ability has been previously demonstrated (e.g. Gosselin & Schyns, 2001).

The difference in response latency between hemi-face and full-face conditions is notably slight (in the region of 25ms), and this again suggests that the bilaterally distributed configural information is not important in familiar face recognition. Although slight, this difference in response latency may conceal some additional processing that is necessary when identifying a hemi-face. Further investigation could ascertain whether this latency is reflective of some perceptual 'filling-in' prior to successful identification, or whether the delay in response is reflective of disruption to face detection processes. Face detection is reliant on the 'first-order' relational information in a face (Maurer et al. 2002), and this information is disrupted by the manipulation in this experiment.

The most striking aspect of the data reported above is that I failed to replicate the effect of mirror-reversal reported in the previous experiment. This may suggest that this earlier finding was a statistical anomaly, however this is not certain. Another likely explanation is that effect was washed out in the present experiment as a consequence of practise effects. The stimuli were each presented twice in each condition and therefore each image appeared 12 times in one form or another. Therefore participants may have been using memory of the experimental stimuli to make their response as opposed to recalling their stored representation from memory, as the experimental design requires.

Experiment 6: Does degree of asymmetry affect the degree of disruption when recognising a mirror-reversed face?

Introduction

The previous two experiments show that there is a slight effect of mirror reversal on response latency in a name verification task. It is uncertain whether this effect is reflective of some cognitive normalisation or whether attentional mechanisms are responsible for the delayed response. In addition, it is currently uncertain whether or not this effect is indeed reliable. The failure to replicate the effect observed in experiment four may suggest that this result was merely a statistical anomaly. The experiment reported here is an attempt to resolve these two concerns.

Primarily this experiment is designed as a replication of experiment four. However, a different set of stimuli is used here, and it is anticipated that these will provide a more direct test of the main hypothesis. One limitation of the previous two experiments was that superficial aspects of the images might have contributed to the effect of reversal. For instance, asymmetry in hairstyle may have contributed to the effect and this would weaken the argument that delayed response is reflective of a retention of the left-right configuration of the face itself, as subjects' delayed response could merely reflect the unusual direction of the hairstyle. In addition, the direction of lighting may have had an effect as it has been previously demonstrated that subjects respond optimally to stimuli that are lit from the left hand side (Sun & Perona, 1998;

McManus et al 2004). The stimuli used in the present experiment are unaffected by these transient image qualities by virtue of the averaging process from which they are derived and therefore will enable a more rigorous test of the hypothesis.

The stimuli used in the present experiment were created for use in a previous study (Burton et al 2005). These images were created by averaging pixel information across a set of twenty photographs of a given celebrity. In all there were fifty famous faces used and twenty images of each were sourced from the World Wide Web. Images were rotated so as to align the pupils to the horizontal, cropped so as to frame the head snugly, resized to 190 X 285 pixels, converted to greyscale and saved in BITMAP format. All images were then morphed to a standard shape using an in-house program based on bi-polar interpolation and these standardised images were then averaged together. An average shape was then calculated by subjecting the 20 files containing shape information and the average image files were morphed back to this average shape to produce the final average (see figure 3.7).

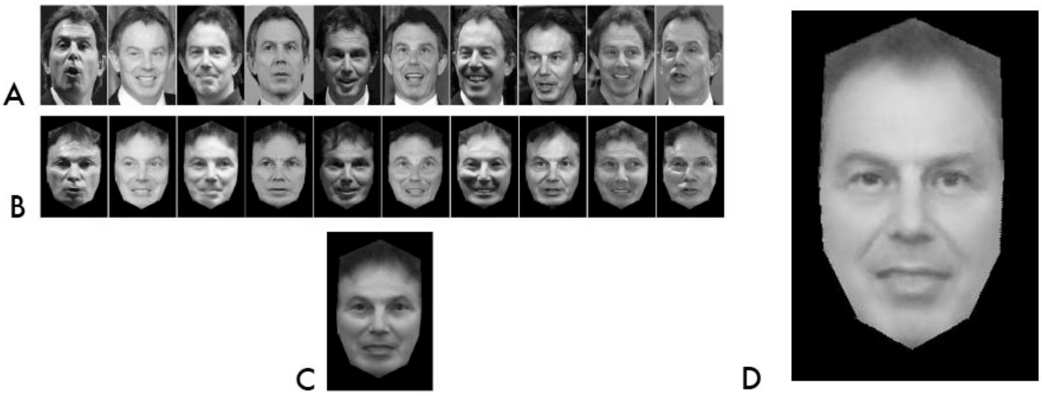

Figure 3.7: An illustration of the image averaging process. Original images (A) were set to a standard shape configuration (B) and then averaged to produce the 'shape-free' texture average (C), this image was then morphed to the average shape to give the final stimulus (D).

It was considered that the use of average faces would enable more accurate symmetry ratings as these representations are unaffected by transient visual characteristics that are prevalent in typical photographs (e.g. lighting direction, expression, head angle and asymmetry in hairstyle). These artefacts would contribute to image-level asymmetries that may bias subjects to make their judgements based on this superficial information instead of the structural symmetry of the face itself.

As the images used in this experiment do not suffer from ambient asymmetries in the way that standard (unconstrained) images do, it enabled symmetry ratings to be collected for the celebrity faces. There are a number of factors that contribute to the asymmetry of a given image (e.g. head angle, expression, lighting direction), however these factors are ‘washed-out’ in the average process described below. For this reason it was anticipated that the average images would allow for more reliable symmetry rating data to be collected than would have been attainable using single images of the celebrities. This symmetry data will be used to assess whether the effect of mirror reversal reported in experiment four is modulated by the degree of symmetry in a face.

There are a number of reasons why the degree of symmetry may modulate this effect. Firstly, if the effect of mirror reversal were due to the increased novelty of the stimulus then one would expect this effect to increase proportionately with the degree of *asymmetry* in the face. The greater the degree of asymmetry, the more different an image will be to its mirror image and therefore the more ‘novel’ it should appear. The manner in which asymmetry would affect a process of cognitive normalisation is less certain however. It could be that increasing asymmetry would actually make this process easier as the left-right orientation is more explicitly apparent in a face that is highly asymmetric. That is, if recognition relies on initially distinguishing the left side of a face from the right, then this would be more readily achieved in a highly asymmetric face, where the direction of asymmetry is less ambiguous.

Method

Symmetry Ratings

All fifty face averages were presented on two sheets of A4 paper, with space below each image for participants to write their symmetry score for each face. Images measured 2.5 cm by 3.6 cm. Participants marked their score on the paper using a pen.

Sixteen undergraduate students [8 male, 8 Female; mean age= 20.6 (sd=1.4)] from the University of Glasgow each rated all fifty experimental stimuli for degree of symmetry on a scale of 1 to 7 where 1 was defined as a very asymmetrical face and where 7 was perfectly symmetrical. Participants were encouraged to use as much of the scale as possible and to inspect the images closely before rating each image. Ratings were subsequently collated and an average rating for each face was obtained.

Symmetry Measurements

Method and Procedure

In addition to the ratings of symmetry it was decided that a physical measure of asymmetry should also be calculated. This was calculated using the four coordinates shown below (figure 3.8), which were taken from the average shape files generated as part of the procedure used to create the average faces.

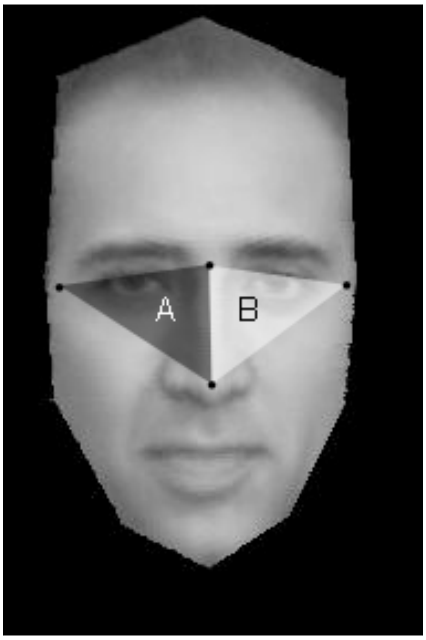

Figure 3.8: An illustration of the method used to derive the physical asymmetry measure reported in experiment 5. Area ‘A’ was subtracted from area ‘B’ and this difference was used to calculate an absolute value of asymmetry.

Name-Verification Task

Stimuli and Materials

The sixteen averages with the highest symmetry rating and the sixteen with the lowest ratings were selected for use in the experiment.

Method and Procedure

Subjects sat in front of a Macintosh workstation running the software Psyscope and the experimenter explained their task to them. Subjects were told that in each trial they would be shown a famous name followed by a famous face and that they were to respond using the keyboard in front of them as quickly and as accurately as possible.

Each trial consisted of a fixation cross in the centre of the screen for 1000ms, followed by a famous name for 1500ms and then by the presentation of a face for 200ms. Subjects' responses and response latencies were recorded. Each of the 32 stimuli appeared twice in each of the two conditions (one positive and one negative trial per condition) giving a total of 128 trials. Order of stimulus presentation was fully randomised. The experiment lasted for approximately 10 minutes.

Results

Symmetry Ratings

The average symmetry rating was 4.57 (sd=0.65) and the sample had a median symmetry rating of 4.63. The sixteen celebrities with highest mean symmetry ratings were assigned to the 'high symmetry' condition [mean rating= 5.26, sd=0.17] and the sixteen with the lowest rating below this were assigned to the 'low symmetry' condition [mean rating= 3.80, sd=0.45]. An independent samples t-test confirmed that the degree of perceived symmetry in these two stimulus groups differed significantly [$t(30)= 2.14$; $p<0.05$].

2AFC Name Verification

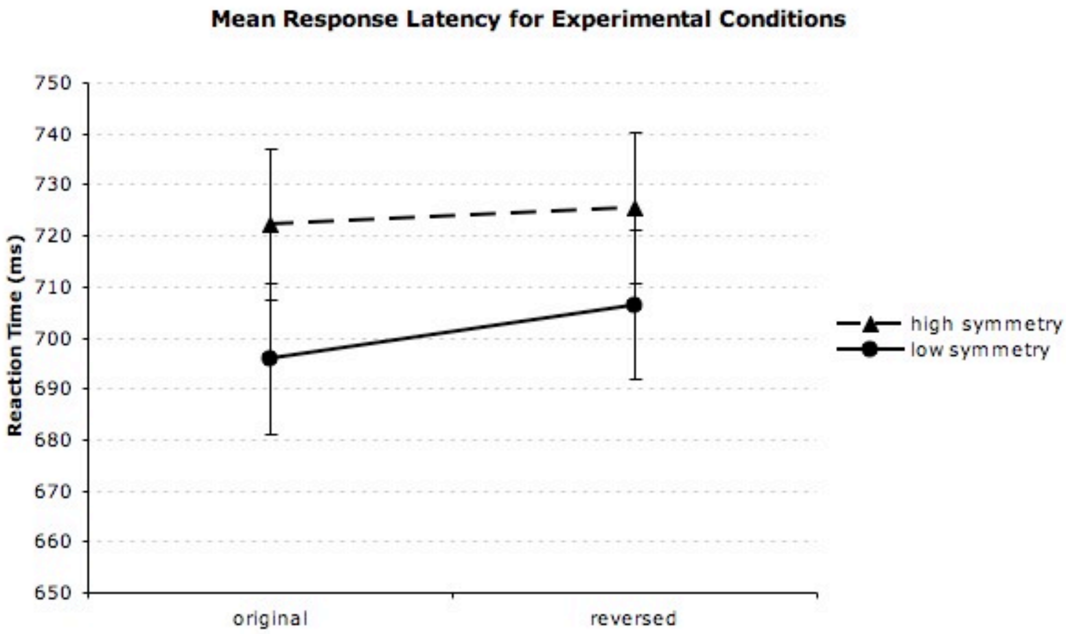

Figure 3.9: Mean RT data for experimental conditions (error bars denote 95% confidence intervals calculated using the method described by Loftus & Masson, 1994).

A within-subject two-factor ANOVA was performed on reaction time data. This analysis revealed no reliable main effect of symmetry [F(1,31)=2.84.; p=0.10] and no effect of reversal [F(1,31)=0.413; p=0.53]. In addition, the interaction between these experimental factors was not statistically reliable [F(1,31)=0.26; p=0.62].

Symmetry Measurements

The direction of asymmetry was not considered to be important for the analysis and so asymmetry measurements were converted to absolute values. Additionally, for each stimulus I calculated the average difference (across subjects) between response latency when the stimulus was presented in its original orientation compared to when it was presented in its reverse orientation.

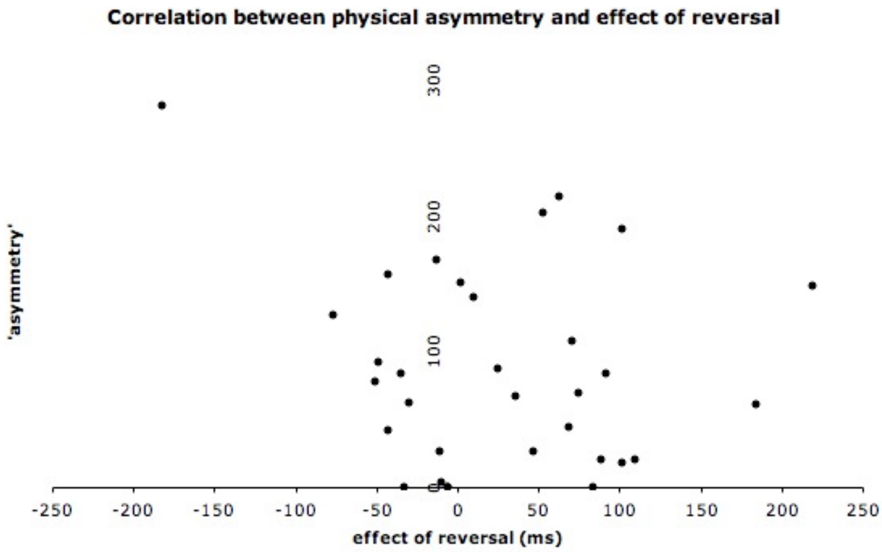

Figure 3.10: **NB** Higher values on the y-axis indicate higher levels of physical asymmetry, and positive values on the x-axis indicate that the original stimulus was responded to more quickly than the mirror-reversed copy.

The correlation between physical asymmetry and the effect of reversal was non-significant (r= -0.19; p>0.05), and this was also true when the differences between mean response latencies in original and reverse orientations were converted to absolute values (r=0.28; p>0.05). The mean difference between original and reversed presentations across items was 26.4 ms, and although this result suggests that reversal might on average decrease reaction time, this difference did not differ significantly from chance (0) when tested using a two-tailed hypothesis test [t(31)=1.89; p=0.068].

Discussion

The null effects reported here represent a second failure to replicate the significant effect reported in experiment three. Again there was found to be no reliable difference in either response latency or response accuracy when recognising faces either in their original or reversed lateral orientation. Although this weakens the certainty at which it can be asserted that the cognitive system is sensitive to mirror reversal, our data do not rule out this possibility. In fact, in the previous three experiments I found a consistent pattern in response latencies, with mirror reversed images being responded to slower than non-reversed images in each experiment. In experiment four this difference (32 ms) was significant and in experiment five and six this difference was not significant (3ms and 9ms respectively), however in each instance mean response latency for non-reversed images was shorter than for reversed images.

It is possible that in the present experiment the effect of reversal was attenuated by the nature of the stimuli. The process of making the average images involved collecting twenty images of each celebrity in the corpus and crucially there was no attempt made to ensure that these images were in their original lateral orientation. Therefore the averages may have included a number of images in their reverse orientation, which would have served to 'wash-out' some of the asymmetry in the faces.

There is no evidence in the data reported here that the symmetry of a face modulates the effect of reversal. However, as no main effect of reversal was observed in the present experiment, this question remains largely unresolved. Using a larger stimulus set may help to resolve this question and such an investigation would improve our understanding of the mechanisms underlying the invariance displayed by recognition systems to mirror reversal. However, as the effect of reversal on response times is evidently slight and has proved difficult to detect reliably thus far, it is anticipated that an extremely powerful design would be required to detect an interaction between reversal and degree of symmetry.

It is very difficult to collect a sufficient number of famous face images that are unambiguous with respect to their mirror orientation and are additionally symmetric with regards to head pose, eye gaze and direction of lighting. Therefore, in order investigate this question properly one would need to capture images of faces under controlled conditions. In addition, symmetry ratings should be replaced by a more objective measure of symmetry based on bipolar measurements of the face. Finally, it would be necessary to familiarise participants with each face prior to a familiarity decision and this would be a lengthy procedure. Collectively, the endeavour necessary to investigate this question exhaustively would be extremely time consuming and it is considered that it would be more revealing to measure the effect of mirror reversal using other established paradigms in face recognition literature.

Experiment 7: Repetition priming from reversed and non-reversed images of celebrities

Introduction

In their 2002 paper, Brooks Rosielle and Cooper demonstrated that priming of familiar faces is invariant to the lateral orientation of the test stimulus. In line with previous research demonstrating the same null-effect in object recognition (Biederman & Cooper 1991) they concluded that our representations of familiar faces are *'invariant to mirror reversal'* (p307) and that *'...no time consuming transformations of mirror orientation occur during the process of face recognition'* (p312). That representations of faces and objects are similar in this respect is surprising given that face and basic-level object processing are dissociable by reference to their susceptibility to the 'inversion effect' and also by the supposed reliance on second-order relations in face recognition (see Maurer et al. 2002 for a review). In addition, previous research has shown the cognitive system to be sensitive to mirror reversal of faces (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003).

The priming paradigm used by Brooks *et al* (2002) was identical to that used in Biedemann and Cooper's (1991) demonstration of reflectional invariance in object recognition. In this task participants are shown a face for a short duration, and they must respond to the subsequent presentation of a name by signalling whether it matched the face or not. The faces at test were either presented in a previous name-verification decision (primed) or had not been seen previously (non-primed). The typical finding is that primed images are responded to quicker than non-primed images, but in their experiment Brooks et al found no additional advantage for the image being presented in the same orientation in priming and test. Crucially, Brooks *et al* (2002) used the same image at priming as they did at test, and so it is uncertain whether their null effect is constrained to same-image priming, or whether an effect of reversal might be detected by using different images in the two phases of the experiment.

In this experiment I sought to replicate Brooks *et al* (2002) using a different repetition-priming paradigm. The paradigm used in this experiment has been used in previous studies of face priming (e.g. Jenkins et al. 2002). Subjects are required to make a nationality decision to images of famous celebrities in the priming phase of the experiment, and priming is measured by the degree to which this exposure facilitates response time in a subsequent familiarity decision relative to non-primed stimuli. In this experiment a different image of the primed identities was used at prime phase and test phase, and in addition the experimental manipulation of reversal was performed on images presented in the prime phase (not the test phase as in Brooks et al, 2002). It is considered that these differences in design will probe the cognitive representations more directly and therefore provide a better test of whether face priming is sensitive to lateral orientation. Repetition priming is modulated by the similarity between prime and test stimuli (e.g. Bruce & Valentine, 1985) and therefore it was hypothesised that reversal would lead to a reduction in the degree of observed priming.

Method

Stimuli and Materials

Images of ubiquitous celebrity faces were sourced from the World Wide Web. Only images in which there was writing visible in either the foreground or the background were selected for use in the experiment, so as to ensure that the images had not been subjected to mirror reversal prior to publication on the Internet. Thirty celebrities were selected for use in the experiment on the criterion that it was possible to obtain two images in which some form of writing was visible (typically at promotional events and outside award ceremonies), and all images were full-face shots with both ears visible. For these thirty faces, a mirror-reversed copy was created for use in the priming phase (experimental manipulation). A further thirty celebrities (of equivalent fame) were used as non-primed stimuli in the test phase of the experiment, and in addition images of sixty unfamiliar faces were taken from an existing database for use in the familiarity decision phase of the experiment.

Each of these images was cropped neatly around the head and resized to 300 by 500 pixels. All images were presented in greyscale and were histogram equalised. Images were centrally presented with a black background on a computer monitor displaying a resolution of 1680 by 1050 pixels and measured roughly 7.5cm by 10cm (pertaining to a visual angle of 7.2 by 9.5 degrees of visual angle at a distance on 60cm).

Design and Procedure

Thirty two (21 female) subjects with an average age of 20.4 (sd=2.8) were recruited from the undergraduate population at the University of Glasgow and they either received course credit or participated on a voluntary basis. Two subjects were removed prior to analysis because of technical problems with the data output.

Each participant was sat in a darkened room so that their eyes were level with the centre of the screen and were told that their task was to decide whether the faces displayed on the screen were American or non-American (two alternative forced choice). Trials consisted of a fixation cross for 1500ms, followed by the face for 500ms, and then a blank screen indicating that they must respond. There were thirty trials in total (15 original orientation/15 mirror reversed) and the priming phase lasted only a minute or so. The identity of the reversed images was counterbalanced across subjects to ensure that effects of distinctiveness and familiarity were controlled.

Subjects were given instructions for the test phase roughly a minute after the prime phase had been completed. Participants were told that they were to be shown a series of faces and that they must respond as quickly as possible to whether the faces were familiar or unfamiliar to them. They were then shown a series of images comprising of 30 new images of the primed identities (all in their veridical orientation), 30 famous unprimed faces, and 60 images of unfamiliar faces. Order of presentation was fully randomised. Responses and response latencies were recorded for subsequent analysis.

Results

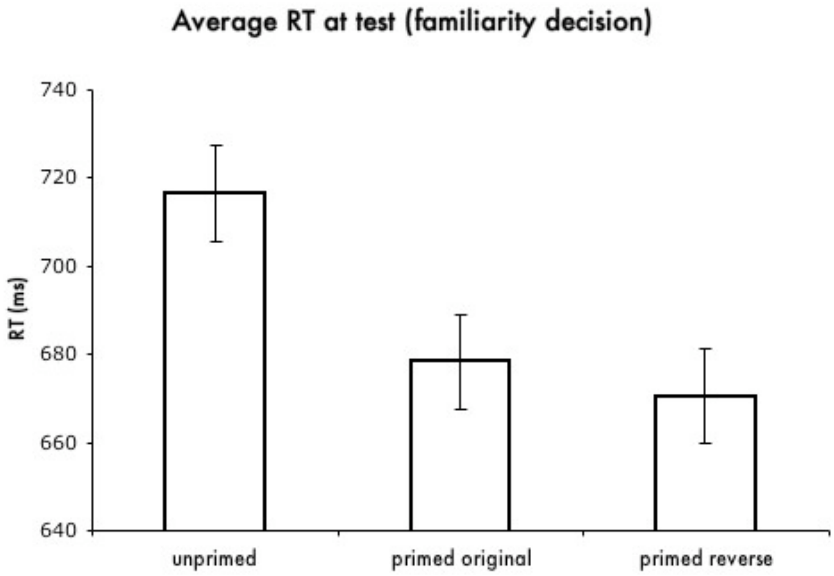

Figure 3.11: Mean reaction time data for 'familiar' responses (error bars denote 95% confidence intervals, calculated using the method described by Loftus & Masson, 1994).

Reaction time data coded as a 'familiar' response in the familiarity decision task was analysed using a one-way ANOVA. There was found to be a main effect of experimental condition [F(2,29)= 22.299; p<0.01]. Planned t-tests confirmed that this effect was reflective of the faster reaction for items primed with either the original orientation [t(29)= 5.20; p<0.01] or the mirror orientation [t(29)= 6.23; p<0.01]. There was no significant difference between reaction time data for items primed by the original orientation relative to items primed in their reverse orientation [t(29)=1.02; p>0.05]. This analysis, whilst confirming that a general priming effect was observed, shows priming to be invariant to the lateral orientation of the prime stimulus.

Discussion

The results reported in this experiment are in agreement with Brooks *et al* (2002), and confirm that face priming is invariant to lateral mirror orientation. It can now be accepted that the null effects reported by Brooks et al were not limited to the paradigm they used. However, in the context of previous demonstrations showing a sensitivity to the lateral orientation (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003), and given that an effect of reversal on name verification was detected in experiment three, it is uncertain whether this invariance in priming studies translates to a more general invariance of our face representations to mirror reversal. More certainly, it appears that whether or not an effect of mirror reversal is detected depends largely on the method used to probe the memory representations.

Whereas I failed to detect a difference in priming with mirror reversed versus veridically orientated stimuli, I did find a difference between these stimulus conditions in response latency when participants were tested using a name-verification paradigm. This is somewhat puzzling given that the differences between the name-verification paradigm used in experiment four and that used in the test phase by Brooks *et al* (2002) are slight. Furthermore, if existing data is to be accepted, it would appear that the veridical orientation of faces is embodied in our cognitive system (e.g. Rhodes, 1986). If this is the case then it is necessary to dissociate mechanisms responsible for priming effects from those enabling the sensitivity to mirror reversal reported in previous research.

Previous research has shown that memory for the left-right orientation of familiar objects (a US penny) is improved when participants are encouraged to visualise the stimulus prior to identifying its mirror orientation, and that when the use of imagery is discouraged performance is typically at chance (Kosslyn & Rabin, 1999). Although this finding has not been replicated with faces it does suggest that for object memory at least there is some dissociation between the memory representations (which appear to preserve left-right orientation) and the process that matches this representation to an incoming stimulus (which is not sensitive to left-right orientation). If this were to be true for face stimuli also it may explain why an invariance to mirror reversal was found in the present experiment but not in the name verification paradigm used in experiment four.

In the name verification procedure used in experiment three the name of a familiar face is shown previously to the presentation of the face and participants make a 2AFC decision (match/mismatch) to the image. In this case it can be assumed that the memory representation of the face is activated prior to the image being presented, much like the procedure used for successful identification of left-right orientation by Kosslyn & Rabin (1999). This procedure is in contrast to that used by Brooks *et al* (2002) where face images were presented prior to presentation of the name and a 2AFC decision was made to the person's name. This paradigm discourages the use of imagery and this may explain recognition was unaffected by reversal in Brooks *et al* (2002) and also why repetition priming is invariant to mirror reflection.

Experiment 8: The effect of mirror reversal on memory recall for faces

Introduction

McKelvie (1983) demonstrated using a classical recognition paradigm that memory for faces is disrupted by mirror reversal. However, given that the data reported thus far is inconclusive as to whether cognitive representations of faces specify lateral orientation, it was decided that this finding should be subjected to further scrutiny. The series of experiments reported by McKelvie (1983) all used the same images at learning as at test, and this limits the scope of their conclusions. Given that face recognition (as defined in the introduction to this chapter) is characterised by the ability to abstract information from our experience of a face in a given environmental context, using exactly the same image at test may enable subjects to use mechanisms other than face recognition to recall the faces from memory. Also, image specific memory can be dissociated from face memory (Schweinberger et al, 2002; Dyer et al. 2002), so a replication of McKelvie's (1983) data using a different image at learning than at test was considered necessary.

Method

Stimuli and Materials

Twenty female and twenty male identities were chosen from an in-house database of unfamiliar faces. Of these 40 identities, 20 (10 male, 10 female) were selected for use in the learning phase of the experiment and the remaining 20 were to be used as ‘new’ stimuli in the test phase. For each of the 20 faces chosen for the learning phase, two images that were taken using different cameras and under slightly different lighting conditions were used as stimuli. All experimental stimuli were presented in 256 levels of greyscale on a 15” computer monitor set to a screen resolution of 1152 X 864 pixels, and each face measured roughly 220 X 330 pixels. Images were all centrally presented.

Design and Procedure

Forty-four people participated in the study and all were recruited from the undergraduate and postgraduate populations at the University of Glasgow. Of the 44 subjects, 28 were female and 16 were male, and it was ensured that an equal number of males and females were assigned to the between-group conditions. The average age of the sample was 21.7 (SD= 2.6).

Participants were sat in front of the computer monitor and were told that the task they were about to take part in would test their memory for unfamiliar faces. All subjects were shown the same 20 images and told to remember the people in the pictures as well as possible because their memory for the faces would be subsequently tested. Twenty seconds after the final learning image had been presented the instructions for the test phase appeared on the screen. Subjects were told that they were to be shown 40 images of 40 different people, half of which they would have seen in the learning phase, and half of them they would not have seen. Participants were to respond as to whether the face was ‘old’ (i.e. appeared at learning) or ‘new’ by pressing the key corresponding to their decision.

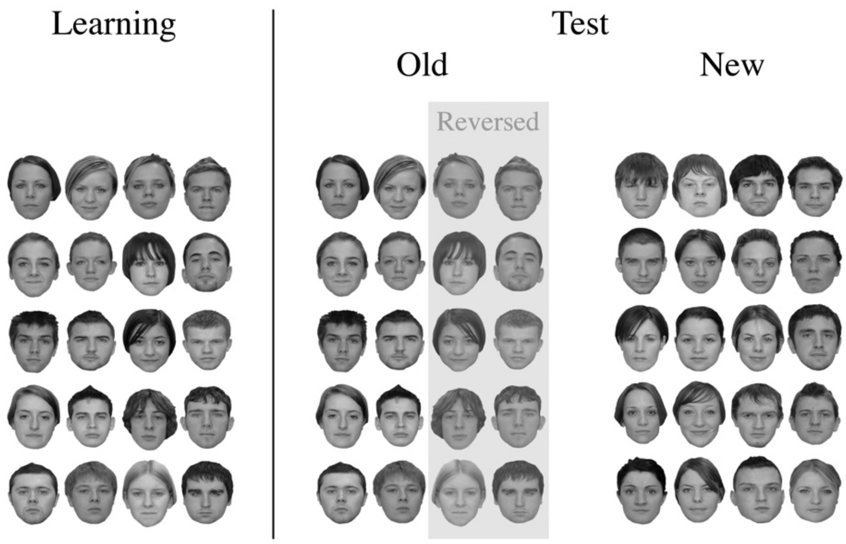

Figure 3.12: An illustration of the experimental design used in Experiment 8

At test phase, half of the participants (8 male, 14 female) were tested using the same images that had been presented in the learning phase, and the other half were tested using images taken by a different camera, under slightly different lighting conditions (see Figure 3.13). In addition, half of the test images were mirror reversed relative to the learning image (within subjects factor). Counterbalancing ensured that the reversal manipulation was not confounded by the distinctiveness of the individual faces used.

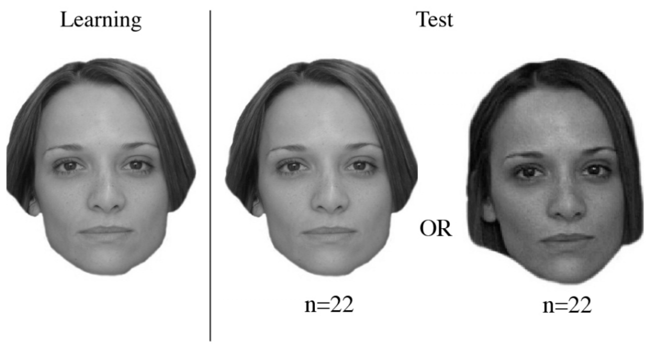

Figure 3.13: The between-group experimental manipulation: Subjects either saw the same image, or a different image of the 'old' identities at test.

Results

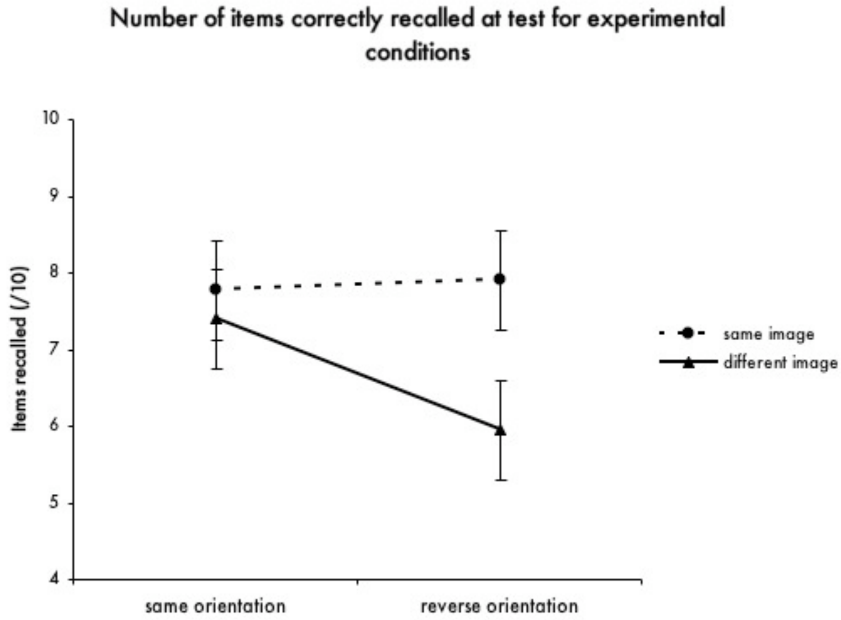

Figure 3.14: Error bars denote 95% confidence intervals, which were calculated using the method described by Loftus & Masson (1994).

A mixed-factor, two way ANOVA showed the effect of test image (same image/ different image) to be statistically reliable [$F(1,21)= 12.66$; $p<0.01$]. The effect of test orientation was non-significant, although only marginally so [$F(1,21)= 3.56$; $p=0.66$]. Crucially, there was also a statistically reliable interaction between the experimental factors [$F(1,21)= 5.39$; $p<0.05$].

An analysis of simple main effects showed that the effect of test image was significant only for images reversed at test [$F(1,21)=17.65$; $p<0.01$], and the effect of test orientation was only significant when a different image was shown at learning and test [$F(1,21)=8.85$; $p<0.01$].

Discussion

These results confirm that memory representations of faces are selectively responsive to lateral orientation. This finding is in agreement with McKelvie's (1983) previous demonstration of disrupted memory performance when test images are presented in the opposite lateral orientation relative to learnt images. In addition, this data combines with previous demonstrations showing cognitive sensitivity to the lateral orientation of faces (e.g. Mita et al. 1977; Rhodes, 1986; Brédart, 2003) to provide evidence that cognitive representations of faces retain the left-right organisation of the veridical percept.

The significant interaction reported does however necessitate some qualifying remarks. Primarily, it is difficult to explain the fact that the effect of reversal did not have a significant effect on memory performance in the same image condition. In all of the five experiments published by McKelvie (1983) a difference between memory recall of reversed versus non-reversed images was reported, and this occurred despite the fact that the same images were used at leaning as in test. There is one clear difference between the design of these experiments and that of the present study however. The delay between learning and test phase was only thirty seconds in this experiment, yet McKelvie (1983) reports a rest duration of ten minutes.

The most straightforward explanation for why this difference in methodology may account for the contradictory results is that the same-image task in the present study was too easy. If the performance on this task is at ceiling level then this may explain why an effect of reversal was not detected. Were participants' memory for the images sufficiently reliable, it may be that seeing the test image in its original orientation provided no detectable advantage. This interpretation of the data is consistent with the findings of McKelvie (1983), as the long retention interval used in this study made for a more difficult memory task. However it should be noted that only 2 out of 22 participants correctly recalled all of the test images in the same image condition, which is not indicative of ceiling effects.

Another proposed explanation is that the representation of orientation is facilitated by short-term memory decay and subsequent abstraction in long-term visual memory. This proposal is somewhat paradoxical, as one would expect that if anything the process of abstraction that occurs in long-term memory would enable the construction of a representation that is invariant with respect to orientation. However, such an explanation is not without precedent. It has been argued that the visual pathway represents the topography of the retina as mirror reflections across cerebral hemispheres (Orton, 1925; Noble, 1966, 1968; Corballis & Beale, 1976; Gross &

Bornstein, 1978), and this theory is supported by studies of commisured animals (Noble, 1966, 1968) and by neuronal investigations (Berlucchi & Marzi, 1970; Rollenhagen & Olson, 2000; Baylis & Driver, 2001).

Some have argued that the tendency for many species to confuse mirror images is a consequence of the bilateral symmetry of the nervous system (Corballis & Beale, 1976; Rollenhagen & Olson, 2000), and behavioural data in support of this assertion has been reported (Bradshaw et al, 1973). As face recognition appears to rely on brain areas predominantly in the right hemisphere (e.g. Sergent et al. 1992), it may be that as visual information 'climbs' the hierarchy of the visual system the representation becomes of a more uni-directional nature due to the increased laterality of function. Any such explanation must be proposed tentatively however, as it is in opposition to established conceptions of object constancy, where invariance to orientation is thought to rely on higher-level cortical processes (e.g. Perrett & Oram, 1993).

Experiment 9: Explicit memory for the left-right orientation of face images

Introduction

In the previous experiment it was shown that mirror reversal significantly reduces our ability to recall faces from memory when a different image is used at test. This suggests that the asymmetry of a face is represented in the mind after just one presentation and that the direction of this asymmetry is also specified. When the same image was used in learning and test phases however, performance was unaffected by mirror reversal.

The present experiment was designed to test whether or not the orientation of images of faces can be explicitly remembered. Previous research has shown that memory recognition for the orientation of familiar objects is around 70% using a same/different orientation memory procedure (Madigan & Rouse, 1974) but as yet this has not been tested using face stimuli. That recognition was not disrupted by mirror reversal in the previous experiment when the same image was used at test as at learning suggests that an image of a face in one orientation is cognitively equivalent to the mirror image of that image. It is predicted therefore that memory for the orientation of an image should be at chance.

Method

Stimuli and Materials

Twenty-six famous and twenty-six non-famous faces were sourced from the Internet. Images were chosen that contained writing in either the foreground or the background so as to ensure that pictures had not been mirror-reversed prior to publication on the world-wide-web. Face images were rotated so as to bring the pupils into alignment with the horizontal plane, cropped around the outline of the head, converted to greyscale and resized to 380 pixels X 570 pixels.

All fifty-two images were used at training, and half of the images in each condition were reversed prior to presentation so that there were no effects of prior memory. The face images measured 8cm by 11cm, pertaining a visual angle of 7.6 by 10.5 degrees (at 60cm), and were centred 13cm apart (visual angle= 6.2 degrees).

Design and Procedure

32 subjects from the undergraduate and postgraduate populations at the University of Glasgow agreed to participate in the study. The sample consisted of 19 female and 13 male subjects ranging from 17 to 24 years of age (mean 21).

Subjects sat in front of a Macintosh workstation running the software Psyscope and the experimenter explained their task to them. They were to be shown the 52 face images (half of which had been mirror reversed prior to presentation) in a random order and the task was to memorise the images as well as possible because their memory for the images would be tested later. Images appeared in the centre of the screen for 2000ms and were preceded by the presentation of a fixation cross for 750ms.

On completion of the training phase subjects were instructed to complete an unrelated filler task, which took three minutes in total. After the three minutes had elapsed the experimenter instructed the subject to begin the test phase. Subjects were told that they were to be shown the same set of images as had been presented to them in the training phase, but that some of the images had been 'mirror reversed'. Their task was to decide if the image was as it had appeared earlier or if it had been reversed. Presentation of the images was preceded by a fixation cross (750ms) and the image was displayed on the screen until the subjects responded.

Results

Participants overall performance accuracy was found to vary significantly above chance levels (50%) using a one sample, one tailed t-test [$t(31)=1.939$; $p<0.05$]. However, memory for image orientation [mean= 53.6%; sd=10.2%] did not vary significantly between experimental conditions [$t(31)=.891$; $p=.380$]. In addition, one-sample, one tailed t-tests were executed on data from the two conditions separately. Image memory performance for images of familiar faces [mean= 54.9%; sd=11%] was found to vary significantly from chance [$t(31)=2.525$; $p<0.01$], however image memory for unfamiliar faces [mean= 52.3%; sd=9.2%] did not vary reliably from chance [$t(31)=1.438$; $p>0.05$].

Discussion

Contrary to the prediction made, memory performance for left-right orientation of face images did vary significantly from chance. Overall performance was very poor however and participants performed above chance only when memorising the orientation of images depicting famous people. It is most likely that this difference resulted from some image-level artefact as opposed to above chance performance being a product of familiarity. As half of the images were reversed prior to presentation in the learning phase the effect cannot be due to residual memory for the veridical orientation of the famous faces. Any effect of familiarity therefore can only be explained by a mechanism that promotes a more robust memory trace for some or all of the visual characteristics of an image when the subject is familiar.

Though this provides a suitable explanation of the data reported here it is considered more likely that an artefact of the familiar stimulus set has caused the above chance memory performance. One difference between the unfamiliar and famous stimuli is there are more instances of gaze aversion in the familiar stimuli (8/26) than in the unfamiliar stimuli (3/26). It could be that memory for the orientation of familiar face images has been facilitated by this asymmetry. Given that recall accuracy for the mirror orientation images of more saliently asymmetric objects has been reported to be 70% (Madigan & Rouse, 1979) it would appear that the degree of image level asymmetry contributes to improved accuracy on this task. Therefore it is likely that hair style, eye gaze and other transient asymmetries would also improve memory for mirror orientation and this may explain the above chance performance reported for familiar faces.

The data reported here is in agreement with the results of preceding experiments in this chapter. The apparent difficulty in remembering the left-right orientation of images of faces reflects the slight effect of mirror reversal on face recognition and also the null effect of image reversal on unfamiliar face memory reported in the previous experiment. It would appear that some underlying process in the representation of faces promotes a perceptual invariance to the mirror orientation of faces. In addition, the ability to detect reversal appears to be confounded by the high degree of symmetry present in faces.

Chapter Summary

The principal hypothesis to be tested in this chapter was that memory representations of familiar faces specify lateral orientation. Previous studies have shown sensitivity to the mirror reversal in affective judgements (Mita et al. 1977), familiarity judgements (Rhodes, 1986; Brédart, 2003), and memory for images of unfamiliar faces (McKelvie, 1983). In the experiments reported in this chapter sensitivity to reversal was additionally detected in name verification (experiment 4) and also memory for unfamiliar faces (experiment 8). This difference in behavioural response to mirror reversed faces was not detected in the name verification responses of experiment 5 and 6, the repetition-priming paradigm in experiment 7, nor in the same image memory test in experiment 8. Whereas the null effect in experiment 5 can most likely be explained by subjects' over-exposure to the experimental stimuli, the null effect reported in experiment 7 would appear to confirm that repetition priming is invariant to the lateral orientation of the stimulus.

In addition, the differences between responses to images in reversed orientations reported in this chapter are small. Presenting images of familiar faces in their reverse orientation (experiment 4) does not affect accuracy of response and the only detectable effect was a mean reaction time difference of 32ms relative to responses to faces in their original orientation. Furthermore, in a second experiment (experiment 5) this difference was not observed, most probably due to practise effects. It would appear that the visual system is minimally responsive to changes in the mirror orientation of faces, and it could be that the null effects of reversal are due to a lack of sensitivity in the measure used to detect the effect of reversal. After all, the asymmetries in a face are small (Farkas, 1994; Farkas & Cheung, 1981; Farkas & Munro, 1978) and therefore one might expect that any sensitivity to reversal is also small. This problem of measurement sensitivity is highlighted in experiment 8 where in the less difficult task of same-image memory there was no effect of reversal, whereas in the more difficult different-image memory task there was a strongly significant effect of reversal.

It is interesting therefore to speculate as to why our sensitivity to mirror reversal is so slight. As has already been mentioned, the visual information disambiguating orientation is typically modest (e.g. Farkas & Cheung, 1981). However as our perceptual mechanisms are tuned to detect minute differences in configuration of facial features (Haig, 1984), and face recognition depends on the ability to discriminate individuals based on these small differences, one might reasonably expect that the effect of reversal would be more easily detectable. Therefore it is likely that there is a further reason for the resilience displayed by cognitive representations to mirror transformation. It may be that the slight effects of mirror reversal are related to a more general tendency to confuse mirror images (e.g. Sutherland, 1960; Sekuler & Houlihan, 1968; Hamilton & Tieman, 1973). If the mirror reflection of a perceived object is represented in the visual system (Noble, 1966, 1968) then this may account for null effects of mirror inversion (e.g. Brooks et al. 2002).

Though this may be the case, invariance to linear transformations of faces in recognition has previously been reported for transformations other than mirror reflection (Hole et al. 2002). However, it has been suggested that the mechanisms enabling this invariance may have evolved out of a necessity for object constancy despite changes in depth rotation (Hole et al. 2002; Sinha et al. 2006), and it may be that a similar explanation can be given for the apparent robustness of face recognition systems to mirror reflection. That reversed faces are recognised despite the alteration of configural information (experiments 2, 4, 5, 6 & 7) may be a result of some automatic normalisation process that has evolved for means more general than the recognition of mirror-reversed faces. It has been shown that 'mirror confusion' most probably results from the neurons in the visual system responding similarly to lateral mirror images (Rollenhagen & Olson, 2000). Tolerance to mirror reversal of familiar (chapter 3) and unfamiliar (chapter 2) faces may be enhanced by such generalisation.

Though face processing is generally unimpaired by mirror reflection, evidence of sensitivity to this transformation has been detected both in this chapter (experiments 4 & 7) and in previous research (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003). Furthermore, though the differences between response latencies to original and reversed images in experiments 5 and 6 were not significant, in both cases a faster mean response time was reported for veridically orientated stimuli. Collectively these results prove not only that the subtle asymmetries present in a face are represented in the mind, but also that the left-right orientation of these asymmetries is preserved to some degree. For the cognitive system to respond differently to a reflection than to a veridical percept some form of coding must specify the veridical orientation in memory. That said, though form and mirror orientation can be considered as mutually exclusive properties (e.g. Rock, 1973; Hummel and Biederman, 1992), it would appear that in the case of face perception they are to some extent associated.

Chapter 4

Explicit memory for the left-right orientation of familiar faces

Introduction

Though the bilateral asymmetry of a given face is typically small, humans are nonetheless sensitive to this information. In the previous chapter two experiments were reported where a sensitivity to mirror reversal was detected. Though I failed to replicate the effect of reversal on reaction time reported in experiment four in two subsequent experiments, the mean response latencies in these experiments were nonetheless faster for images in their original orientation than it was for mirror reversed images. In addition, previous research has shown recognition systems to be sensitive to mirror reversal (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003). It would appear therefore that although the representation of reflectional asymmetry in the mind is a non-salient quality, the asymmetry is nonetheless represented and its left-right orientation is specified. In the present chapter three experiments are reported which test whether or not this sensitivity is sufficiently robust to enable reliable identification of the veridical orientation of a familiar face.

Previous research has shown that it can be surprisingly difficult to explicitly recall directional information relating to objects which we encounter everyday. Nickerson and Adams (1979) demonstrated that people have significant difficulty in recalling the appearance of a US coin and that they were particularly poor at remembering the orientation of Abraham Lincolns head, a finding which has been since replicated (Kosslyn & Rabin, 1999). Further research has shown this difficulty to be strikingly profound, with Rubin and Kontis (1983) reporting that a significant majority of participants *incorrectly* recalled the head's orientation and studies using UK coins have reported similarly poor performance (Jones, 1990; Martin & Jones, 1997). This peculiar finding is thought to occur due to a generalised bias towards leftward head direction in images, coins and paintings (e.g. McManus & Humphreys 1973; McKelvie & Aikins, 1993).

Though this research suggests that the left-right orientation of objects is not specified in memory of common objects, a more recent study has demonstrated above-chance performance at recalling their orientation. Kelly *et al* (2001) asked British subjects to identify the real-world orientation of a two-pence coin and a postage stamp using a two-alternative forced choice (2AFC) paradigm and found performance on this task to be 70%. The relatively high performance level obtained by Kelly *et al* (2001) is thought to be a product of the sensitivity inherent in the 2AFC paradigm, and the authors report accuracy of 80% when testing Japanese subjects on asymmetrical stimuli ubiquitous within their culture (see figure 4.1).

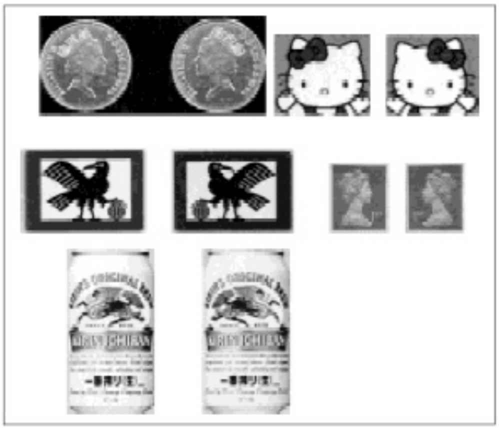

Figure 4.1: Japanese and British stimuli used by Kelly *et al* (2001).

Given the sensitivity of this approach it provides a good test of whether participants can explicitly identify the correct orientation of a familiar face. One would anticipate this task to be very difficult however as the asymmetries in faces are slight in comparison to the gross asymmetries displayed by the stimuli used by Kelly *et al* (2001), where the direction is disambiguated by coarse elements of the image. It is surprising therefore that the ability to discriminate between mirror images of faces has been previously demonstrated. Rhodes (1986) examined whether or not participants would chose the original orientation of a familiar face when asked to indicate which of two simultaneously presented images (one original orientation, one reversed) appeared most 'like' the person. Using this mirror discrimination paradigm (MD) she reported that subjects correctly chose the original image on 72% of occasions for highly familiar faces and 58% of occasions for moderately familiar faces. This was taken as evidence that the mind preserves the subtle configural information of faces and furthermore that the lateral direction of this information is also retained.

The experiments reported in the present chapter will ascertain whether this memory for lateral orientation affords subjects reliable performance on a task where they are explicitly asked to identify the 'real-world' orientation of a face. The studies reported in the previous chapter demonstrated that effects of mirror reversal on the cognitive processing of faces are difficult to detect and when they are detected these effects are typically slight. For this reason it is anticipated that participants will find it difficult to recall the real-world orientation of faces from memory when explicitly instructed to do so. Furthermore, the tendency for mirror reversals of faces to be treated as equivalent stimuli (as reported in experiment 2) is predicted to additionally contribute to the difficulty of this task.

Experiment 10: Perceptual sensitivity to mirror reversal of human faces

Introduction

Adapting the 2AFC paradigm used by Kelly *et al* (2001) this experiment sought to determine whether the sensitivity to reversal of familiar faces reported in chapter 3 would enable reliable performance on an explicit test of memory for reflectional orientation. Given that previous research has shown performance on this task to be poor even with grossly asymmetric stimuli that are perceived in everyday life, it was expected that participants would find it very difficult to explicitly identify the correct orientation of familiar faces.

Method

Stimuli and materials

Twenty-four famous and twenty-four non-famous faces were sourced from the Internet. Images were chosen that contained writing in either the foreground or the background so as to ensure that pictures had not been mirror-reversed prior to publication on the world-wide-web.

Face images were rotated so as to bring the pupils into alignment with the horizontal plane, cropped around the outline of the head, converted to greyscale and resized to 380 pixels X 570 pixels. For each image a mirror-reversed copy was created by flipping the image horizontally using Adobe Photoshop. Displays were then created by randomly allocating a position (left/right) to the mirror- reversed image for each identity. Original and mirror-reversed images were arranged side by side on a white background (see figure 4.2) measuring 1000 by 600 pixels.

Design and Procedure

18 subjects (14 Female) between the ages of 18 and 32 (mean= 21) from the undergraduate population at the University of Glasgow consented to participate in the study. Participants were paid in cash or, if they were first year Psychology students, received course credits.

Figure 4.2: An example of the stimulus presentation (original on the left).

Subjects were sat in front of a Macintosh workstation running the software Psyscope and the experimenter explained their task to them. Participants were told that in each trial they would be presented with a pair of images and that these images would be of the same person and would be identical in everyway but that one of the images would have been mirror-reversed. Their task, they were instructed, was to decide which was the original image and to make their response by pressing the key on the left if the original appeared on the left and pressing the key on the right if the original appeared on the right. Subjects were told that they could take as long as they liked before making a response, and that they may find the task difficult to the extent they might feel as though they were guessing.

Presentation was fully randomised. There were 48 trials in total (24 familiar/ 24 unfamiliar), and the task took an average of six minutes to complete. For each trial, a fixation cross appeared for 750ms, prior to the presentation of the image-pair.

Results & Discussion

Mean performance (i.e. percentage correct) was higher for familiar (mean= 59.8; sd=9.84) than for unfamiliar faces (mean=47.6; sd=6.84). A related measures t-test showed this difference to be reliable [$t(17)=3.83$; $p<0.001$]. A one-way, one tailed t-test confirmed performance for familiar faces was significantly greater than chance [$t(17)=4.762$; $p<0.01$] however performance on unfamiliar faces did not vary significantly from chance levels [$t(17)=1.752$; $p>0.05$].

A follow-up-analysis was carried out to ascertain whether certain image level qualities of the familiar face stimulus set were affecting performance on the discrimination task. It is possible that participants were performing the MD task by relying on asymmetries other than those present in the structure of the faces themselves, and so additional analyses were deemed necessary. Participants were found to be no more accurate at discriminating between mirror images of celebrities with asymmetric hairstyles (n=11) relative to those with symmetric hairstyle (n=13) using a two-tailed significance test [$t(23)=0.784$; $p>0.05$]. Likewise subjects did not perform more accurately when the gaze of the celebrity was averted (n=12), [$t(23)=0.049$; $p>0.05$] or when the angle of their head was averted (n=13), [$t(23)=0.851$; $p>0.05$].

It is unlikely therefore that participants were making their decision based on unstable factors such as the hairstyle of the person in order to disambiguate the left-right orientation, and if they were it did not afford them reliable accuracy. In replication of Rhodes (1986) it was found that participants did not correctly identify the veridical orientation of celebrities with asymmetric hairstyles more accurately than those with symmetrical hair, and furthermore head angle and gaze direction did not affect accuracy. In addition, the celebrities used did not have defining surface features (e.g. birthmarks) that could be used by participants to disambiguate the mirror orientation of the images.

Discussion

The results of this experiment show that when a face is familiar, subjects can identify its veridical mirror orientation with accuracy significantly superior to chance performance. This finding supports earlier contentions (e.g. Rhodes, 1986) that the cognitive representations underlying face recognition are sufficiently detailed to code fine-grain configural asymmetries in the face. Though accuracy significantly exceeded that which would be predicted if subjects were not selectively sensitive to the original orientation of familiar faces, an average accuracy of 60% represents a poor level of performance. Accuracy in the experiment reported here is notably less than the 70% reported by Rhodes (1986), and there are a number of methodological differences that may explain these incongruous observations.

Firstly, it may be that the stimuli used in the present experiment were not as familiar to participants as those used by Rhodes (1986) where the faces used were of participants' colleagues with whom they have daily face-to-face contact. It could be that both the quantity and quality of prior exposure to the faces are important factors in successful performance on the MD procedure. The familiar faces used in this task are typically encountered with less frequency than colleagues and familiarity may be additionally compromised by the nature of this exposure. Brédart (2003) contended that the competition between representations of our own face as perceived in photographs (veridical orientation) and when looking in the mirror (mirror reversed orientation) causes poorer performance with our own faces relative to other highly familiar faces on a MD task. As was discussed in the previous chapter, famous faces are often reversed prior to publication in print media, and the prevalence of mirrored stimuli may therefore be an additional source of confusion for participants in this task.

A further possibility is that the task instructions given to participants in the present experiment has in some way impaired performance relative to the performance of subjects in the task set by Rhodes (1986). In the present experiment participants were explicitly aware that one of the images had been mirror reversed and were asked to indicate which image was in the 'real world' orientation, whereas in Rhodes' paradigm participants were unaware of the image manipulation and were simply asked to decide which image looked 'most like' the person. It may be that by encouraging participants to use their knowledge of the veridical orientation in an implicit manner Rhodes has made the task easier for them. In the present experiment all participants reported finding the task very difficult and that they felt as though they were guessing, and another common response was that they often 'second-guessed' themselves. When making explicit judgements participants may be more likely to 'second-guess' themselves whereas the 'most like' decision might encourage participants to use their initial impression more confidently.

In a previous experiment using the same explicit MD as used here Kelly *et al* (2001) found that 70% of participants could correctly identify the veridical orientation of commonly perceived objects, and for some more salient objects in their visual environment a success rate of 80% was reported. This performance is notably superior to the accuracy displayed by participants in the present experiment where the same judgement was made to images of famous faces. This difference is not surprising as the mirror orientation of the stimuli used by Kelly *et al* (2001) was specified by salient and stable visual characteristics (i.e. the direction of the eagles head, see figure 4.1).

In the present study however, subjects had to rely on the subtle morphological asymmetry of the faces. Therefore, given that the ability to mirror discriminate would appear to result from the structural asymmetries of the face stimuli used in this experiment, it is little wonder that performance on this task was inferior to that reported by Kelly *et al* (2001). The structural asymmetries of faces are typically slight (Farkas & Cheung, 1981) and this is in contrast to the grossly asymmetric figures used in the MD task reported by Kelly and colleagues. Perhaps more surprising is the fact that subjects could perform this task at all given that the two images presented to participants were so similar.

Experiment 11: Perceptual sensitivity to mirror reversal of human faces II

Introduction

A limitation of the previous experiment is that it failed to confirm that the above chance performance in the familiar condition was a product of familiarity *per se* or whether some artefact of the celebrity images was responsible for this effect. Though no evidence was found to suggest that either hairstyle (e.g. directional partings) or gaze direction differed between the experimental conditions, it was considered that a further experiment was necessary in order to isolate the effect of familiarity more convincingly.

With direction of lighting, facial asymmetry and other transient asymmetries all potentially causing the effect reported in the previous experiment, an experiment was designed that ensured such superficial qualities would not obscure any difference between familiar and unfamiliar conditions. Using a design in which all face stimuli can be used as familiar items (for participants who are their workmates) and unfamiliar items (for participants from a different workplace) allowed for a more valid investigation. In addition, using images of participants' colleagues will help clarify whether the use of celebrity identities in the previous experiment contributed to the MD performance reported in the previous experiment.

Method

Stimuli and Materials

A digital camera was used to collect images of professors, lecturers, research staff and postgraduate students from the psychology departments at The University of Stirling and the University of Glasgow. From this database of images the twenty-four identities from each department that were anticipated to be most familiar to the postgraduate students and research staff in their respective departments were selected for use in the experiment.

Using Adobe Photoshop CS, each image was cropped around the contour of the head, converted to greyscale, and resized proportionately to 300 pixels in width. For each of the forty-eight images an array was created (figure 4.3) which contained the image in its original orientation and a mirror-reversed copy of the original placed side by side. The position of the original (left/right) was randomised, and the arrays measured 1000 by 600 pixels.

Design and Procedure

Mirror Discrimination Task

Subjects were recruited from the psychology departments at The University of Stirling and from The University of Glasgow. 14 Subjects were recruited from each department so as to counterbalance the images used within the two experimental conditions. The sample consisted of 13 female and 15 male subjects ranging from 21 to 45 years in age (mean= 25.8).

Figure 4.3: Example of experimental stimuli (original on left).

Subjects were sat in front of a Macintosh workstation running the software Psyscope and the experimenter explained their task to them. Participants were told that in each trial they would be presented with a pair of images and that these images would be of the same person and would be identical in everyway but that one of the images would have been mirror-reversed. Their task, they were instructed, was to decide which was the original image and to make their response by pressing the key on the left if the original appeared on the left and pressing the key on the right if the original appeared on the right. Subjects were told that they could take as long as they liked before making a response, that they may find the task difficult and that they were not to worry if they felt like they were guessing.

Presentation was fully randomised. There were 48 trials in total (24 familiar/ 24 unfamiliar) and for each trial a fixation cross appeared for 750ms prior to the presentation of the image-pair. On completion of the test phase subjects were shown each of the forty-eight faces used in the experiment and were asked to indicate whether or not they had been familiar with the faces prior to taking part in the experiment. On average the experiment lasted roughly six minutes

Symmetry Ratings

Twenty undergraduate students [9 female; mean age= 19.7 (sd=2.2)] from the University of Glasgow volunteered to rate the experimental stimuli. Psychology undergraduates were not used so as to ensure that participants were unfamiliar with the faces, and subjects were asked after they had completed the ratings if they had been familiar with any of the faces they had rated. Each participant was given the experimental stimuli and asked to rate them for asymmetry on a scale of 1 (very symmetrical) to 7 (very asymmetrical). The stimuli were presented in greyscale format on a sheet of white paper. Eight stimuli were presented per page, and the order of the six pages was randomised across subjects. Under each stimulus there was a rating scale and participants were asked to circle the number they had chosen to rate the image. The rating task typically took subjects 5 to 10 minutes to complete.

Results

Mirror Discrimination Task

Only data from trials where the intended condition and the actual condition were congruent was analysed (e.g. for participants from Stirling only the identities that were responded to as familiar and were from the Stirling subset of images would contribute towards the mean for the familiar condition, and only those confirmed to be unfamiliar from the Glasgow subset contributed to the unfamiliar data). Some overlap was expected as the two departments have a history of academic collaboration, and it was not certain whether the stimuli collected from each department would be familiar to all participants from the respective department.

Mean performance was higher for familiar (mean= 53.8; sd=2.09) than for unfamiliar faces (mean=50.2; sd=1.63). However, a related measures t-test showed this difference to be unreliable [$t(1,27)=1.58$; $p=0.126$]. When subjected to a one way, one tailed t-test performance data for the familiar condition was found to vary significantly from chance [$t(27)=1.788$; $p<0.05$]. Performance in the unfamiliar condition did not vary significantly from chance levels [$t(27)=0.108$; $p>0.05$].

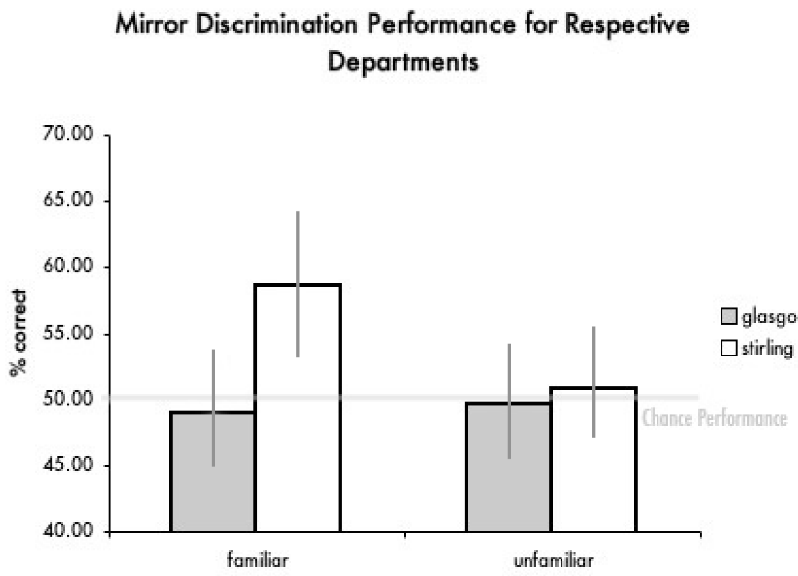

Figure 4.4: Bar graph showing the MD performance by participants from the two departments (error bars denote 95% confidence intervals).

When subjects' performance was split according to the department they belonged to, it was found that participants from Stirling performed significantly above chance for familiar [$t(1,13)=2.842$; $p<0.01$] but not for unfamiliar faces [$t(1,13)=0.135$; $p>0.05$], whereas the participants from Glasgow did not perform significantly above chance levels for either familiar [$t(1,13)=0.391$; $p>0.05$] or unfamiliar faces [$t(1,13)=0.150$; $p>0.05$].

Symmetry Ratings

An average asymmetry score was calculated for each stimulus used in the experiment using rating data collected from the 20 independent raters. The difference in mean asymmetry scores between the two departments [Stirling mean= 3.56 (SD=0.74); Glasgow mean= 3.39 (SD= 0.83)] was not significant [$t(38)= 0.69$; $P>0.05$]. Furthermore, there was found to be no correlation between the asymmetry of a face and the proportion of correct responses made to that face in the MD task [$r=0.119$; $p>0.05$].

Discussion

Mirror discrimination performance in this experiment was only marginally superior to that which would be predicted if participants were making guess decisions. Not only did overall performance for familiar faces differ only very slightly from chance, but also in one of the departments performance was statistically equivalent to chance. Given that the familiar stimuli used were images taken of the participants colleagues this is somewhat surprising, as one would expect these faces to be encountered as commonly as the famous faces used in the previous experiment. It is possible however that we spend more time inspecting the faces of celebrities than we do our colleagues as it may be deemed socially inappropriate to stare at a colleagues face for an extended period of time. Regardless, even if the faces were minimally familiar to participants a more robust level of performance would have been expected given that Rhodes (1986) reported an accuracy of 58% in a MD paradigm to faces of low familiarity. As previously discussed however, it could be that the explicit instructions used in the experiments reported here has made this task more difficult.

Performance was found to be significantly above chance when discriminating mirror images of familiar faces for participants from the University of Stirling but not for participants from the University of Glasgow. This difference does not appear to be reflective of a difference in the asymmetry of the stimuli themselves and therefore the most likely explanation for this is that participants from the University of Stirling were more familiar with the faces in their department relative to the participants from the University of Glasgow. Though these two psychology departments are similar in size (e.g. number of staff), it is possible that some other factor such as the closer proximity of colleagues' offices at Stirling may have caused more interaction between these colleagues. As no familiarity ratings were collected as part of this study it is difficult to assess the effect familiarity on performance, other than to say that a certain level of familiarity is necessary to perform the MD task above chance.

Experiment 12: The effect of familiarity on performance in a mirror discrimination task

Introduction

In the previous experiment performance of the mirror discrimination task was above chance for one group of participants but not the other. It is uncertain why this might be and it would not appear that this difference results from differences in the morphology of the faces themselves, or in the asymmetry of other aspects of the images. Using a different paradigm Rhodes (1986) has shown that increased familiarity improves MD performance and it is possible that the chance performance level displayed by participants from Glasgow University is a result of these participants being less familiar with their 'familiar' set relative to the participants from the University of Stirling. As no familiarity ratings were collected in the previous experiment, another experiment was designed to assess the effect of familiarity on performance in the MD task.

Method

Stimuli and Materials

Images of 14 contestants from the UK television show *Big Brother 8* were sourced from the worldwide web (http://uk.tv.yahoo.com/big-brother/photos/). These images were originally captured from live streams of the television footage and were in their original lateral orientation. In addition, images of the television show's presenters Davina McCall and Dermot O'Leary were used as stimuli. For each image a mirror reversed copy was generated using Adobe Photoshop CS. This reversal was then placed alongside the original image and the subsequent pairs were presented to participants, with the originals appearing an equal amount of times in the left-hand image as in the right. Image pairs measured 600 by 400 pixels and were presented in full colour (see figure 4.5).

Figure 4.5: Example of a experimental stimuli (original on left).

Design and Procedure

Sixty-eight participants [37 female, mean age= 18.2] were recruited from the Psychology summer school at the University of Glasgow and they were tested during the introductory lecture. Prior to testing participants were asked to specify whether or not they had watched the current series of Big Brother and if so to indicate the average amount of time spent watching the programme each week. They were also asked to give familiarity ratings to the names of the contestants that were to be presented. When all participants had given answers to these questions they were shown the series of face pairs in a random order (same order for all subjects) and asked to specify, by marking the appropriate box on their answer sheet, which image had *not* been mirror reversed. Each image pair was presented for five seconds each on a projector screen in a lecture theatre and the sixty-eight participants were tested in three separate sessions. After the task had been completed the students were thanked for their participation and their response sheets were collected.

Results

Participants were split into three groups based on how often they reported having watched *Big Brother 8.* The unfamiliar group consisted of participants who had never watched the series (n=20), the low familiarity group contained those participants that reported having watched the programme for less than an hour per week (n=19) and the high familiarity group was defined as those participants that had reported watching the series for an hour or more per week (n=29). The probability of correctly choosing the veridical orientation from the image pair was then calculated for each subject.

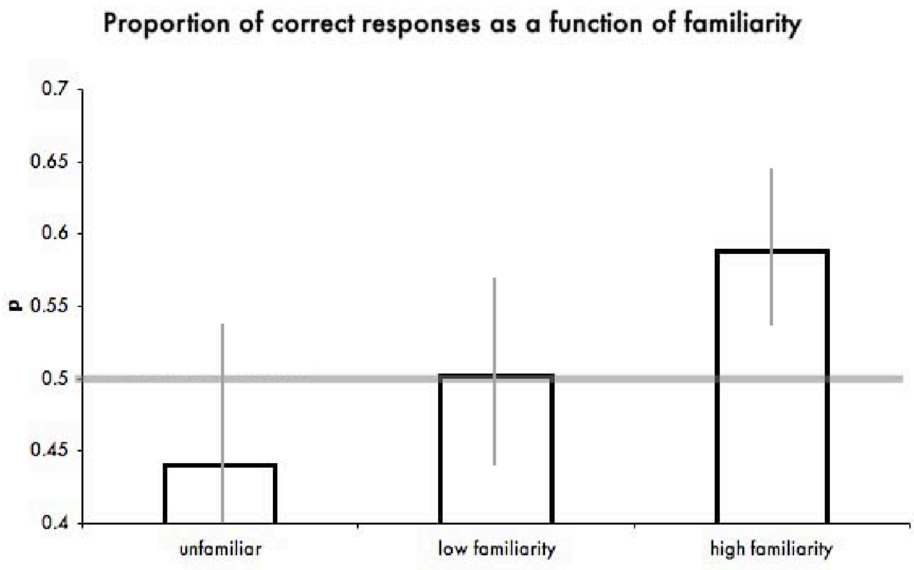

Figure 4.6: NB The horizontal grey line represents chance performance and error bars denote 95% confidence intervals.

Performance did not exceed chance level (.5) for both unfamiliar [t(19)=-1.410; p>0.05] and low familiarity groups [t(18)=0.081; p>0.05]. Performance did however differ significantly from chance for the high familiarity group [t(28)=3.157; p<0.001], suggesting that the ability to perform the MD task successfully is facilitated by increased familiarity. However, though the overall accuracy data for experimental groups would suggest that familiarity improves performance, the correlation between average performance on an image pair and the average familiarity rating was not significant [r=0.217; p=0.457].

Discussion

This experiment provides evidence that performance on the explicit mirror discrimination task is improved as exposure to a face accumulates, and data from the previous two experiments are also in support of this contention. In both previous experiments performance on a MD task was shown to be above chance only for faces that were familiar to the subjects. In addition, chance level performance was reported in experiment 9 for participants from the University of Glasgow psychology department. In discussing this finding the most likely explanation was thought to be that the participants in the Glasgow psychology department were less familiar with the Glasgow stimulus set that the Stirling participants were with the Stirling stimulus set. The fact that there was not a statistically reliable difference between the stimulus sets in terms of perceived asymmetry would suggest that the difference in performance was not due to some physical property of the stimuli.

This finding replicates Rhodes' (1986) previous demonstration of superior mirror discrimination performance for highly familiar faces, and is also in agreement with studies showing a relationship between familiarity and the ability to identify the correct left-right orientation of common objects (Kelly et al. 2001) and works of art (Blount et al. 1975). That increased familiarity is associated with improved performance on this task suggests that as a face becomes familiar the cognitive representation of its shape becomes sufficiently detailed to enable participants to discriminate between a mirror image and a normally oriented image of that face. For this to occur the cognitive representation must not only be finely tuned to the relational information in the familiar face but it must also retain information that specifies the left-right orientation of this configural information.

An interesting question arising naturally from the data reported in this experiment is whether performance on the MD task will improve proportionately with increasing exposure to a given face or whether there is some upper-limit that constrains performance. Though the correlation between average familiarity ratings and average accuracy on the MD task was not reliable, it is considered that this is reflective of the imprecise measure of familiarity used as opposed to a functional independence between familiarity and MD performance.

Chapter Summary

The data reported in this chapter provides further evidence that the left-right orientation of highly familiar faces is explicitly coded in the mind. In the previous two chapters the evidence supporting this contention remained equivocal, as invariance to mirror reversal had been demonstrated as often as sensitivity to this manipulation. However, given that it has now been shown that participants can explicitly identify the normal orientation of a familiar face there can be little doubt that our stored representations of faces are detailed enough to code for the subtle differences between the left and right sides of a face. In addition to this, the results reported in these three experiments demonstrate an ability to recall the left-right orientation of this information from memory.

Though these two conclusions can be asserted with greater conviction given the results reported in the present chapter, it remains unclear which property of the cognitive representations causes MD performance to be so poor. Given that the process of identification requires the representation of precise spatial information (e.g. Farkas & Munro, 1987) it is unlikely that our memory for the configuration of faces is not detailed enough to detect the alteration in visual appearance induced by mirror reversal. Rather it seems more plausible that low MD performance is reflective of some difficulty in recalling the orientation of this information. This explanation would concur with demonstrations of poor memory for the left-right orientation of common objects (e.g. Nickerson & Adams, 1979) and would also connote that the cognitive representation of faces is to some degree ‘object-centred’ (see Biedermann, 1987), or at least that such representations are not exclusively viewer-centred in their orientation.

Another explanation is that participants perform MD with face stimuli by somehow judging the relative familiarity of the two images. If one assumes that recognition involves perceived familiarity exceeding a certain threshold (e.g. Burton et al. 1990) then subjects may be basing their judgement on the relative familiarity of each stimulus. It has been demonstrated already that recognition accuracy is unaffected by mirror reversal (experiment 4) and therefore this method would involve making subtle introspective judgements regarding 'above-threshold' sensations of familiarity. This would be a quite unnatural process and one which participants may find difficult.

Though the performance reported in this chapter is poor, it would appear that familiarity improves the ability to identify correct mirror orientation, and this is in agreement with previous demonstrations of successful MD performance to images of familiar faces (Rhodes, 1986), common objects (Kelly et al. 2001) and works of art (Blount et al. 1975). Why familiarity might improve performance on this task remains uncertain however. It could be that exposure encourages the refinement of cognitive representations which results in a more precise mapping of the faces shape thus allowing for the correct orientation to be detected more readily. However it could also be that familiarity encourages the left-right orientation of this information to be more unambiguously specified in memory (i.e. with successive exposure the representation becomes increasingly viewer-centred). If these two processes are indeed separable then their relative contribution to the exposure-driven improvement in MD accuracy is considered to be a question worthy of further investigation.

A more general question arising from the data reported here is whether MD performance can provide a reliable index of familiarity. The results of experiment twelve suggest that above-chance performance is dependent on a sufficient level of exposure to the faces prior to testing however this investigation is somewhat incomplete. Since the participants in the 'high familiarity' group had relatively little exposure to the majority of the faces (some of the contestants had only been on the show for a matter of days) it is considered that this condition was not reflective of the highest level of familiarity attainable. It would be interesting to measure the MD performance of groups varying in frequency of exposure at incremental stages throughout the course of a television series. Using this method it would be possible to determine whether the rate of improvement at MD varies as a function of the frequency of exposure, whether this improvement is accumulative and whether an upper-limit of performance constrains this growth.

Chapter 5

Implicit memory for the left-right orientation of unfamiliar faces

Introduction

The two experiments in chapter 4 show that the probability of choosing the veridical orientation of an unfamiliar face is equivalent to chance. This finding is corroborated by previous research requiring participants to mirror discriminate images of unfamiliar faces (e.g. Rhodes, 1986) and unfamiliar works of art (Blount et al. 1975). Both of these studies show that familiarity with the stimuli allows for above-chance performance but in neither study does performance with unfamiliar faces exceed chance. This is taken as evidence that mirror discrimination is not achievable on the basis of extractable rules relating to asymmetry in general, such as the tendency in portraiture for subjects to be depicted with their left cheek showing to the viewer (McManus & Humphrey, 1973).

As well as improving mirror discrimination performance it has been demonstrated that familiarity promotes higher ratings of preference to objects (Zajonc, 1968) and faces (Mita et al., 1977). In what has become a classical research paper Zajonc (1968) demonstrated that objects to which we have been exposed (even if we are not consciously aware of this exposure) are more favourably judged than objects that we have not previously encountered. Although some recent research has suggested that this effect may not occur independently of recognition (e.g. Newell & Shanks, 2007), the prevailing account remains that positive affect can be induced by both conscious and non-conscious exposure (see Zajonc 2001).

In the present chapter, two experiments are reported that investigate whether exposure to directional asymmetry in human faces results in a preference towards the veridical left-right orientation of unfamiliar faces. This would be expected if, contrary to previous reports (Rhodes, 1986; McManus & Humphrey, 1973), extractable rules relating to the asymmetry of faces in general are cognitively represented. It has previously been shown that positive affect is associated with frequently encountered grammatical rules (Newell & Bright, 2001; Zizak & Reber, 2004). This is thought to be an example of the more general "structural" mere exposure effect that increases positive affect to novel stimuli that conform to implicitly acquired rule systems (e.g. Gordon & Holyoak, 1983). It is possible therefore that a similar preference will be observed towards unfamiliar faces shown in their veridical orientation when the face displays asymmetry conforming to an underlying rule system.

Facial asymmetry has been shown to conform to general rules of asymmetry that occur in all biological organisms. The term directional asymmetry (DA) is used to refer to any anatomical asymmetry present in an organism that results from non-random developmental and/or genetic processes (Klingenberg, 2003; Palmer & Strobeck, 2003). This is distinguished from fluctuating asymmetry (FA) which results from the organism being unable to resist environmental and/or genetic stressors during the course of development (e.g. Parsons, 1990a, 1990b; Polak, 2003), and as such has been implicated in attractiveness research emphasising the importance of perceived genetic health in mate preference (Møller & Thornhill, 1998; Thornhill & Møller, 1998). It is generally agreed that there is a mild degree of the former source of asymmetry (DA) present in human faces, with the right side of the face generally larger than the left (Farkas & Cheung, 1981; Peck et el. 1991; Sackeim, 1985; Simmons et al. 2004).

Simmons *et al* (2004) made measurements of 172 male faces and 205 female faces and reported that there was a general tendency for faces to be larger on the right-hand side. Interestingly, they also found that when participants were asked to make asymmetry or attractiveness judgements to these faces they appeared to do so independently of this general asymmetry. The degree of DA in the faces did not affect asymmetry or attractiveness judgements towards them but participants instead based their judgements on the asymmetry variance around DA. This finding shows that the cognitive system is sensitively tuned to perceive the subtle asymmetries in faces that are thought to reflect the underlying developmental instability (FA). In addition, they show that the cognitive system is capable of adapting to consistent qualities of facial structure, which is in line with previous research (Rhodes, 1996; Rhodes et al. 2003a, 2003b; Webster & MacLin, 1999). This study implies that population-wide asymmetries in facial structure are to some extent cognitively represented.

Previous research has demonstrated incidental learning of real world regularities for common objects (Kelly et al. 2001). In addition, it has been shown that participants can readily abstract prototypical representations after exposure to a series of visually similar objects (e.g. Posner & Keele, 1968). It is possible therefore that the cognitive representation of DA present in faces will result in a bias towards choosing the veridical orientation of an unfamiliar face in a 2AFC paradigm. The 'mirror discrimination' task used in the previous experiment failed to detect such a bias, however it is possible that this was because the measure lacked sensitivity. In the present chapter the 'likeability paradigm' (Zajonc, 1968) is used to determine whether our cognitive systems are sufficiently impressible to allow for the representation of DA.

Previous research has shown that the mere exposure effect (as outlined earlier) is capable of generalising to an average representation of previously exposed stimuli even if the prototype stimulus has not itself been encountered. This effect has been demonstrated for artificial grammar (Gordon & Holyoak, 1983; Manza et al. 1998; Zizak & Reber, 2004), matrices of coloured squares (Gordon & Holyoak, 1983) and human faces (Rhodes et al. 2001). An abstractive process of this kind would be necessary for the cognitive representation of DA, which has been previously demonstrated by Simmons *et al* (2004), and also for participants to show preference to veridically presented unfamiliar faces relative to their mirror reversal. However, in more recent investigation the same researchers failed to find evidence for an increase in likeability ratings towards composites of previously seen faces despite finding a preference towards the individual faces that had been previously encountered (Rhodes et al. 2005). This was taken as evidence that increased positive affect towards familiar faces does not generalise across individual faces to prototypical representations and therefore that the attractiveness of average faces cannot be explained in terms of a generalised mere exposure effect.

Although the mere exposure effect has rarely been demonstrated following naturalistic exposure (Cutting, 2003) the learning of incidental regularities in lateral orientation has previously been established (Kelly et al. 2001). Given that information relating to DA in faces is to some extent available to our cognitive systems (Simmons et al. 2004), and that the Mere Exposure Effect has been shown to generalise across stimuli (e.g. Gordon & Holyoak, 1983; Rhodes et al. 2001) the experiments in the present chapter were designed to test for behavioural sensitivity to the mirror reversal of unfamiliar faces. Using the paradigm designed by Zajonc (1968), veridically oriented images of unfamiliar faces were presented alongside a mirror reversed copy of the image and participants were required to decide which image they 'liked' more. If DA is cognitively represented then a bias towards preferring the originally oriented image should be detected, as the form of this image should (in general) be more familiar to participants relative to the mirror reversed copy.

A further reason to expect that veridically oriented stimuli would be more likely to stimulate preference is that prototypical faces, amongst other object classes (see Halberstadt & Rhodes, 2003) are rated as being more attractive (e.g. Rhodes & Tremewan, 1996; Rhodes et al. 1999). Given that the 'prototypical face' is asymmetrical (e.g. Farkas & Cheung, 1981; Farkas & Munro, 1987; Sackeim, 1985; Simmons et al. 2004) reversing the orientation of faces will (on the majority of occasions) make the resultant image less similar to this average representation. This should translate to a detectable bias in response behaviour, with participants perceiving the original orientation as more attractive. As this prediction is the same as is predicted by the generalised mere exposure hypothesis (see Zizak & Reber, 2004) it is anticipated that dissociating the underlying concepts of attraction and likeability will not be possible here.

Experiment 13: Implicit sensitivity to the mirror reversal of unfamiliar faces I

Introduction

Faces have been shown to display an increasing degree of asymmetry with age, and female faces are thought to undergo a more gradual change in morphology during the process of aging than male faces (Farkas, 1994). For this reason the stimuli used in this experiment were split into four groups according to the age and the sex of the person in the photograph. Were directional asymmetry to be represented in the mind then a response bias towards the veridical stimulus would most likely be detected for images of elderly male faces, where DA is at its most pronounced.

Method

Stimuli and Materials

Unfamiliar face images were taken from an existing database (Minear & Park, 2004). This database categorises faces by age, expression and sex. In the present study a sample of 44 faces was selected from each of the following categories: young female, young male, old female, and old male. In the "young" subset ages ranged from 18 to 29, the age-range in the "old" subset was 69 to 93. All faces used were of neutral expression.

Images were cropped and rescaled to a standard size (365 pixels horizontally), however no attempt was made to remove the external features (hair, ears, jewellery etc.). For every image in the experimental corpus, a mirror reversed (x axis) copy was generated. Example stimuli are shown in figure 5.1.

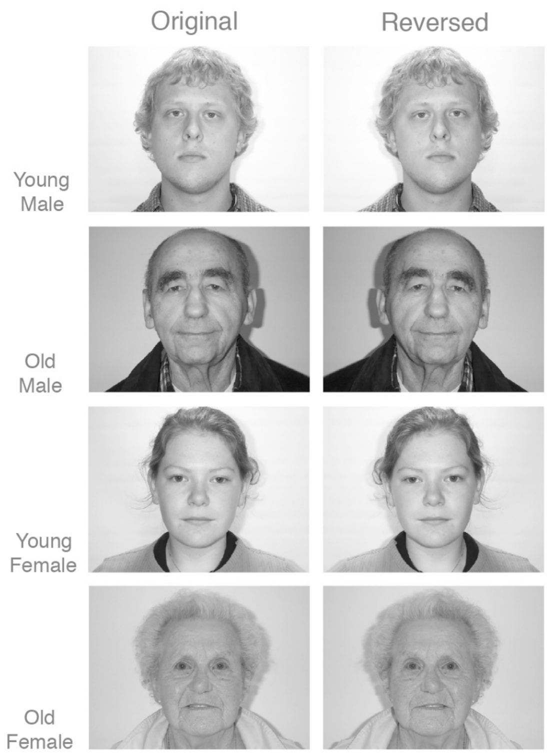

Figure 5.1: Example stimuli from experiment 13.

Design and Procedure

36 subjects from the undergraduate and postgraduate populations at the University of Glasgow agreed to participate in the study. The sample consisted of 21 female and 15 male subjects ranging from 19 to 32 years in age (mean 23).

Subjects were sat in front of a Macintosh workstation running the experimental software package Psyscope X. For each trial, one face appeared in both its original and reversed form and the presentation was preceded by a fixation cross in the centre of the screen (750ms). The images were presented at either side of the screen and the original image appeared on each side an equal amount of times in each condition. The images remained on screen until the subject made a response.

Subjects were told that they would be presented (in each of the 176 trials) with two images of the same person, one at either side of the screen. They were told that they must indicate which of these two images they "liked" better. Furthermore, they were instructed not to take too long to make these decisions and to rely on their initial instinct. Two sets of response keys corresponded to the two sides of the screen. Three keys corresponded to the left side, and three to the right. These three keys allowed the subject to indicate the strength of their preference to the respective image. The red key indicated strong preference, blue medium preference and green weak preference. The positioning of these keys was symmetrical with respect to the centre of the screen (fixation). Responses and response latencies were recorded for subsequent analysis.

Results

Mean proportions of responses to the veridically oriented stimulus are shown in the graphs below, with chance level performance being indicated by the grey line. One-sample t-tests were carried out to determine whether experimental conditions varied from chance levels (mean population value assumed to be $p = .5$). Using two-tailed significance testing, a bias towards the original (non-reversed) stimulus was detected, but only in the old-male stimulus condition, and only for the strongest confidence level [$t(1,35)=2.357$; $p<0.05$]. No other conditions reached statistical significance.

Thus, of the twelve t-tests carried out on the data only one was found to vary significantly from chance with a probability of less than 5%.

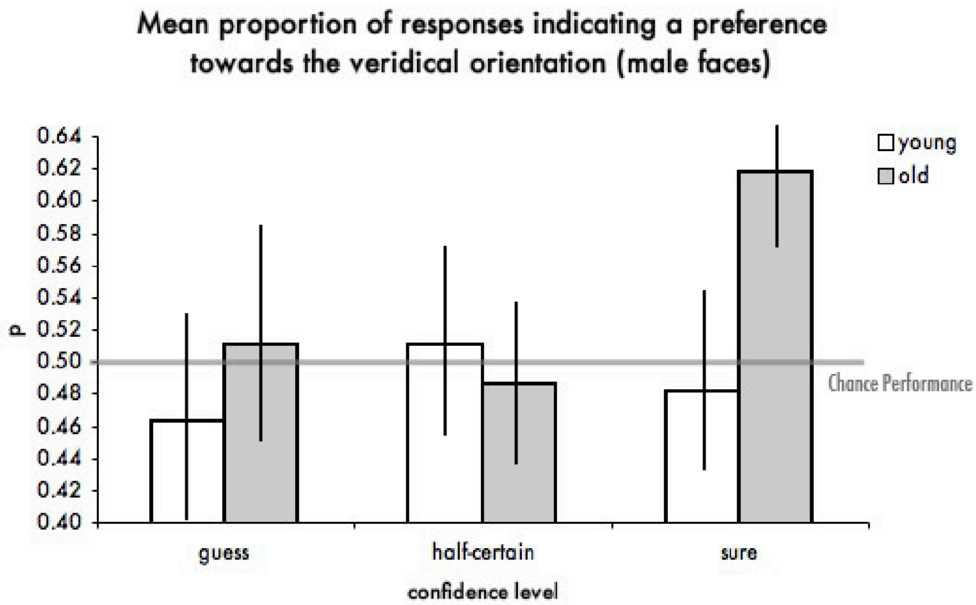

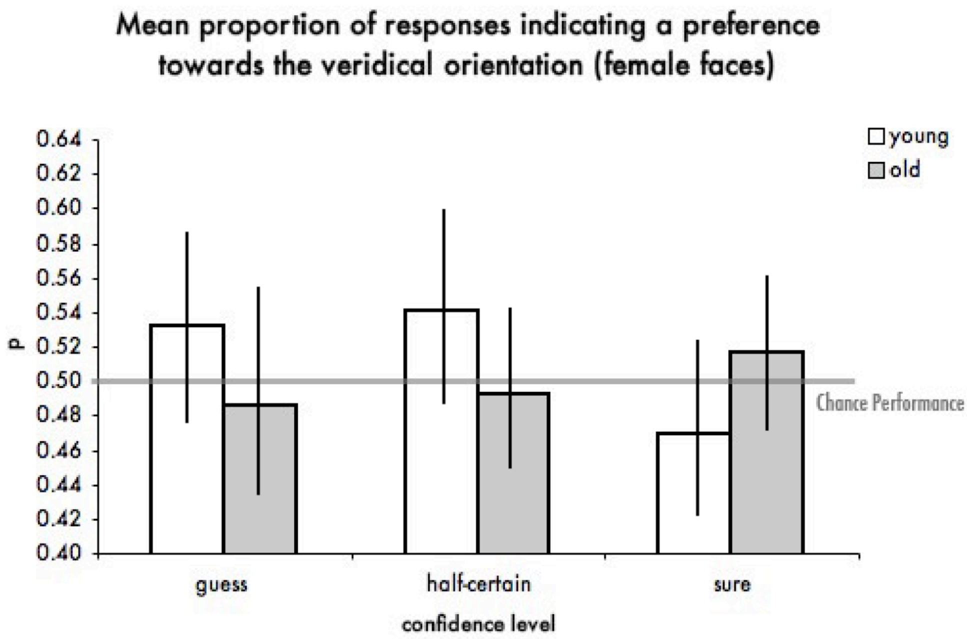

Figures 5.3 & 5.4: Bar graphs showing the proportion of responses made in preference to the veridically oriented stimulus for male (fig. 5.2) and female (fig. 5.3) faces. **NB** Error bars denote 95% confidence intervals.

It was considered that a two-tailed test was most suitable in this case, because it was possible that some stimulus factors may in fact bias response in the opposite direction to that which was hypothesised (for instance a systematic asymmetry in lighting direction was present in the stimulus set). However even when using one-tailed tests of significance the only result to exceed chance performance was for ‘sure’ responses towards elderly male faces [$t(1,35)=2.357$; $p<0.01$].

Discussion

The present experiment has provided very limited evidence for the cognitive representation of Directional Asymmetry in faces. In the old male face condition a preference towards the veridical orientation of the stimuli was detected in responses made with the highest level of confidence. Although it was hypothesised that this group of stimuli would display the highest degree of directional asymmetry and therefore that participants would be most likely to prefer the veridical orientation in this condition, the possibility of this data being anomalous cannot be overlooked. Given that twelve t-tests were carried out in this analysis the likelihood of one of these tests being significant due to chance is quite high and for this reason the null hypothesis in this condition should be rejected with the necessary caution. Furthermore it could be that the higher incidence of asymmetric hairstyles (side-partings) in the images of elderly male faces contributed to this observed bias.

That said, it was hypothesised that because male faces are typically subject to a more profound change in morphology during the course of aging, directional asymmetry would be more pronounced in this stimulus set. Any effect of generalised familiarity for veridical lateral orientation relative to mirror reversals would presumably be enhanced by high levels of DA and therefore a preference towards the non-reversed image would be most likely to occur in this condition. If this finding were to be replicated using another class of stimuli displaying high levels of DA then it would be possible to reject the null hypothesis with greater certainty.

Experiment 14: Implicit sensitivity to the mirror reversal of unfamiliar faces II

Introduction

In addition to DA in facial morphology, faces become additionally asymmetric during the expression of emotion: a phenomenon that results from the lateralisation of brain function responsible for the production of facial expression (e.g. Borod et al. 1997). For this reason a preference towards veridically oriented stimuli is more likely to be detected in images of faces that are expressing emotion relative to images of non-expressive faces. Should such a bias be detected it could be concluded that the significant bias towards normally oriented images of old male faces reported in the previous experiment was not simply a statistical anomaly, but instead is reflective of a generalised Mere Exposure Effect to DA in faces. If Mere Exposure Effects can generalise to consistent patterns in the appearance of faces then a preference towards normally oriented images of smiling faces should be detected, as their form in general will be more familiar to participants than their mirror reversals.

Method

Stimuli and Materials

200 unfamiliar face images were selected from the database of images used to compile the stimulus set used in experiment one (Minear & Park, 2004). These images comprised of 50 female and 50 male identities, with each identity being pictured displaying both neutral and smiling expressions. Images were cropped and rescaled to a standard size (365 pixels horizontally), however no attempt was made to remove the external features (hair, ears, jewellery etc.). For every image in the experimental corpus, a mirror reversed (x axis) copy was generated.

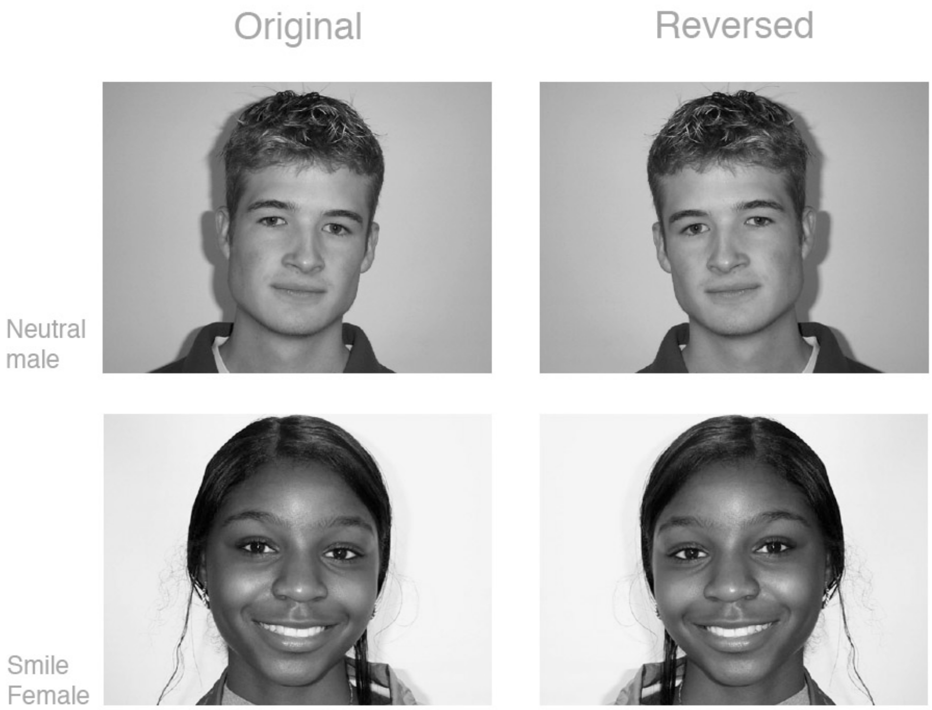

Figure 5.5: Example stimuli from two of the four experimental conditions.

Design and Procedure

30 subjects from the undergraduate and postgraduate populations at the University of Glasgow agreed to participate in the study. The sample consisted of 19 female and 11 male subjects ranging from 19 to 32 years in age (mean 22). The methodology used was identical to the previous experiment, however with 50 presentations per condition there were slightly more stimulus presentations (200 trials).

Results

Mean proportions of responses to the veridically oriented stimulus are shown in the graphs below, with chance level performance being indicated by the grey line. Two tailed one-sample t-tests were used to analyse whether each of the experimental conditions differed (at each confidence level) from chance performance (population mean= 0.5). Performance did not differ significantly from chance levels in any of the experimental conditions.

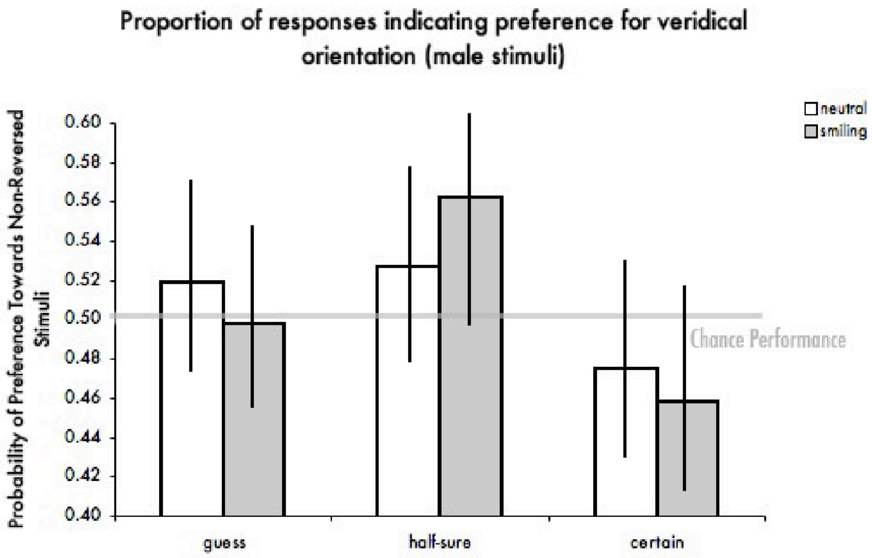

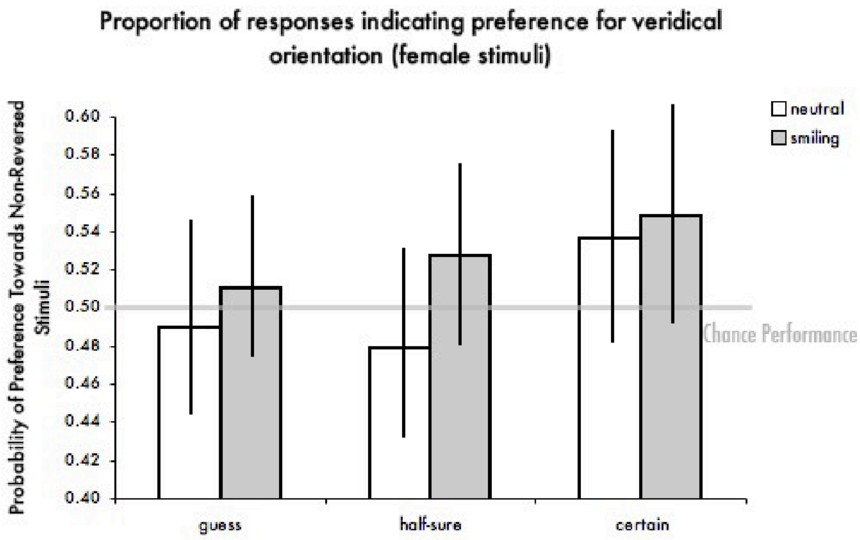

Figures 5.6 & 5.7: Bar graphs showing the proportion of responses made in preference to the veridically oriented stimulus for male (fig. 5.2) and female (fig. 5.3) faces. **NB** Error bars denote 95% confidence intervals.

Discussion

Given the null effects reported in this experiment it would appear unlikely that Mere Exposure Effects are suitably robust to generalise to population-wide patterns in facial asymmetry. In addition, failure to replicate the above-chance bias reported for old male faces in experiment twelve would suggest that a reappraisal of this result is necessary. If the significance of this bias were indeed due to the relatively high degree of DA in the stimulus condition then it would be predicted that a 'veridical orientation bias' would also be detected when participants made likeability judgements towards mirror reversals of expressive faces. This was not found to be the case and therefore it would appear unlikely that this previous finding was reflective of cognitive sensitivity to DA.

Further to the structural asymmetry present in faces and the underlying rules that this structure conforms to, there may have been other sources of systematic variation in the asymmetry of the experimental stimuli. As no attempt was made to remove hairstyle, jewellery or other such stimulus artefacts it might have been expected that sensitivity to regularities in (for instance) the ear in which earrings are most commonly worn may have induced a bias towards the normal orientation. People most commonly wear their wristwatches on their left wrist, and a similar pattern may also be true for jewellery. Additionally, patterns in spontaneous head angle (see McManus et al. 2004) could also have been expected to induce bias towards the original orientation. However no bias was detected in any of the stimulus conditions suggesting that these regularities (in addition to structural regularities) did not contribute to a generalised mere exposure effect.

Chapter Summary

Given that the prototypical face is asymmetrical (Farkas & Cheung, 1981; Peck et el. 1991; Sackheim, 1985; Simmons et al. 2004) and that prototypical faces are rated as being more attractive (e.g. Rhodes & Tremewan, 1996; Rhodes et al. 1999) it was expected that a preference towards the veridical orientation of unfamiliar faces would be detected using the methodology described in the present chapter. This hypothesis was not supported by the empirical data reported here. Although a preference was detected for the veridical orientation of elderly male faces, this result is reported in the context of 23 other non-significant statistical tests, suggesting that this was most probably a type one error.

Given the apparent sensitivity of the experimental paradigm at detecting changes in affective response resulting from exposure to visual form (Zajonc, 2001) and the reported ability for such sensitivity to extend to prototypical examples of stimulus classes (Gordon & Holyoak, 1983; Rhodes et al. 2001) it is perhaps surprising that a veridical orientation bias was not detected. After all, previous research has shown that DA is cognitively represented (Simmons et al. 2004).

One explanation for this discrepancy is simply that Mere Exposure Effects are very difficult to detect following naturalistic exposure. As only one example of this has been reported in the literature (Cutting, 2003) it is probable that the classical result is more difficult to observe when the experimental procedure does not include a controlled exposure period. It is considered that by employing a methodology in which exposure is experimentally induced the likelihood of observing the hypothesised effect would be increased. Such an experimental method would require that participants were exposed to stimuli under controlled conditions so that the effect of exposure to DA on subsequent preference decisions (between veridical and mirror oriented faces) could be systematically explored. Exposure to varying levels of DA (i.e. the degree to which a face conforms to the structural norms, or doesn't) may differentially effect subsequent preference towards normally oriented and reversed stimuli and using the classical MEE paradigm would enable such an investigation.

The question of whether a preference for the veridical orientation can be induced by exposure to varying proportions of DA in faces is however a deviation from the question posed in the present chapter. The question posed here was whether or not exposure to consistencies in the structural asymmetry of faces in everyday lives leads to the cognitive representation of these regularities. Although a previous study suggested that attractiveness judgements are made independently of DA (thus providing good evidence for the implicit learning and subsequent cognitive representation of the rules underlying structural asymmetry in faces, Simmons et al. 2004), this remains the only demonstration of cognitive representation of this kind. For this reason further demonstrations of additional psychological consequences are necessary to support this hypothesis.

Chapter 6

Perceptual asymmetry in face processing

Introduction

To this point, the experiments reported have investigated the manner in which physical asymmetry in faces is psychologically represented. However, there is another factor in addition to the morphological asymmetry of the human face that may contribute to asymmetrical cognitive representations of faces. It has been repeatedly shown that the perceptual importance of the two sides of the face is asymmetrically distributed (e.g. Wolff, 1933; Heller & Levy, 1981). It has further been shown, in free viewing conditions, that the left side of the face is more commonly attended to (Yarbus, 1967; Butler et al. 2005), and information to the left of the face has been said to be more 'diagnostic' of a faces identity (Vinette et al. 2004).

In studies using 'chimeric faces' it is typically found that chimeras made using the half of the face to the viewers left are judged as being more 'like' the person than are the right chimeric faces (e.g. Wolff, 1933). In addition, these images are thought to express more emotion (Heller & Levy, 1981; David, 1993; Ferber & Murray, 2005), and are more readily classified by gender (Luh et al. 1991; Butler et al. 2005). Although there is some debate as to the universality of this bias across cultures (Gilbert & Balkan, 1973; Vaid & Singh, 1989), the overwhelming consensus is that face perception is prone to a powerful left perceptual bias, and this phenomenon is generally ascribed to a right hemisphere bias for face processing (De Renzi et al. 1994; Dutta & Mandal, 2002; Hugdahl et al. 1993).

Importantly, this perceptual bias exists independently of morphological asymmetry and contributes to asymmetrical mental representations (Brady et al. 2005). In the context of the questions posed so far in this thesis, the existence of this perceptual asymmetry should be considered an important point of discussion. I have demonstrated (under certain conditions) that participants are commonly sensitive to the mirror reversal of familiar faces, and have so far argued that this sensitivity stems from the representation of asymmetrical facial morphology. However, given that faces are perceived asymmetrically it is likely that this bias interacts with the physical asymmetry of a face to produce asymmetrical cognitive representations.

In this final empirical chapter two experiments are reported that demonstrate an increased reliance on the left eye[1] relative to the right eye in a face recognition procedure. These experiments are designed essentially as a replication of a previous study investigating the manner in which faces are cognitively represented (Cooper & Wojan, 2000). In this paper, the researchers addressed the question of whether faces are represented using 'coordinate' or 'categorical' relations. In order to address this question they manipulated images of famous people by repositioning either one eye or both eyes and recorded the extent to which these alterations affected recognition (i.e. coordinate relations are more disrupted by moving both eyes, but relational is less compromised by this manipulation). Cooper and Wojan (2000) found that face recognition was detrimentally affected by moving both eyes but not one eye in isolation (see figure 6.1) and offered this finding as support for a 'coordinate system' of representation for face recognition.

Figure 6.1: Reproduction of stimulus conditions as described by Cooper & Wojan (2000).

In this chapter I demonstrate that the experimental design employed by Cooper and Wojan (2000), and the analysis they report, does not provide a valid test of their hypothesis. The analysis overlooks existing literature demonstrating perceptual asymmetry in face perception. Given that a perceptual bias towards the left side of the viewed face has been repeatedly demonstrated, it is likely that the null effect of moving one eye reported in their research disguises an asymmetrical pattern in their 'one-eye move' condition. Data for left-eye displacement and right-eye displacement

[1] Henceforth any reference to the 'left eye' refers to the eye appearing in the viewers left visual field

are not reported separately, but instead both contribute to the overall response measures associated with the 'one eye' move condition. This is considered to be a flaw in their experimental design and (in light of the existing research noted previously) one would expect that these two manipulations would differentially affect recognition. The two experiments reported in this chapter are designed to investigate whether the effect of displacing (experiment 14) and masking (experiment 15) the eye appearing in the viewers left visual field is more detrimental to face recognition than identical manipulations made to the eye displayed to the viewers right.

Experiment 15: The effects of configural manipulations of facial features on face recognition: Left eye Vs Right eye

Introduction

The experiment reported here replicates the 'one-eye move' condition detailed by Cooper and Wojan (2000) however the procedure used here is slightly different. In the present experiment an alternative 2AFC name-verification paradigm is employed, one which has become established within our research group as a reliable way in which to test the congruency between internal representations and presented stimuli (see Burton et al. 2005). Whereas the presentation of the face precedes the presentation of the name (test stimulus) in the task favoured by Cooper and Wojan, the order of presentation is reversed in the present study, with 'match/mismatch' responses being made to the face images. Given that a robust left-bias has been established in the face perception literature it is expected that face recognition performance will be more detrimentally effected by moving the left eye of the test images (i.e. the eye on the left of the presented stimulus from the viewers perspective) relative to the effect of an identical manipulation made to the right eye.

Method

Stimuli and materials

Sixty high-resolution images of famous celebrities were sourced from the World Wide Web. Only images where the entire face was visible in a full-face portrait were used. Images were subject to a process of standardisation. Firstly the pupils were brought into alignment with the horizontal. The image was then cropped so as to leave only the head and the area directly surrounding the head. Finally the image was converted to greyscale and resized to 190 X 285 pixels. The twenty images allocated to the 'original' condition were to receive no further treatment. Images allocated to the 'left eye move' condition were to have the eye appearing to the left of the image moved upwards by 1/20 of the total height of the image. Images assigned to the 'right eye move' condition were to have an identical manipulation performed on the eye appearing in the right side of the stimulus. Counterbalancing ensured that, across subjects, the famous identities appeared an equal amount of times in each stimulus condition.

The movement of the eye was achieved using the Adobe Photoshop CS software package. For each eye manipulation, an area around the eye was selected so as to include the eye, the eyebrow, and any wrinkles or shadows cast by the eye socket. This selection was then moved so that the centre of the pupil was moved upwards by a distance equivalent to 5% of the image, and so that its horizontal coordinate remained constant. The blank area left by this manipulation was filled in using the 'clone stamp' tool, which has the ability to copy areas of the image and is useful for copying surfaces with gradual tonal variation such as skin. Examples of the three stimulus conditions are shown in figure 6.2.

Figure 6.2: Examples of the three stimulus conditions in experiment 15.

Design and Procedure

32 participants from the undergraduate and postgraduate populations at the University of Glasgow took part in the study, and were either paid in cash or received course credit. A within subjects design was used. The experiment consisted of 120 trials, with twenty positive and twenty negative trials in each of the three conditions. The experiment lasted a total of five minutes. Counterbalancing ensured that, across subjects, the famous identities appeared an equal amount of times in each stimulus condition.

Each trial consisted of the presentation of a fixation cross for 750ms, this was followed by the presentation of a famous name for 1000ms, and then the presentation of a famous face (belonging to one of the three conditions) for 1000ms. Participants were asked to respond as quickly as possible after the onset of presentation of the face whether it was the person named previously, and the presentation terminated on depression of the response keys. For positive trials the name matched the face, and for negative trials the name and face were mismatched with a random identity of the same sex (taken from the experimental stimulus set). Responses and response times were recorded, and counterbalancing ensured that there was no effect of lateralisation of response (left key press/ right key press).

Stimuli were presented on a black background. Names were presented in bold, white, 32-point typeface. Face images measured 190 X 285 pixels and were centrally presented on a monitor display measuring 1152 by 864 pixels. Each participant was instructed to maintain fixation in the middle of the screen, to make their responses as quickly and accurately as possible and warned that the faces had been distorted in various ways. Participants were encouraged not to focus on these distortions but instead to concentrate on recognising the face as quickly as possible.

Results

Error rates are calculated from positive trials only (i.e. when the face matched the preceding name). Please note that "right eye moved" refers to the eye appearing on the right of the viewed image.

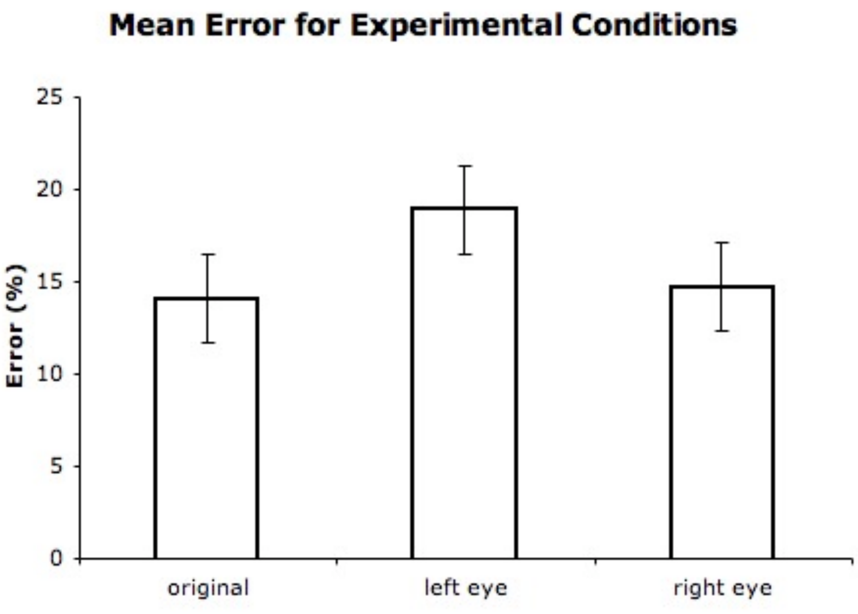

Figure 6.3: NB Error bars denote 95% confidence intervals which were calculated using the method described by Loftus & Masson (1994).

A one factor ANOVA was performed on subjects mean error rates for the three experimental conditions. The main effect of moving the eye was found to be of statistical significance [$F(1,31)=4.823$; $p<0.05$]. Subsequent post-hoc analysis by way of a Tukey's HSD test showed significant differences between left eye move and right eye move ($p<0.05$) and between left eye move and original images ($p<0.05$). The difference between right eye move and original conditions was not found to be significant.

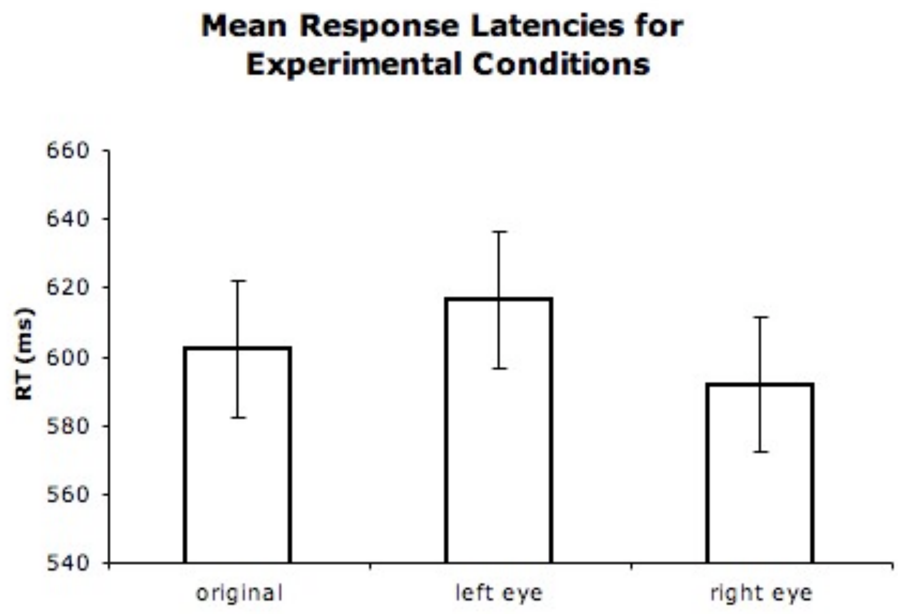

Figure 6.4: NB Error bars denote 95% confidence intervals which were calculated using the method described by Loftus & Masson (1994).

A one factor ANOVA showed no significant main effect of eye movement on response latency [$F(1,32)=2.215$; $p=.118$]. Post-hoc Tukey's HSD tests showed no significant effects between the three levels of eye movement.

Discussion

As predicted, the experimental manipulations of moving the left and right eyes of the face stimuli had different effects on recognition performance. Whereas moving the left eye was detrimental to recognition, moving the right eye had no effect on participants' response accuracy. In light of this demonstration, Cooper and Wojan's decision to report data from their 'one-eye' move condition as an aggregation of these two conditions would appear unjustified. In addition, this data provides a further demonstration of perceptual asymmetry in face processing.

The data reported here show that moving the right eye to a new vertical position has no effect on recognition performance. This finding is somewhat surprising as the stimulus produced by this manipulation is plainly altered in both first and second order configuration (see Maurer et al. 2002). However, similar results have previously been reported in this thesis. Unfamiliar face matching is performed equally well whether half faces or whole faces are displayed (experiment 1) and showing half faces results in analogous recognition accuracy when compared to whole-face presentation. These results are consistent with the notion that face recognition proceeds with equal success when participants are presented with just half of a familiar face.

It is not possible to argue that the processing of faces is entirely unhindered by removing or distorting one side of the face however. For instance, in experiment five it was demonstrated that removing either the left or right sides of the face increases response latency in a name verification decision, and in the present experiment recognition accuracy was detrimentally affected by moving the left eye. Therefore, the integrity of bilaterally distributed information in familiar faces contributes to optimal conditions for recognition.

There exists a clear inconsistency between the results reported here and the findings of experiment five however. Whereas there was not found to be any asymmetry in the effect of removing information to the left and right of the meridian in experiment five, in the present experiment moving the left eye worsened performance, but moving the right eye did not. The apparent incongruence of these two findings appears to suggest that the asymmetrical effects reported in this experiment are a result of the featural configuration being altered, as the same effect is not produced when half of the face is removed entirely. It is possible for instance that the left eye plays an important role in describing the configural relationships between the facial features, but that the right eye is not as heavily implicated in this process. As removing visual information does not change configural relations between the perceived facial elements, this might explain why no behavioural asymmetry was reported in experiment five.

Experiment 16: The effects of masking the left eye and the right eye on recognition performance

Introduction

This experiment was aimed at replicating the pattern of results reported in the previous experiment by using a different stimulus manipulation. Whereas the eyes were moved in the previous experiment to produce a different configuration of facial features, in this experiment the eyes were instead occluded. Occlusion does not alter the relationships between the features and therefore it is hoped that this will provide a test of whether the asymmetric effects reported in the previous experiment are restricted to positional alteration or whether similar effects are observed when the visual information is simply removed.

Method

Stimuli and materials

The same sixty high-resolution images of famous celebrities that were used in the previous experiment were used here, and these were subjected to the process of standardisation reported previously. The twenty images allocated to the 'original' condition received no further treatment. Images allocated to the 'left eye mask' condition had the eye appearing to the left of the image masked by a solid black rectangle. The rectangle varied slightly in size from image to image, as it was sized so as to cover the area of the face containing the eye, the area directly around the eye, and the eyebrow. Images assigned to the 'right eye erase' condition were subjected to an identical manipulation performed on the eye appearing in the right side of the stimulus. Examples of the stimulus conditions are provided in Figure 6.5.

Design and procedure

24 participants from the undergraduate and postgraduate populations at the University of Glasgow took part in the study, and were either paid in cash or received course credit. A within subjects design was used. The experiment consisted of 120 trials, with twenty positive and twenty negative trials in each of the three conditions. The experiment lasted a total of five minutes. Counterbalancing ensured that, across subjects, the famous identities appeared an equal amount of times in each stimulus condition.

Each trial consisted of the presentation of a fixation cross for 750ms, this was followed by the presentation of a famous name for 1000ms, and then the presentation of a famous face (belonging to one of the three conditions) for 1000ms. Participants were asked to respond as quickly as possible after the onset of presentation of the face as to whether it was the person named previously, and the presentation terminated on depression of the response keys.

Figure 6.5: Examples of the three stimulus conditions in experiment 16.

For positive trials the name matched the face, and for negative trials the name and face were mismatched with a random identity of the same sex (taken from the experimental stimulus set). Responses and response times were recorded, and counterbalancing ensured that there was no effect of response lateralisation.

Stimuli were presented on a black background. Names were presented in bold, white, 32-point. Face images measured 190 X 285 pixels and were presented centrally on a monitor display measuring 1152 by 864 pixels. Each participant was instructed to maintain fixation in the middle of the screen, to make their responses as quickly and accurately as possible and warned that the faces had been distorted in various ways. Participants were encouraged not to focus on these distortions but instead to concentrate on recognising the face as quickly as possible.

Results

Error rates are calculated from positive trials only (i.e. when the face matched the preceding name). Please note that "right eye moved" refers to the eye appearing on the right of the viewed image.

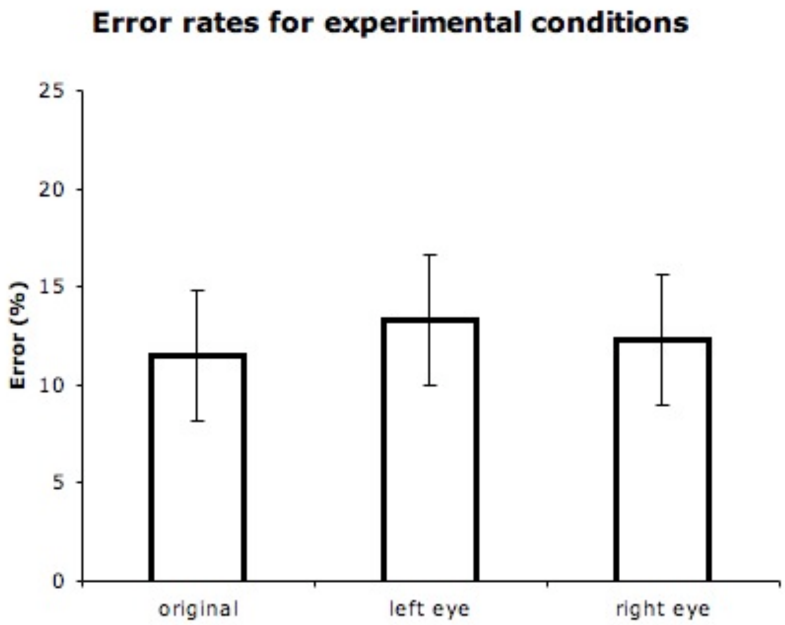

Figure 6.6: NB Error bars denote 95% confidence intervals which were calculated using the method described by Loftus & Masson (1994).

A one factor ANOVA was performed on subjects mean error rates for the three experimental conditions. The main effect of masking the eye was found to be unreliable [$F(1,23)=0.320$; $p>0.05$].

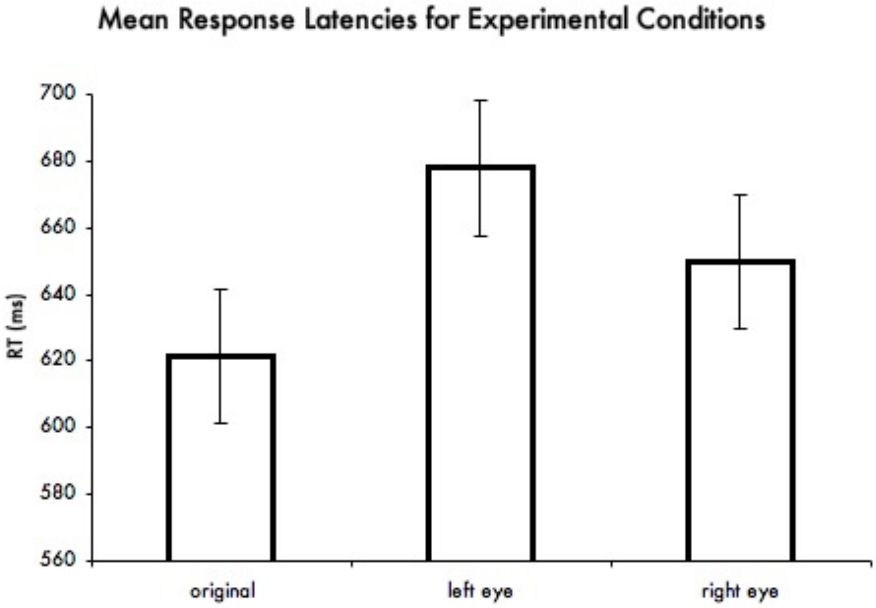

Figure 6.7: NB Error bars denote 95% confidence intervals which were calculated using the method described by Loftus & Masson (1994).

Subjects median response latencies for the three experimental conditions were subjected to a one factor ANOVA, which showed a reliable main effect of eye masking [$F(1,23)=7.942$; $p<0.01$]. A Tukey post-hoc analysis showed a significant difference between the left eye mask and the original image conditions ($p<0.001$), but no other reliable differences were detected.

Discussion

In the experiment reported above there was found to be no difference between accuracy performance after removing left and right eye information from familiar faces. The effect of removing the left eye produced a significant increase in reaction time however, with responses being significantly slower in this condition relative to when no information was removed. The difference of 28ms between this manipulation and an identical treatment to the opposite side of the image was not significant however. This is in contrast to the previous experiment, where recognition was shown to be significantly more error-prone towards stimuli where the left eye had been moved relative to images in which the right eye had been displaced.

It is considered that this difference can be explained by the fact that featural configuration was not altered in the present experiment. If one assumes that the first feature participants commonly look at in the face is the left eye (e.g. Gallois et al. 1989), then it can be reasonably anticipated that this initial orientation would produce different perceptual effects in the two experiments reported in this chapter. In experiment 15, when the left eye has been moved, our perception of the configural relations becomes grossly altered (i.e. the absolute distance and angular relationships between the left eye and every other feature are altered), and recognition is therefore detrimentally affected. However, when we look at this eye and it is the right eye that has been moved, only the inter-ocular relationship is compromised (relative to the first fixation) and therefore recognition is largely unaffected.

In the present experiment our tendency is also to look at the left eye first and to judge the face relative to this reference point, however as no visual information exists in this location our attention is merely shifted towards a more informative region, but our perception of the face remains uncompromised. There exists no misleading information in this location, just no information at all, and therefore this manipulation is less disruptive towards recognition. This reasoning explains why the only manipulation (in experiments 4, 15 and 16) that had a detrimental effect on accuracy was relocation of the left eye in experiment 15, and why reaction time data in the present experiment is suggestive of a tendency for left eye occlusion to be more disruptive than right eye occlusion (i.e. this reflects the time taken to shift attention towards more informative regions).

Chapter Summary

It has previously been demonstrated that memory representations of faces are asymmetric with respect to the prominence of the facial hemispheres (Brady et al. 2005), and one might expect this given that encoding is influenced by a robust leftward bias. It is not possible to confirm however, based on the data reported in this chapter, whether the asymmetric effect of moving the left and right eyes is reflective of an asymmetric representation of leftward information in memory, or whether it merely shows that when viewing a face we initially look to the left. Instead, it is considered that the data reported here demonstrate an interaction between a leftward perceptual bias characteristic of face processing and the manner in which familiar faces are represented in memory.

It has been repeatedly demonstrated that faces are processed in a manner that affords particular salience to the configuration of facial features (e.g. Yin, 1969). Furthermore, it has been demonstrated that manipulating these spatial relationships can disrupt face recognition and that we are extremely sensitive to such alterations (Haig, 1984), especially when the internal features are displaced (Hosie et al. 1988). In this chapter I have demonstrated that the disruption caused by manipulating the spatial relationships of familiar faces varies across the left and right sides of the face. That the effect of spatial displacement does not conform to the bilateral symmetry common to the structure of faces suggests that there exists an asymmetry in the manner in which faces are perceived by the visual system, and this finding is in agreement with previous research (e.g. Heller & Levy, 1981; Wolff, 1933). It is commonly accepted that this asymmetry is most likely to be caused by the asymmetric distribution of the face processing across the cerebral hemispheres (e.g. Vinette et al. 2004; Butler et al. 2005).

It is of interest however that this perceptual asymmetry was only apparent in one of the two experiments reported in this chapter. Occlusion, it would appear, does not affect the accuracy with which we identify familiar faces. Displacement of facial features does however cause poorer recognition performance, but it would appear that certain displacements are more disruptive than others. The stimulus manipulations used in experiment 15 were chosen because they could be similarly applied to the right and left sides of the face without compromising the appearance of any one feature. Any detrimental effect can therefore be attributed to the resultant alteration of spatial relationships between facial features. As feature displacement was only detrimental to performance when the left eye was moved, it is argued that the left eye is more heavily implicated in configural processing than the right eye.

In discussing the results of experiment 15 and 16 it was proposed that this asymmetry could be explained by the fact that the left eye is more commonly fixated than the right (e.g. Yarbus, 1967). However it is not certain that this effect relies on eye-movements *per se*. On the basis of the data reported here it is equally plausible that asymmetry in the memory representations of configural relationships is responsible for the observed effects. It has been shown that the leftward perceptual bias demonstrated in the perception of chimeric faces exists even in the absence of eye movements (Butler & Harvey, 2005; David, 1993), and it is possible that the perceptual asymmetry reported here may occur independently of eye movements also.

This is considered to be an interesting research question that could be addressed quite simply. If cognitive representations of faces contain more information describing the left side of the face than the right side, then it would be expected that an asymmetric effect would still be observed at presentation times of less than 100ms. This question could be pursued in future research, however the data are not reported here. It is likely that asymmetry would still be observed, as asymmetric representations are presumably generated and maintained by the asymmetric scanning strategy characteristic of face perception.

Given that the data reported here is suggestive of a memory representation (or encoding strategy) that promotes an asymmetric reliance on the left eye relative to the right, it would appear that cognitive representations of faces are necessarily asymmetric. Furthermore, it would appear that the asymmetry in these representations is a product of both physical and perceptual asymmetries. As perceptual asymmetries proceed independently of physical asymmetry however (Brady et al. 2005), it is unlikely that this contributes towards the sensitivity to reversal reported in chapters 2 and 3. Rather, it would appear that perceptual asymmetry results from an automatic bias that is driven by underlying functional asymmetry and therefore should not influence the detection (implicit or otherwise) of a mirror reversed face.

Chapter 7

General Discussion

The first experiment reported in this thesis demonstrated that when participants are presented with two opposite halves of an unfamiliar face they commonly report them as belonging to two different people. Having established that the asymmetry in a face is sufficient in degree to produce such an effect, the remainder of the thesis (excluding chapter 6) concentrated on trying to detect cognitive effects produced by mirror reversing faces. It was considered that this approach would allow conclusions to be drawn regarding the importance of facial asymmetry in familiar face recognition. However, the results reported in the experimental chapters indicate that any cognitive effects are largely undetectable and that when such effects are detected they are commonly slight.

Invariance to mirror reversal is commonly thought to facilitate object recognition despite changes in viewpoint (Vetter et al. 1994) and this may also be true for face recognition (Troje & Bülthoff, 1998). Faces can be recognised after just one presentation even when the face is subjected to large changes in orientation (e.g. Troje & Bülthoff, 1996) and therefore recognition systems would appear to be relatively robust to changes in viewpoint. Troje and Bülthoff (1998) present data that suggests a form of 'cognitive mirror reversal' supports this ability. Recognition accuracy of three-quarter views of faces was found to be superior when a mirror reversal of this view was presented at test compared with when the opposite three-quarter angle was shown (see figure 7.1). This data suggests that mirror approximations of three-quarter views are generated by perceptual systems and that identity judgements are made based on these representations when a face is originally encountered in the opposite angle.

Although recognition performance for mirror reversed images was worse than when tested with the original orientation of study images, this difference was slight, and in one of the experiments Troje and Bülthoff (1998) report this difference to be non-significant. Thus, the implications of their findings are twofold. Firstly, this study shows that generalisation across viewpoint (for a given image of an unfamiliar face) is more dependent on an image based transformation than any sophisticated abstraction of this information. Furthermore, it demonstrates that this image-based transformation occurs without significant loss of information and without any great degree of cognitive effort.

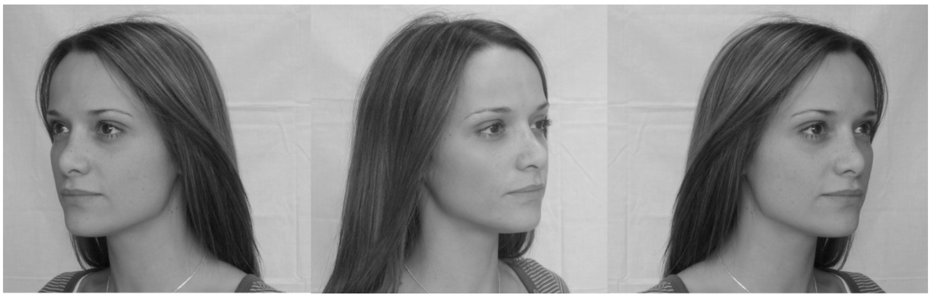

Figure 7.1: An illustration of Troje & Bülthoff's (1998) finding. The study face (left) was recognised more reliably when its mirror image (right) was presented at test, compared with when the opposite three quarter view was presented (centre). NB Troje & Bülthoff used computer-generated stimuli in their experiment.

This result is reflective of data reported in this thesis. It would appear likely from the data reported here and from previous research (e.g. Cooper et al. 2002), that face recognition is unaffected by mirror reversal. Although this invariance has previously been demonstrated in object recognition (Biedermann & Cooper, 1992), it was anticipated that this same result would not be found when using faces as stimuli. It was considered that memory for the visual form of faces would be represented in a manner which retained the structural asymmetry of a face as this would be necessary for the perceptual system to successfully match the incoming stimulus to a stored representation. In general, this assumption was not supported by the data reported in this thesis and these null effects require explanation in this final chapter.

As early as experiment two there was a strong indication that sensitivity to mirror reversal of faces would be difficult to detect behaviourally. In this experiment it was reported that mirror reversing one image in a face pair had no effect on performance in an unfamiliar face matching procedure. As it had been previously confirmed that the two halves of a face were sufficiently asymmetric to produce a tendency for subjects to report these stimuli as representing two separate identities, this data suggests that face processing is subject to a process of 'mirror normalisation' prior to identification. As both accuracy and response time were unaffected by this manipulation it appeared likely that performance on this task was mediated by two forms of internal representation: one which represents the perceived face in its original orientation and one which represents it in its mirror reversed orientation.

The invariance to mirror orientation demonstrated in experiment two was replicated a number of times in chapter three using familiar faces as stimuli. Recognition accuracy was found to be unaffected by mirror reversal in experiments four, five, six and seven. Furthermore in three of these experiments the time taken to recognise a familiar face in its mirror orientation was not significantly slower than the time taken to recognise the same face in its original orientation. In combination these results provide strong evidence that face recognition processes depend upon a form of representation that is tolerant to mirror reversal. Given that the effect of mirror reversal on face recognition would appear negligible, it is of value to speculate as to why this may be, and also to consider what this result tells us about the nature of representations underpinning face recognition.

It has been previously hypothesised that the visual system represents stimuli in both their original and mirror reversed orientation by the left and right hemispheres respectively (Orton, 1925). Orton proposed this theory as an explanation for 'mirror confusion', which he argued was a probable cause of reading disorders. That language processing is highly lateralised in the brain is offered as support for this hypothesis, which predicts that highly lateralised functions allow for less ambiguous representation of orientation. As reading is dependent upon a uni-directional parsing strategy, representation of mirror orientations should cause disruption, and Orton proposed that dyslexia is caused by sub-normal development of left-hemisphere specialisation (Orton, 1928). This theory was further developed by Corballis and Beale (1976) who also suggested that the two hemispheres represent visual stimuli as mirror symmetric representations. According to their theory however, this organisation is not a product of the initial representation of visual stimuli (they conceded that the visual system is not structured in such a way as to enable this phenomenon) but instead is created by the nature of neural connections between the hemispheres, which connect symmetrical points.

Although these theories have an intuitive appeal, they have very little evidence to support them. As Gross and Bornstein (1978) point out, severe left-right confusion is most commonly caused by damage to the left-hemisphere only (Critchley, 1953). In addition, when stimuli are presented to one hemisphere in isolation, images are not perceived as being identical to their mirror reversal when this stimulus is presented to the opposite hemisphere (Bradshaw et al. 1973). Further studies have demonstrated this result in monkeys with severed corpus callosum (Hamilton & Tieman, 1973; Lehman & Spencer, 1973), and these results combine to disprove both Ortons (1925) and Corballis and Beales' (1976) hypotheses.

Gross and Bormstein (1978) argue convincingly that the tendency for mirror images to be treated as equivalent stimuli does not stem from the symmetrical organisation of the nervous system but instead evolved to complement the organisation of our visual environment. Mirror images, they contest, are treated as equivalent because often in the natural world they represent opposite views of the same subject. Vertebrates generally are bilaterally symmetric and therefore mirror images accurately approximate to the opposite profile view of an animal. This view is advocated, and empirically supported, in the research of Troje and Bülthoff (1998). Furthermore, more recent studies have provided a neurological basis for this hypothesis, showing that neurons in the inferotemporal cortex of monkeys generalise mirror images (Baylis & Driver, 2001; Rollenhagen and Olson, 2000).

If the brain generalises mirror images, this may cause problems for face recognition. A point that has been emphasised in this thesis is that asymmetry in faces is diagnostic of identity, and this is evidenced by the discriminatory power of this information (e.g. Burke & Healey, 1993) in addition to the usefulness of this biometric in automatic face recognition systems (e.g. Mitra et al. 2007). Furthermore, an implicit assumption has been that for a face to be recognised the perceived face must be successfully matched to some form of stored representation. To ascertain whether or not the asymmetry of a face is retained in this stored representation was the principal concern of the experiments reported in this thesis.

The way in which the mind creates representations of familiar faces is an issue that has received a surprisingly small amount of attention. Though the nature of these representations is often probed by researchers, relatively little time has been spent trying to understand the way in which these representations are derived from experience. Burton *et al* (2005) provide a rare attempt to explain this process, proposing that face recognition is underpinned by an 'average' representation of the visual exposure that an individual receives. These representations are said to provide a more reliable basis for recognition as they negate sources of information present in images that are unimportant for the recognition process (e.g. lighting direction). Furthermore average images of faces have been shown to improve face recognition in humans relative to individual ‘naturalistic’ images (Burton et al. 2005), and recently have been shown to enable perfect performance in an automatic recognition system (Jenkins & Burton, 2008). Given that a form of object constancy that allows mirror reversed images of faces to be treated as equivalent would appear to underpin face processing, it is interesting to speculate as to how this might affect such internal representations.

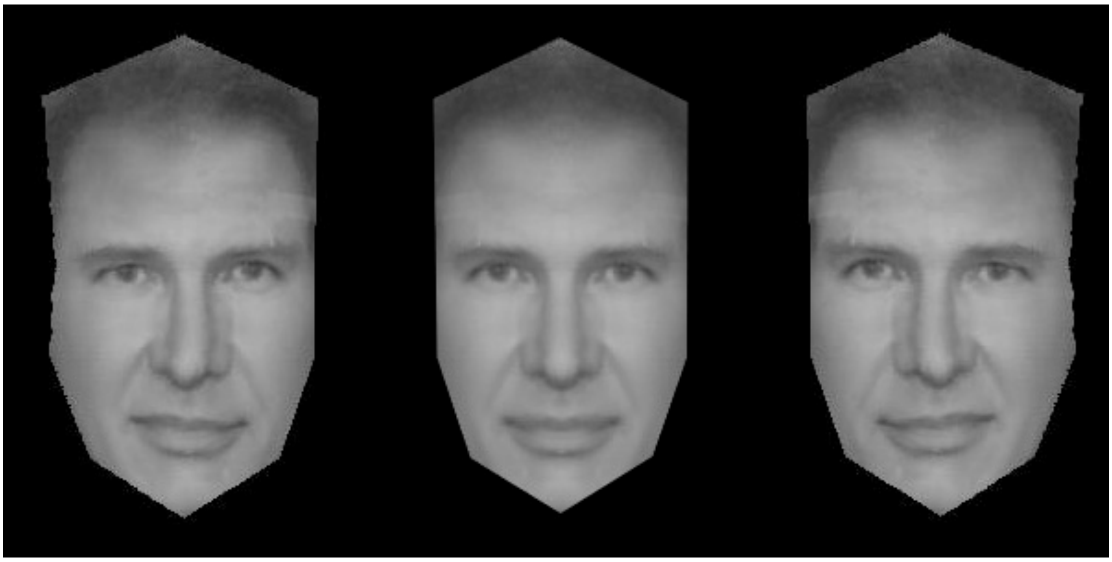

Figure 7.2: The averaging technique described by Burton *et al* (2005) was applied to images of Harrison Ford in their veridical (left) and reversed (right) orientations. The image in the middle shows the symmetrical image generated by averaging these representations together.

Should both veridical and mirror reversed orientations of faces contribute to the memory of a face, then according to Burton *et al* (2005) the resultant representation should be symmetrical (Figure 7.2, centre image). Intuitively speaking, this representation does not look very much like Harrison Ford as losing the asymmetries present in the face would appear to have removed aspects of the face that describe his appearance (e.g. his crooked nose and lopsided smirk). Although the hypothesis has not been tested it would seem highly likely that this stimulus would be recognised less well by both humans and computers than the original asymmetric stimulus. In experiment seven I demonstrated that ‘face averages’ shown in reversed orientation are recognised just as quickly as when they are presented in the original orientation. It is likely that this result is reflective of some normalisation process that allows the perceived asymmetries to be reversed as opposed to the existence of a representation that is a composite of both orientations.

Although mirror image generalisation is a ubiquitous phenomenon, the tendency to treat mirror images as equivalent is rarely absolute. For instance, Rollenhagen and Olsen (2000) found that neurons responded more similarly to lateral mirror images than to vertical mirror images, however these neurons were still less responsive to mirror images than to identical images. In addition, Troje and Bülthoff (1997) found that in two memory experiments, recognition of a face was superior when an identical image was presented at test compared to when a mirror reversed copy was presented in the test phase (though this difference was only significant in one of the experiments). Indeed, the complete inability to discriminate between two simultaneously presented mirror images is only found in rare instances of neurological damage (e.g. Davidoff & Warrington, 2001). It is likely therefore that although recognition systems are tolerant to lateral reversal, they still retain the ability to separate representations of the veridical and reversed orientations, thus enabling the representation of facial asymmetry.

In general, the results of the experiments carried out in this thesis demonstrate that despite an apparent invariance to lateral orientation in face recognition, facial asymmetry is to some degree represented in memory. Though the observed sensitivity to mirror reversal of familiar faces was typically slight, there are a number of instances in which this sensitivity was statistically significant. In experiment 4 recognition time was found to be slower for reversed relative to veridical presentations of famous faces. Though this effect was not replicated in two subsequent experiments (experiments 5 and 6), responses made to reversed images in these experiments were on average marginally slower than the responses made to veridically oriented stimuli. Furthermore, in the same chapter reversal was found to impede memory performance for previously unfamiliar faces, a result that replicated previous research (McKelvie, 1983).

That cognitive processes are impeded by the reversal of a face demonstrates that facial asymmetry is retained in memory. In chapter four this finding was confirmed using a 2AFC procedure whereby participants were required to identify the real-world orientation of a familiar face (i.e. to discriminate it from its mirror-image). In accordance with a previous report (Rhodes, 1986), participants could perform this task, though performance was only marginally superior to chance (experiments 10, 11 and 12) and performance on mirror discrimination was facilitated by increasing familiarity (experiment 12). Thus, the marginal cognitive sensitivity detected in the previous chapter was replicated in chapter four. From this data it is clear that explicit knowledge of veridical orientation is very difficult to recall, and performance on this task is poorer than in the 'likeness' judgements used by Rhodes (1986). It is possible that this difference reflects the implicit nature of memory for lateral orientation, with Rhodes' measure being more sensitive to this form of implicit representation.

Though the data reported here, in conjunction with previous research, serves to confirm that facial asymmetry is represented in memory, it is perplexing why effects of mirror reversal are found in some cases but not others. As has already been discussed, these effects are likely to be constrained by the fact that the visual system readily generalises mirror images. Furthermore, the fact that faces display a high degree of symmetry means that mirror reflections differ only in small ways from the original stimulus, and therefore only slight effects of this manipulation could realistically have been predicted. Indeed, that any effects of reversal have been detected in this thesis is testament to the accuracy with which spatial relationships on the human face are represented, and I have demonstrated that this accuracy improves with exposure to a face (experiment 12).

Although the lack of a consistent effect across experiments may simply reflect the slightness of the reversal effect, it is considered that some attempt should be made to explain why reversal did not produce behavioural effects in all the paradigms employed. The stimulus conditions used in experiment seven provide an example of a case in which the reversal effect was detected and also an example of a case in which it was not. When participants were tested using the same image that they had been exposed to in the learning phase the proportion of images correctly identified was unaffected by mirror reversal of the test stimulus. However, when different images of the same identity were used across learning and test, participants recalled significantly less mirror reversed test images than non-reversed images. This difference is evidence that the asymmetry in a face has been reliably represented by cognitive systems and that this asymmetry exists within a representation that enables face recognition to occur despite intervening image-level variation. However, just why there was no detectable effect of mirror reversal on the recollection of images when participants were tested using the same images remains unclear.

As has already been discussed in chapter three, the most likely explanation for the interaction between test image and mirror reversal is that the same image memory test was too easy for the subtle effect of reversal to be detected. Such an explanation implies that performance on the same image memory task had reached ceiling level, however the support for this assertion is rather weak. Indeed, only two out of the twenty-two participants in the 'same image' test condition correctly recalled all twenty of the learnt images. Another possible explanation is that the visual system supports two separate forms of representations, one that is invariant to mirror reversal, and another that is not. The representation of image-level 'surface' characteristics can be separated from a more abstract form of object representation (Hitch, Brandimonte & Walker, 1995) and if this framework can be equally applied to memory for faces then it could be that somewhere in the abstractive process lateral orientation becomes specified.

That separate forms of internal representation can be differentially sensitive to the lateral orientation of objects has previously been demonstrated by Kosslyn and Rabin (1999). In this study participants either selected the correct left-right orientation of a coin from a pair (2AFC), or had to form a mental image before performing this task. Kosslyn and Rabin (1999) report that in the 2AFC procedure performance was not significantly superior to chance, but when using the imagery method 22 out of 30 participants correctly recalled the orientation of the coin head. However, this result was only obtained when participants were instructed to imagine a specific coin and not when asked to bring to mind a general representation of a coin. In explaining their findings, Kosslyn and Rabin (1999) argue that the availability of the mirror reversals to perceptual systems in the 2AFC task make the task more difficult but that this difficulty could be remedied by recalling the form of a coin form memory prior to being presented with the two alternatives. Long-term memory representations, they contend, represent left-right orientation, but the process that matches a perceived object to these representations is invariant to reflection.

This dissociation may provide a resolution to an inconsistency between the findings reported in experiment four of this thesis and previous published research (Brooks et al. 2002). The name verification paradigms used in these two experiments are almost identical, however they generated different results. In experiment four of this thesis mirror reversal of a face was found to significantly increase response latency when verifying that this face belonged to the person whose name the participant had just previously seen. Brooks *et al* (2002) presented the stimuli in the opposite order (i.e. a face in either veridical or reversed orientation was presented prior to presentation of a name) and found that there was no effect of reversal. It could be that the design employed in the present thesis encouraged participants to recall the face from long-term memory more than the stimulus order used in the previous study, and therefore engaged processes that are sensitive to lateral orientation. This same methodological difference may also account for why there were no reported differences in recognition responses between moving the left eye and right eye in Cooper and Wojan (2000), yet these differences were detected in experiments 15 and 16.

Further research could attempt to determine whether the effect of mirror reversal is modulated by task difficulty, or by the degree to which a given representation can be said to be 'abstract' (e.g. see experiment 7). Also, in light of the study by Kosslyn and Rabin (1999), the effect of imagery on the ability to correctly identify left-right orientation should be investigated. If the use of imagery improves performance on the coin-head orientation task it may be that it also can improve performance in the 2AFC procedure used in chapter four. Indeed, a dissociation between perceptual and imagery ability has been demonstrated in prosopagnostic patients (Michelon & Biedermann, 2002), and this raises the possibility that separate mechanisms mediate performance on these two types of task. Whether or not imagery can improve performance on mirror discrimination task is an empirical question that could also be addressed using the paradigm described in chapter 5 to ascertain whether the 'general schema' of a face is asymmetrical.

The course of investigation followed in this thesis was designed to improve understanding of the way in which familiar faces are cognitively represented. Slight effects of mirror reversal were observed in a number of face processing tasks and sensitivity to the lateral orientation of familiar faces was therefore established. In line with previous research (e.g. Mita et al., 1977; Rhodes, 1986; McKelvie, 1983; Brédart, 2003), this finding shows that the mind represents faces in a detailed enough manner to allow subtle structural asymmetries to be represented. Given that the effects of reversal were most probably obscured by a general tendency for mirror reversals to be treated as equivalent, in addition to the high volume of redundant information present in faces in the task of recognition (e.g. Vinnete et al. 2004), it is likely that the structural asymmetry of familiar faces is coded in memory in a fairly robust manner.

It has previously been proposed that mechanisms responsible for detecting symmetry interact with higher-level cognitive processes responsible for face recognition (Rhodes et al. 2005). The contribution of symmetry detection mechanisms to face processing, in addition to the close proximity of cortical areas supporting these functions (Chen et al. 2007), is assumed to enable reliable preferences for symmetrical faces (Little & Jones, 2006). In light of the findings reported in this thesis however, it remains possible that brain mechanisms responsible for symmetry processing may also support face recognition. Further research manipulating the degree of asymmetry in faces would help to clarify the contribution of facial asymmetry to face recognition.

Rhodes (1986) argues that cognitive sensitivity to the reversal of familiar faces implicates a configural coding strategy in face memory and sensitivity to this manipulation has been demonstrated on a number of occasions in this thesis. Whether this implies that familiar faces are stored in memory primarily in terms of the configural relations between features remains uncertain however. Just because asymmetry can be expressed as a holistic property of the configural relations of a face doesn't mean that perceptual sensitivity to reversal of this information demonstrates that faces are described configurally in memory. As has already been discussed in chapter one, definitions of 'configural' and 'featural' representation are inherently circular, however mirror orientation of face images can clearly be disambiguated by both configural and featural information (Harrison Ford's crooked nose provides an example of the latter possibility).

Given the high degree of redundant information in the task of recognising faces, it is not possible to conclude whether 'reversal effects' result from an underlying system that represents faces as a product of their 'configuration' or as a sum of the facial features. However, the data reported in chapter 6 would suggest that the configuration of facial features forms at least part of this representation, and that the contribution of inter-featural relations to the overall representation of a face is not uniformly distributed. The salience of relational metrics would appear to be asymmetrically distributed, with the eye to the left of fixation being more heavily implicated in configural processing.

The effects of mirror reversal reported in this thesis are slight in comparison to the profound decrement in face recognition performance that result from vertical inversion (e.g. Yin, 1969). This difference can be explained in terms of the relative instability of the left-right axis compared with the up-down axis in the visual environment as well as in neuronal processing (e.g. Davidoff & Warrington, 2001). Furthermore, because first-order featural relations are preserved in mirror transformation, perceptual processes developed for parsing facial stimuli are not interrupted as severely by this manipulation. It was hypothesised that mirror reversal would nonetheless disrupt recognition processes because the second-order relations pertaining to identity are altered. The disruption caused by this transformation was however found to be slight, and participants typically find it extremely difficult to identify the correct left-right orientation of a face.

Finally, the data reported here serves to highlight the uniqueness of the inversion effect. Global transformation of face stimuli by stretching along the x-axis and y-axis does not impair face recognition performance (Hole et al. 2002). As has been shown here, neither does reflection across the y-axis. Reflection across the x-axis severely impairs face processing however, but this effect is mediated by the angle of rotation. It would appear therefore that uni-dimensional geometric operations on face stimuli are disruptive to recognition only in as much as they cause the stimulus to be oriented away from upright.

References

Bahrick, H. P., Bahrick P. O. and Wittlinger, R. P. (1975). 50 years of memory for names and faces– Cross-sectional approach. *Journal of Experimental Psychology: General* 104(1): 54-75.

Barlow, H. B. and Reeves B. C. (1979). Versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research* 19(7): 783-793.

Baylis, G. C. and Driver, J. (2001). Shape-coding in IT cells generalizes over contrast and mirror reversal, but not figure-ground reversal. *Nature Neuroscience* 4(9): 937-942.

Berlucchi, G. and Marzi C. A. (1970). Veridical interocular transfer of lateral mirror-image discriminations in split-chiasm cats. *Journal of Comparative and Physiological Psychology* 72(1), 1-7.

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, 94, 115-147.

Biederman, I. and Cooper E. E. (1991). Evidence for complete translational and reflectional invariance in visual object priming. *Perception* 20(5): 585-593.

Blount, P., J. Holmes, Rodger, J., Coltheart, M. and McManus, C. (1975). Ability to discriminate original from mirror-image reproductions of works of art. *Perception* 4(4): 385-389.

Bonner, L., Burton, A.M. & Bruce, V. (2003). Getting to know you: How we learn new faces. *Visual Cognition*, 10, 527-536.

Bornstein, M. H., Ferdinandsen K. and Gross, C. G. (1981). Perception of symmetry in infancy. *Developmental Psychology* 17(1): 82-86.

Bornstein, M. H., C. G. Gross, and Wolf, J. Z. (1978). Perceptual similarity of mirror images in infancy. *Cognition* 6(2): 89-116.

Borod, J. C., Haywood, C. S. and Koff, E. (1997). Neuropsychological aspects of facial asymmetry during emotional expression: A review of the normal adult literature. *Neuropsychology Review* 7(1), 41-60.

Bradshaw, J. L., Nettleto, N. C. and Patterson, K. (1973). Identification of mirror-reversed and nonreversed facial profiles in same and opposite visual-fields. *Journal of Experimental Psychology* 99(1): 42-48.

Brady, N., Campbell, M. and Flaherty, M. (2005). Perceptual asymmetries are preserved in memory for highly familiar faces of self and friend. *Brain and Cognition* 58(3): 334-342.

Brédart, S. (2003). Recognising the usual orientation of one's own face: the role of asymmetrically located details. *Perception* 32(7): 805-811.

Broca, P. P. (1861) Perte de la Parole, Ramollissement Chronique et Destruction Partielle du Lobe Antérieur Gauche du Cerveau, *Bulletin de la Société Anthropologique*, *2*, 235-238

Brooks, B. E. and Cooper, E. E. (2006). What types of visual recognition tasks are mediated by the neural subsystem that subserves face recognition? *Journal of Experimental Psychology: Learning Memory and Cognition* 32(4): 684-698.

Brooks, B. E., Rosielle, L. J., and Cooper, E. E. (2002). The priming of face recognition after metric transformations. *Perception* 31(3): 297-313.

Bruce V. and Humphreys G. W. (eds.) (1994). *Object and face recognition.* Lawrence Erlbaum Associates.

Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, A. M. and Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied* 5(4): 339-360.

Bruce, V. & Valentine, T. (1985). Identity priming in face recognition. *British Journal of Psychology*, 76, 373-383.

Bruyer, R., C. Galvez, C. and Prairial, C. (1993). Effect of disorientation on visual analysis, familiarity decision and semantic decision on faces. *British Journal of Psychology* 84: 433-441.

Burke P.H. and Healy M.J. (1993). A serial study of normal facial asymmetry in monozygotic twins. *Annals of Human Biology* 20(6): 527–534.

Burt, D. M. and Perrett, D. I. (1997). Perceptual asymmetries in judgements of facial attractiveness, age, gender, speech and expression. *Neuropsychologia* 35(5): 685-693.

Burton, A. M., Bruce, V. and Johnston, R. A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology* 81: 361-380.

Burton, A. M., Jenkins, R., Hancock, P. J. B. and White, D. (2005). Robust representations for face recognition: The power of averages. *Cognitive Psychology* 51(3): 256-284.

Burton, A. M., Wilson, S., Cowan, M. and Bruce, V. (1999). Face recognition in poor-quality video: Evidence from security surveillance. *Psychological Science* 10(3): 243-248.

Busey, T. A. and Vanderkolk, J. R. (2005). Behavioral and electrophysiological evidence for configural processing in fingerprint experts. *Vision Research* 45(4): 431-448.

Butler, S., Gilchrist, I. D., Burt, D. M., Perrett, D. I., Jones, E. and Harvey, M. (2005). Are the perceptual biases found in chimeric face processing reflected in eye-movement patterns? *Neuropsychologia* 43(1): 52

Butler, S. and Harvey, M. (2005). Scanning bias enhancement of the perceptual biases found in chimeric face processing. *Journal of Cognitive Neuroscience*: Suppl. 44-45.

Butler, S. H. and Harvey, M. (2006). Perceptual biases in chimeric face processing: Eye-movement patterns cannot explain it all. *Brain Research* 1124: 96-99.

Calder, A. J., Young, A. W., Keane, J., Dean, M. (2000). Configural information in facial expression perception. *Journal of Experimental Psychology: Human Perception and Performance* 26(2): 527-551.

Chen, C. C., Kao, K. L. C. and Tyler, C. W. (2007). Face configuration processing in the human brain: The role of symmetry. *Cerebral Cortex* 17(6): 1423-1432.

Cohen J. D., MacWhinney B., Flatt M., & Provost J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments and Computers,* 25, 257-271.

Collishaw, S. M. and Hole, G. J. (2000). Featural and configurational processes in the recognition of faces of different familiarity. *Perception* 29(8): 893-909.

Collishaw, S. M. and Hole, G. J. (2002). Is there a linear or a nonlinear relationship between rotation and configural processing of faces? *Perception* 31(3): 287-296.

Cooper, A. C. G. and Humphreys, G. W. (2000). Task-specific effects of orientation information: neuropsychological evidence. *Neuropsychologia* 38(12): 1607-1615.

Cooper, E. E. and Brooks, B. E. (2004). Qualitative differences in the representation of spatial relations for different object classes. *Journal of Experimental Psychology: Human Perception and Performance* 30(2): 243-256.

Cooper, E. E. and Wojan, T. J. (2000). Differences in the coding of spatial relations in face identification and basic-level object recognition. *Journal of Experimental Psychology: Learning Memory and Cognition* 26(2): 470-488.

Corballis, M. C. (2003). Laterality and human speciation. *Speciation of Modern Homo Sapien* 106: 137-152.

Corballis, M. C., and Beale, I. L. (1976) *The Psychology of Left and Right.* New York, Halstead.

Corballis, P. M., Funnell, M. G. and Gazzaniga, M. S. (2000). An evolutionary perspective on hemispheric asymmetries*. Brain and Cognition* 43(1-3): 112-117.

Cornell, J. M. (1985). Spontaneous mirror-writing in children. *Canadian Journal of Psychology-Revue Canadienne De Psychologie* 39(1): 174-179.

Cornish, K. M. and McManus, I. C. (1996). Hand preference and hand skill in children with autism. *Journal of Autism and Developmental Disorders* 26(6): 597-609.

Critchley, M. (1926). The significance of mirror writing. *Proceedings of the Royal Society of Medicine* **20** (1926), pp. 397–404.

Critchley, M. (1953) *The Parietal Lobes.* London.

Cutting, J. E. (2003). Gustave Caillebotte, French impressionism, and mere exposure. *Psychonomic Bulletin & Review* 10(2): 319-343.

Dahl, C. D., N. K., Logothetis, N. K. and Hoffman, K. L. (2007). Individuation and holistic processing of faces in rhesus monkeys. *Proceedings of the Royal Society B-Biological Sciences* 274(1622): 2069-2076.

David, A. S. (1993). Spatial and selective attention in the cerebal hemispheres in depression, mania and scizophrenia. *Brain and Cognition* 23(2): 166-180.

Davidoff, J. and Warrington, E. K. (1999). The bare bones of object recognition: implications from a case of object recognition impairment. *Neuropsychologia* 37(3): 279-292.

Davidoff, J. and Warrington, E. K. (2001). A particular difficulty in discriminating between mirror images. *Neuropsychologia* 39(10): 1022-1036.

Davidson, H. P. (1935). A study of the confusing letters B, D, P and Q. *Pedagogical Seminary and Journal of Genetic Psychology* 47(2): 458-468.

Deregowski, J. B. (1971). Symmetry, Gestalt and Information Theory. *Quarterly Journal of Experimental Psychology* 23: 381-395.

DeRenzi, E., D. Perani, D., Carlesimo, G. A., Silveri, M. C. and Fazio, F. (1994). Prosopagnosia can be associated with damage confined to the right-hemisphere – An MRI and PET study and a review of the literature. *Neuropsychologia* 32(8): 893-902.

Diamond, R. and Carey, S. (1986). Why faces are not special: an effect of expertise. *Journal of Experimantal Psychology: General* 115(2): 107–117.

Desimone, R. (1991). Face-selective cells in the temporal cortex of monkeys. *Journal of Cognitive Neuroscience* 3(1): 1-8.

Duchaine, B. C. and Nakayama, K. (2006). Developmental prosopagnosia: a window to content-specific face processing. *Current Opinion in Neurobiology* 16(2): 166-173.

Dutta, T. and. Mandal, M. K (2002). Visual-field superiority as a function of stimulus type and content. *International Journal of Neuroscience* 112(8): 945-952.

Dyer, A. G., Neumeyer, C. and Chittka, L. (2005). Honeybee (Apis mellifera) vision can discriminate between and recognise images of human faces. *Journal of Experimental Biology* 208(24): 4709-4714.

Ekman, P., Hager, J. C. and Friesen, W. V. (1981). The symmetry of emotional and deliberate facial actions. *Psychophysiology* 18(2): 101-106.

Ellis, H.D. (1981). Theoretical aspects of face recognition. In G. Davies, H. Ellis, & J. Sheperd (Eds.), *Perceiving and remembering faces*. London: Academic Press.

Falkai, P., Honer, W. G., Alfter, D., Schneider-Axmann, T., Bussfeld, P., Cordes, J., Blank, B., Schonell, H., Steinmetz, H., Maier, W., Tepest, R. (2002). The temporal lobe in schizophrenia from uni- and multiply affected families. *Neuroscience Letters* 325(1): 25-28.

Farah, M. J., Tanaka, J. W. and Drain, H. M. (1995). What causes the face inversion effect. *Journal of Experimental Psychology: Human Perception and Performance* 21(3): 628-634.

Farah, M. J., Wilson, K. D., Drain, M. and Tanaka, J. N. (1998). What is "special" about face perception? *Psychological Review* 105(3): 482-498.

Farkas LG. (1994). *Anthropometry of the head and face*. Raven Press, New York.

Farkas, L. G. and G. Cheung (1981). Facial asymmetry in healthy North American Caucasians – An anthropometrical study. *Angle Orthodontist* 51(1): 70-77.

Farkas L. G and Munro I. R. (1987) *Anthropometric facial proportions in medicine.* Charles C Thomas, Illinois.

Ferber, S., Murray L. (2005). Are perceptual judgments dissociated from motor processes? - A prism adaptation study. *Cognitive Brain Research*, 23: 453-456

Fellous, J. M. (1997). Gender discrimination and prediction on the basis of facial metric information. *Vision Research* 37(14): 1961-1973.

Ferrario, V. F., Sforza, C., Ciusa, V., Dellavia, C. and Tartaglia, G. M. (2001). The effect of sex and age on facial asymmetry in healthy subjects: A cross-sectional study from adolescence to mid-adulthood. *Journal of Oral and Maxillofacial Surgery* 59(4): 382-388.

Ferrario, V. F., Sforza, C., Pizzini, G., Vogel, G. and Miani, A. (1993). Sexual dimorphism in the human face assessed by Euclidian distance matrix analysis. *Journal of Anatomy* 183: 593-600.

Ferrario, V. F., Sforza, C., Poggio, C. E., Schmitz, J. H. and Colombo, A. (1997). Soft tissue facial morphology related to headform: A three-dimensional quantitative analysis in childhood. *Journal of Craniofacial Genetics and Developmental Biology* 17(2): 86-95.

Freire, A., Lee, K. and Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception* 29(2): 159-170.

Gallois, P., Buquet, C., Charlier, J., Paris, V., Hache, J. C. and Dereux, J. F. (1989). Asymmetry in the strategy for visual-perception of faces and emotional facial expressions. *Revue Neurologique* 145(8-9): 661-664.

Gangestad, S. W. and Thornhill, R. (1997). The evolutionary psychology of extrapair sex: The role of fluctuating asymmetry. *Evolution and Human Behavior* 18(2): 69-88.

Gauthier, I. and Tarr , M. J. (1997). Becoming a "greeble" expert: Exploring mechanisms for face recognition. *Vision Research* 37(12): 1673-1682.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P. and Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience* 2(6): 568-573.

Gazzaniga, M. S. (2005). Forty-five years of split-brain research and still going strong. *Nature Reviews Neuroscience* 6(8): 653-659.

Geschwind. N and Levitsky, W. (1968). Human brain – Left-right asymmetries in temporal speech region. *Science* 161(3837): 186-187.

Gilbert, C. and Bakan, P. (1973). Visual asymmetry in perception of faces. *Neuropsychologia* 11(3): 355-362.

Giurfa, M., Eichmann, B. and Menzel, R. (1996). Symmetry perception in an insect. *Nature* 382(6590): 458-461.

Glick, S. D., Ross, D. A. and Hough, L. B. (1982). Lateral asymmetry of neurotransmitters in human-brain. *Brain Research* 234(1): 53-63.

Goffaux, V., Sorger, B., Schiltz, C., Goebel, R., & Rossion, B. (2007). Investigation of featural versus configural processing of faces in the middle fusiform gyrus [Abstract]. *Journal of Vision*, *7*(9): 625a

Gordon, P. C. and. Holyoak, K. J (1983). Implicit learning and generalisation of the mere exposure effect. *Journal of Personality and Social Psychology* 45(3): 492-500.

Gosselin, F. and Schyns, P. G. (2001). Bubbles: a technique to reveal the use of information in recognition tasks. *Vision Research* 41(17): 2261-2271.

Grammer, K. and Thornhill, R. (1994). Human (homo-sapiens) facial attractiveness and sexual selection – The role of symmetry and averageness. *Journal of Comparative Psychology* 108(3): 233-242.

Grattan, M. P., Devos, E., Levy, J. and McClintock, M. K. (1992). Asymmetric action in the human newborn – Sex-differences in patterns of organization. *Child Development* 63(2): 273-289.

Gross, C. G., & Bornstein, M. H. (1978), Left and right in science and art. *Leonardo,* 11, 29-38.

Grüsser, O. J., Selke, T. and Zynda, B. (1988). Cerebral lateralization and some implications for art, aesthetic perception and aristic creativity. In I. Rentschler, B. Herzberger, & D. Epstein (Eds.), *Beauty and the brain. Biological aspects of aesthetics* (pp. 257-293). Boston: Birkhauser.

Güntürkün, O. (2003). Adult persistence of head-turning asymmetry. *Nature* 421(6924): 711-711.

Haig, N. D. (1984). The effect of feature displacement on face recognition. *Perception* 13(5): 505-512.

Halberstadt, J. and Rhodes, G. (2003). It's not just average faces that are attractive: Computer-manipulated averageness makes birds, fish, and automobiles attractive. *Psychonomic Bulletin & Review* 10(1): 149-156.

Halgren, E., Dale, A. M., Sereno, M. I., Tootell, R. B. H., Marinkovic, K. and Rosen, B. R. (1999). Location of human face-selective cortex with respect to retinotopic areas. *Human Brain Mapping* 7(1): 29-37.

Hamilton, C. R. and Tieman, S. B. (1973). Interocular transfer of mirror-image discriminations by chiasm-sectioned monkeys. *Brain Research* 64 (21): 241-255.

Hancock, P.J.B., Bruce, V. & Burton, A.M. (2000). Recognition of unfamiliar faces. *Trends in Cognitive Science*, 4(9), 330-337.

Hardie, S., Hancock, P. J. B., Rodway, P., Penton-Voak, I., Carson, D. & Wright, L. (2005). The enigma of facial asymmetry: Is there a gender-specific pattern of facedness? *Laterality*, 10 (4), 295-304

Harris, I. M., Harris, J. A. and Caine, D. (2001). Object orientation agnosia: A failure to find the axis? *Journal of Cognitive Neuroscience* 13(6): 800-812.

Haxby, J. V., Ungerleider, L. G., Clark, V. P., Schouten, J. L., Hoffman, E. A. and Martin, A. (1999). The effect of face inversion on activity in human neural systems for face and object perception. *Neuron* 22(1): 189-199.

Hayes, A. (1988). Identification of 2-tone images – Some implications for high-spatial-frequency and low-spatial-frequency processes in human-vision. *Perception* 17(4): 429-436.

Heller, W. and Levy, J. (1981). Perception and expression of emotion in right-handers and left-handers. *Neuropsychologia* 19(2): 263-272.

Hill H, Schyns P G, Akamatsu S, 1997 Information and viewpoint dependence in face recognition *Cognition* 62: 201–222.

Hitch, G. J., Brandimaonte, M. A. & Walker, P. (1995). Two types of representation in visual memory: Evidence from the effects of stimulus contrast on image combination. *Memory & Cognition*, 23(2), 147-154.

Hobert, O., Johnston, R. J. and Chang, S. (2002). Left-right asymmetry in the nervous system: The Caenorhabditis elegans model. *Nature Reviews Neuroscience* 3(8): 629-640.

Hole, G. J., George, P. A., Eaves, K. and Rasek, A. (2002). Effects of geometric distortions on face-recognition performance. *Perception* 31(10): 1221-1240.

Hosie, J. A., Ellis, H. D. and Haig, N. D. (1988). The effect of feature displacement on the perception of well-known faces. *Perception* 17(4): 461-474.

Hugdahl, K., Heiervang, E., Nordby, H., Smievoll, A. I., Steinmetz, H., Stevenson, J. and Lund, A. (1998). Central auditory processing, MRI morphometry and brain laterality: Applications to dyslexia. *Scandinavian Audiology* 27: 26-34.

Hugdahl, K., Iversen, P. M. and Johnsen, B. H. (1993). Laterality for facial expressions – does the sex of the subject interact with the sex of the stimulus face. *Cortex* 29(2): 325-331.

Hugdahl, K. & Davidson, R. J. (eds.) (2003). *The Asymmetrical Brain.* Massachusetts: MIT Press.

Hummel, J. E. and Biederman I. (1992). Dynamic binding in a neural network for shape-recognition. *Psychological Review* 99(3): 480-517.

Jenkins, R., Burton, A.M. (2008). 100% Accuracy in Automatic Face Recognition. *Science*, 319, 435

Jenkins, R., Burton, A. M. and Ellis, A. W. (2002). Long-term effects of covert face recognition. *Cognition* 86(2): B43-B52.

Johnson, L. N. (2005). Asymmetry at the molecular level in biology. *European Review.* 13 (S2): 77-95.

Jokel, R. and Conn D. (1999). Mirror reading, writing and backward speech in a woman with a head injury: a case of conversion disorder. *Aphasiology* 13(6): 495-509.

Jones, D. and Hill K. (1993). Criteria of facial attractiveness in 5 populatiuons. *Human Nature-an Interdisciplinary Biosocial Perspective* 4(3): 271-296.

Jones, G. V. (1990). Misremembering a common object- when left is not right. *Memory & Cognition* 18(2): 174-182.

Jones, G. V. and Martin M. (1997). Handedness dependency in recall from everyday memory. *British Journal of Psychology* 88: 609-619.

Karnath, H. O., Ferber, S. and Bulthoff, H. H. (2000). Neuronal representation of object orientation. *Neuropsychologia* 38(9): 1235-1241.

Kelly, S. W., Burton, A. M., Kato, T. and Akamatsu, S. (2001). Incidental learning of real-world regularities. *Psychological Science* 12(1): 86-89.

Kemp, R., McManus, C. and Pigott, T. (1990). Sensitivity to the displacement of facial features in negative and inverted images. *Perception* 19(4): 531-543.

Kennedy, D. N., O'Craven, K. M., Ticho, B. S., Goldstein, A. M., Makris, N. and Henson, J. W. (1999). Structural and functional brain asymmetries in human situs inversus totalis. *Neurology* 53(6): 1260-1265.

Klingenberg C. P. (2003). A developmental perspective on developmental instability: Theory, models, and mechanisms. In: *Developmental instability: causes and consequences* (Polak M, ed).Oxford: Oxford University Press; 14–34.

Koffka, K. (1935). *Principles of Gestalt Psychology*. Harcourt (New York)

Kosslyn, S. M. and Rabin, C. (1999). The representation of left-right orientation: A dissociation between imagery and perceptual recognition. *Visual Cognition* 6(5): 497-508.

LambonRalph, M. A., Jarvis, C. and Ellis, A. W. (1997). Life in a mirrored world: Report of a case showing mirror reversal in reading and writing and for non-verbal materials. *Neurocase* 3(4): 249-258.

Langlois, J. H., Roggman, L. A. and Musselman, L. (1994). What is average and what is not about attractive faces. *Psychological Science* 5(4): 214-220.

Le Grand, R., Mondloch, C. J., Maurer, D. and Brent, H. P. (2001). Neuroperception - Early visual experience and face processing. *Nature* 410(6831): 890-890.

Le Grand, R., Mondloch, C. J., Maurer, D. and Brent, H. P. (2004). Impairment in holistic face processing following early visual deprivation. *Psychological Science* 15(11): 762-768.

Leder, H. and Bruce, V. (1998). Local and relational aspects of face distinctiveness. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology* 51(3): 449-473.

Leder, H. and Bruce, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology* 53(2): 513-536.

Leder, H., Candrian, G., Huber, O. and Bruce, V. (2001). Configural features in the context of upright and inverted faces. *Perception* 30(1): 73-83.

Lehman RA, Spencer DD (1973). Mirror image shape discriminations: interocular reversal of responses in the optic chiasm sectioned monkey. *Brain Research* 52: 233–241.

Lehrer, M., Horridge, G. A., Zhang, S. W. and Gadagkar, R. (1995). Shape vision in bees – Innate preference for flower-like patterns. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences* 347(1320): 123-137.

Little, A. C. and Jones, B. C. (2006). Attraction independent of detection suggests special mechanisms for symmetry preferences in human face perception. *Proceedings of the Royal Society B: Biological Sciences* 273(1605): 3093-3099.

Liu Y., Schmidt K., Cohn J., and Weaver R.L. (2002). Human facial asymmetry for expression-invariant facial identification. In: *Proceedings of the Fifth IEEE International Conference on Automatic Face and Gesture Recognition (FG'02).*

Liu Y., Schmidt K., Cohn J., and Mitra S. (2003). Facial asymmetry quantification for expression-invariant human identification. *Computer Vision and Image Understanding Journal* 91(1/2): 138–159.

Locher, P. J., & Nodine, C. F. (1973). Influence of stimulus symmetry on visual scanning patterns. *Perception & Psychophysics* 13*,* 408- 412.

Loftus, G. R. & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review,* 1, 476-490.

Luh, K. E., Rueckert, L. M. and Levy, J. (1991). Perceptual asymmetries for free viewing of several types of chimeric stimuli. *Brain and Cognition* 16(1): 83-103.

Mach, E. (1914). *The Analysis of Sensations and the Relation of the Physical to the Psychical,* S. Waterlow (ed.). Chicago: Open Court.

Madigan, S. and Rouse, M. (1974). Picture memory and visual-generalisation processes. *American Journal of Psychology* 87(1-2): 151-158.

Malcolm, G. L., Leung, C. and Barton, J. J. S. (2004). Regional variation in the inversion effect for faces: Differential effects for feature shape, feature configuration, and external contour. *Perception* 33(10): 1221-1231.

Marr, D. and Nishihara, H. K. (1978). Representation and recognition of spatial-organization of 3-dimensional shapes. *Proceedings of the Royal Society of London Series B: Biological Sciences* 200(1140): 269-294.

Mather, K. (1953). Genetical control of stability in development. *Heredity* 7(3): 297-336.

Mathewson, L. (2004). Mirror writing ability is genetic and probably transmitted as a sex-linked dominant trait: it is hypothesised that mirror writers have bilateral language centres with a callosal interconnection. *Medical Hypotheses* 62(5): 733-739.

Manza, L., Reber, A. S., & Zizak, D. (1998). Artificial grammar learning and the mere exposure preference tasks and the implicit processes. In M. A. Stadler & P. A. Frensch (Eds.), *Handbook of implicit learning* (pp. 201–222). Thousand Oaks, CA: Sage Publications.

Maurer, D., Le Grand, R. and Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences* 6(6): 255-260.

Maurer, D. O'Craven, K. M. Le Grand, R. Mondloch, C. J. Springer, M. V. Lewis, T. L. Grady, C. L. (2007). Neural correlates of processing facial identity based on features versus their spacing. *Neuropsychologia*, 45(7): 1438-1451.

McKelvie, S. J. (1983). Effects of lateral reversal on recognition memory for photographs of faces. *British Journal of Psychology* 74: 391-407.

McKelvie, S. J. (1987). Vertical reversal and lateral reversal as components of inversion in recognition memory for upside-down faces. *Perceptual and Motor Skills* 65(2): 388-390.

McKelvie, S. J. and Aikins, S. (1993). Why is coin head orientation misremembered – Tests of schema interference and handedness hypotheses. *British Journal of Psychology* 84: 355-363.

McKone, E., Kanwisher, N. and Duchaine, B. C (2007). Can generic expertise explain special processing for faces? *Trends in Cognitive Sciences* 11(1): 8-15.

McKone, E., Martini, P. and Nakayama, K. (2001). Categorical perception of face identity in noise isolates configural processing. *Journal of Experimental Psychology: Human Perception and Performance* 27(3): 573-599.

McManus, C. (2005). Reversed bodies, reversed brains, and (some) reversed behaviors: Of zebrafish and men. *Developmental Cell* 8(6): 796-797.

McManus, I. C., Buckman, J. and Woolley, E. (2004). Is light in pictures presumed to come from the left side? *Perception* 33(12): 1421-1436.

McManus, I. C. and Humphrey, N. K. (1973). Turning left cheek. *Nature* 243(5405): 271-272.

McManus, I. C., Martin, N., Stubbings, G. F., Chung, E. M. K. and Mitchison, H. M. (2004). Handedness and situs inversus in primary ciliary dyskinesia. *Proceedings of*

the Royal Society of London Series B- Biological Sciences 271(1557): 2579-2582.
Megreya, A. M. and. Burton, A. M (2006). Unfamiliar faces are not faces: Evidence from a matching task. *Memory & Cognition* 34(4): 865-876.

Megreya, A.M. & Burton, A.M. (2007). Hits and false positives in face matching: A familiarity-based dissociation. *Perception and Psychophysics*, 69, 1175-1184.

Michelon, P, & Biederman, I. (2003). Less impairment in face imagery than face perception in early prosopagnosia. *Neuropsychologia,* 41, 421-441.

Minear, M. and Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments, and Computers*, 36(4), 630-633.

Mita, T. H., Dermer, M. and Knight, J. (1977). Reversed facial images and mere-exposure hypothesis. *Journal of Personality and Social Psychology* 35(8): 597-601.

Mitra, S., Lazar, N. A. and Liu, Y. X. (2007). Understanding the role of facial asymmetry in human face identification. *Statistics and Computing* 17(1): 57-70.

Møller, A. P. (1995). Bumblebee preference for symmetrical flowers. *Proceedings of the National Academy of Sciences of the United States of America* 92(6): 2288-2292.

Møller, A. P. and Sorci, G. (1998). Insect preference for symmetrical artificial flowers. *Oecologia* 114(1): 37-42.

Møller, A. P. and Thornhill, R. (1998). Bilateral symmetry and sexual selection: A meta-analysis. *American Naturalist* 151(2): 174-192.

Newell, B. R. and Bright, J. E. H. (2001). The relationship between the structural mere exposure effect and the implicit learning process. *Quarterly Journal of Experimental Psychology Section a-Human Experimental Psychology* 54(4): 1087-1104.

Newell, B. R. and Shanks, D. R. (2007). Recognising what you like: Examining the relation between the mere-exposure effect and recognition. *European Journal of Cognitive Psychology* 19(1): 103-118.

Nicholls, M. E. R., Clode, D., Wood, S. J. and Wood, A. G. (1999). Laterality of expression in portraiture: putting your best cheek forward. *Proceedings of the Royal Society of London Series B- Biological Sciences* 266(1428): 1517-1522.

Nickerson, R. S. and Adams, M. J. (1979). Long-term-memory for a common object. *Cognitive Psychology* 11(3): 287-307.

Noble, J. (1966). Mirror-images and forebrain commissures of monkey. *Nature* 211(5055): 1263-1266.

Noble, J. (1968). Paradoxical interocular transfer of mirror-image discriminations in the optic chiasm sectioned monkey. *Brain Research*, 10(2), 127-151.

Oinorten, K. A. and Mamanian, D. (2007). Facial symmetry detection ability changes across the menstrual cycle. *Biological Psychology* 75(2): 136-145.

Orton, S. T. (1925). World-blindness in school children. *Archives of Neurology and Psychiatry* 14(5): 581-615.

Orton, S. (1928). Specific reading disability – strephosymbolia. *Journal of the American Medical Association*, 90, 1095-1099.

Palmer A. R. and Strobeck, C. (2003). Fluctuating asymmetry analyses revisited. In: *Developmental instability: causes and consequences* (Polak M, ed). Oxford: Oxford University Press; 279–319.

Parsons, P. A. (1990). Fluctuating asymmetry – An epigenetic measure of stress. *Biological Reviews of the Cambridge Philosophical Society* 65(2): 131-145.

Parsons, P. A. (1990). Fluctuating asymmetry and stress intensity. *Trends in Ecology & Evolution* 5(3): 97-98.

Pavlova, M. and Sokolov, A. (2000). Orientation specificity in biological motion perception. *Perception & Psychophysics* 62(5): 889-899.

Peck, S., Peck, L. and Kataja, M. (1991). Skeletal asymmetry in aesthetically pleasing faces. *Angle Orthodontist* 61(1): 43-48.

Perrett, D. I. and Oram, M. W. (1993). Neurophysiology of shape processing. *Image and Vision Computing* 11(6): 317-333.

Petty, R. G. (1999). Structural asymmetries of the human brain and their disturbance in schizophrenia. *Schizophrenia Bulletin* 25(1): 121-139.

Pflugshaupt, T., Nyffeler, T., von Wartburg, R., Wurtz, P., Luthi, M., Hubl, D., Gutbrod, K., Juengling, F. D., Hess, C. W. and Muri, R. M. (2007). When left becomes right and vice versa: Mirrored vision after cerebral hypoxia. *Neuropsychologia* 45(9): 2078-2091.

Phillips, P. J., Moon, H., Rizvi, S. A., and Rauss, P. (2000). The FERET evaluation methodology for face-recognition algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 1090–1104.

Phillips, P. J., Grother, P., Michaelis, R. J., Blackburn, D. M., Tabassi, E., and Bone, J. M. (2003). *Face recognition vendor test 2002: Evaluation report NISTIR 6965.* Available online at http://www.frvt.org.

Phillips, R. J. and Rawles, R. E. (1979). Recognition of upright and inverted faces – Correlation study. *Perception* 8(5): 577-583.

Polak, M. E. (2003). *Developmental Instability. Causes and Consequences.* New York: Oxford University Press.

Posner, M. I. and Keele, S. W. (1968). On genesis of abstract ideas. *Journal of Experimental Psychology* 77(3P1): 353-&.

Priftis, K., Rusconi, E., Umilta, C. and Zorzi, M. (2003). Pure agnosia for mirror stimuli after right inferior parietal lesion. *Brain* 126: 908-919.

Ramachandran, V. S., Altschuler, E. L. and Hillyer, S. (1997). Mirror agnosia. *Proceedings of the Royal Society of London Series B-Biological Sciences* 264(1382): 645-647.

Ramsdell, A. F. and Yost, H. J. (1998). Molecular mechanisms of vertebrate left-right development. *Trends in Genetics* 14(11): 459-465.

Reynolds G. P. (1983) Increased concentrations and lateral asymmetry of amygdala dopamine. in schizophrenia. *Nature* 305: 527-529.

Rhodes, G. (1985). Perceptual asymmetries in face recognition. *Brain and Cognition* 4(2): 197-218.

Rhodes, G. (1986). Memory for lateral asymmetries in well-known faces – Evidence for configural information in memory representations of faces. *Memory & Cognition* 14(3): 209-219.

Rhodes, G. (2006).The evolutionary psychology of facial beauty. *Annual Review of Psychology* 57: 199-226.

Rhodes, G., Brake, S. and Atkinson, A. P. (1993). What's lost in inverted faces. *Cognition* 47(1): 25-57.

Rhodes, G., J. Halberstadt, J., Jeffery, L. and Palermo, R. (2005). The attractiveness of average faces is not a generalized mere exposure effect. *Social Cognition* 23(3): 205-217.

Rhodes, G., Jeffery, L., Watson, T., Clifford, C. and Nakayama, K. (2003a). Face attractiveness aftereffects: Fitting the mind to the world. *Australian Journal of Psychology* 55: 89-89.

Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W. G. and Nakayama, K. (2003b). Fitting the mind to the world: Face adaptation and attractiveness aftereffects. *Psychological Science* 14(6): 558-566.

Rhodes, G., Peters, M., Lee, K., Morrone, M. C., & Burr, D. (2005). Higher-level mechanisms detect facial asymmetry. *Proceedings of the Royal Society of London Series B- Biological Sciences* 272: 1379-1384.

Rhodes, G., Proffitt, F., Grady, J. M. and Sumich, A. (1998). Facial symmetry and the perception of beauty. *Psychonomic Bulletin & Review* 5(4): 659-669.

Rhodes, G., Sumich, A. and Byatt, G. (1999). Are average facial configurations attractive only because of their symmetry? *Psychological Science* 10(1): 52-58.

Rhodes, G. and Tremewan T. (1996). Averageness, exaggeration, and facial attractiveness. *Psychological Science* 7(2): 105-110.

Rhodes, G., Zebrowitz, L. A., Clark, A., Kalick, S. M., Hightower, A. & McKay, R. (2001). Do facial averageness and symmetry signal health? *Evolution and Human Behavior* 22, 31–46.

Riesenhuber, M., Jarudi, I., Gilad, S. and Sinha, P. (2004). Face processing in humans is compatible with a simple shape-based model of vision. *Proceedings of the Royal Society of London Series B- Biological Sciences* 271: S448-S450.

Robbins, R. and McKone E. (2007). No face-like processing for objects-of-expertise in three behavioural tasks. *Cognition* 103(1): 34-79.

Roberts, K. (2006). *Making Faces.* Great Britain: Norfolk Contemporary Arts Centre.

Rock I. (1973). *Orientation and Form*. Academic Press, New York

Rock I. (1983). *The Logic of Perception*. Cambridge, MA: MIT Press

Rollenhagen, J. E. and Olson, C. R. (2000). Mirror-image confusion in single neurons of the macaque inferotemporal cortex. *Science* 287(5457): 1506-1508.

Ronnqvist, L. (1995). A critical-examination of the moro response in newborn-infants – Symmetry, state relation, underlying mechanisms. *Neuropsychologia* 33(6): 713-726.

Rossion, B., Caldara, R., Seghier, M., Schuller, A. M., Lazeyras, F. and Mayer, E. (2003). A network of occipito-temporal face-sensitive areas besides the right middle fusiform gyrus is necessary for normal face processing. *Brain* 126: 2381-2395.

Rossion, B., Gauthier, I., Tarr, M. J., Despland, P., Bruyer, R., Linotte, S. and Crommelinck, M. (2000). The N170 occipito-temporal component is delayed and enhanced to inverted faces but not to inverted objects: an electrophysiological account of face-specific processes in the human brain. *Neuroreport* 11(1): 69-74.

Rubin, D. C. and Kontis T. C. (1983). A schema for common cents. *Memory & Cognition* 11(4): 335-341.

Rudel, R. G. and Teuber H. L. (1963). Discrimination of direction of line in children. *Journal of Comparative and Physiological Psychology* 56(5): 892-905.

Sackeim, H. A. (1985). Morphologic asymmetries of the face – A review. *Brain and Cognition* 4(3): 296-312.

Sackeim, H. A., Gur, R. C. and Saucy, M. C. (1978). Emotions are expressed more intensely on the left side of face *Science* 202(4366): 434-436.

Scheib, J. E., Gangestad, S. W. and Thornhill, R. (1999). Facial attractiveness, symmetry and cues of good genes. *Proceedings of the Royal Society of London Series B- Biological Sciences* 266(1431): 1913-1917.

Schmidt, K. L., Liu, Y. X. and Cohn, J. F. (2006). The role of structural facial asymmetry in asymmetry of peak facial expressions. *Laterality* 11(6): 540-561.

Schweinberger, S. R., Pickering, E. C., Jentzsch, I., Burton, A. M. and Kaufmann, J. M. (2002). Event-related brain potential evidence for a response of inferior temporal cortex to familiar face repetitions. *Cognitive Brain Research* 14(3): 398-409.

Sekuler, R. and Houlihan, K. (1968). Discrimination of mirror-images: Choice time analysis of human adult performance. *Quarterly Journal of Experimental Psychology* 20: 204-207.

Sergent, L. (1986). Microgenesis of face perception. In H. D. Ellis, M. A. Jeeves, F. Newcombe, & A. W. Young (Eds.), *Aspects of face processing*. Dordecht: Martinus NijhoV.

Sergent, J., S. Ohta, S. and Macdonald, B. (1992). Functional neuroanatomy of face and object processing – A positron emission tomography study. *Brain* 115: 15-36.

Shapleske, J., Rossell, S. L., Woodruff, P. W. R. and David, A. S. (1999). The planum temporale: a systematic, quantitative review of its structural, functional and clinical significance. *Brain Research Reviews* 29(1): 26-49.

Shirakawa, O., N. Kitamura, Lin, X. H., Hashimoto, T. and Maeda, K. (2001). Abnormal neurochemical asymmetry in the temporal lobe of schizophrenia. *Progress in Neuro-Psychopharmacology & Biological Psychiatry* 25(4): 867-877.

Simmons, L. W., Rhodes, G., Peters, M. and Koehler, N. (2004). Are human preferences for facial symmetry focused on signals of developmental instability? *Behavioral Ecology* 15(5): 864-871.

Sinha, P., Balas, B., Ostrovsky, Y. and Russell, R. (2006). Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE* 94(11): 1948-1962.

Snyder, T. J. (1991). Self-rated right left confusability and objectively measured right left discrimination. *Developmental Neuropsychology* 7(2): 219-230.

Sorger, B., Goebel, R., Schiltz, C. and Rossion, B. (2007). Understanding the functional neuroanatomy of acquired prosopagnosia. *Neuroimage* 35(2): 836-852.

Sperry, R. W. (1961). Cerebral organisation and behaviour – Split brain behaves in many respects like 2 separate brains, providing new research possibilities. *Science* 133(346): 1749-1753.

Stevenage, S. V. and. Osborne, C. D (2006). Making heads turn: The effect of familiarity and stimulus rotation on a gender-classification task. *Perception* 35(11): 1485-1494.

Sun, J. and Perona, P. (1998). Where is the sun? *Nature Neuroscience* 1(3): 183-184.

Sun, T. and Walsh, C. A. (2006). Molecular approaches to brain asymmetry and handedness. *Nature Reviews Neuroscience* 7(8): 655-662.

Sutherland, N. S. (1960). Visual-discrimination of orientation by octopus – Mirror images. *British Journal of Psychology* 51(1): 9-18.

Sutherland, N. S. (1961). Visual-discrimination of horizontal and vertical rectangles by rats on a new discrimination-training apparatus. *Quarterly Journal of Experimental Psychology* 13(2): 117-121.

Swaddle, J. P. and Cuthill, I. C. (1994). Preference for symmetrical males by female zebra finches. *Nature* 367(6459): 165-166.

Tan, U., Ors, R., Kurkcuoglu, M. and Kutlu, N. (1992). The lateralization of the grasp reflex in human newborns. *International Journal of Neuroscience* 62(1-2): 1-8.

Tanaka, J. W. and Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology* 46(2): 225-245.

Tanaka, J. W. and Sengco, J. A. (1997). Features and their configuration in face recognition. *Memory & Cognition* 25(5): 583-592.

ten Cate, C. (2002). Posing as professor: Laterality in posing orientation for portraits of scientists. *Journal of Nonverbal Behavior* 26(3): 175-192.

Thornhill, R. and Gangestad, S. W. (1994). Human fluctuating asymmetry and sexual-behavior. *Psychological Science* 5(5): 297-302.

Thornhill, R. and Moller, A. P. (1998). The relative importance of size and asymmetry in sexual selection. *Behavioral Ecology* 9(6): 546-551.

Trevarthen, C. (1996). Lateral asymmetries in infancy: Implications for the development of the hemispheres. *Neuroscience and Biobehavioral Reviews* 20(4): 571-586.

Troje, N.F. and Bülthoff, H.H. (1996), What is the basis for good performance to symmetric views of faces? *Investigative Ophthalmology & Visual Science* 37, 194-199.

Troje, N. F. and Bulthoff, H. H. (1998). How is bilateral symmetry of human faces used for recognition of novel views? *Vision Research* 38(1): 79-89.

Turnbull, O. H., Laws, K. R. and McCarthy, R. A. (1995). Object recognition without knowledge of object orientation. *Cortex* 31(2): 387-395.

Turnbull, O. H. and McCarthy, R. A. (1996). Failure to discriminate between mirror-image objects: A case of viewpoint-independent object recognition? *Neurocase* 2(1): 63-71.

Vaid, J. and M. Singh (1989). Asymmetries in the perception of facial affect – Is there an influence of reading habits. *Neuropsychologia* 27(10): 1277-1287.

Valentine, T. (1988). Upside-down faces – A review of the effect of inversion upon face recognition. *British Journal of Psychology* 79: 471-491.

Valentine, T. and V. Bruce (1988). Mental rotation of faces. *Memory & Cognition* 16(6): 556-566.

Ververs, I. A. P., Devries, J. I. P., Vangeijn, H. P. and Hopkins, B. (1994a). Prenatal head position from 12-38 weeks. 1. Developmental aspects. *Early Human Development* 39(2): 83-91.

Ververs, I. A. P., Devries, J. I. P., Vangeijn, H. P. and Hopkins, B. (1994b). Prenatal head position from 12-38 weeks. 2. The effects of fetal orientation and placental localization. *Early Human Development* 39(2): 93-100.

Vetter, T., Poggio, T. and Bulthoff, H. H. (1994). The importance of symmetry and virtual views in 3-dimensional object recognition. *Current Biology* 4(1): 18-23.

Vinette, C., Gosselin, F. and Schyns, P. G. (2004). Spatio-temporal dynamics of face recognition in a flash: it's in the eyes. *Cognitive Science* 28(2): 289-301.

Wagemans, J. (1997). Characteristics and models of human symmetry detection. *Trends in Cognitive Science* 1: 346–352

Warrington, E. K. and Davidoff J. (2000). Failure at object identification improves mirror image matching. *Neuropsychologia* 38(9): 1229-1234.

Webster, M. A. and MacLin O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review* 6(4): 647-653.

Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt II. *Psycologische Forschung*, 4, 301-350.

Wolff, W. (1933). The experimental study of forms of expression. *Character & Personality*, 2, 168–176.

Wolpert, L. (2005). Development of the asymmetric human. *European Review.* 13 (S2): 97-103.

Yarbus, A. L. (1967). Eye movements during perception of complex objects, in L. A. Riggs, (ed.) *Eye Movements and Vision.* Plenum Press, New York, chapter VII, pp. 171-196.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology* 81(1): 141-145.

Young, A. W., Hellawell, D. and Hay, D. C. (1987). Configurational information in face perception. *Perception* 16(6): 747-759.

Young, A. W., Hellawell, D. J. and Welch, J. (1992). Neglect and visual recognition. *Brain* 115: 51-71.

Yovel, G. and B. Duchaine (2006). Specialized face perception mechanisms extract both part and spacing information: Evidence from developmental prosopagnosia. *Journal of Cognitive Neuroscience* 18(4): 580-593.

Yovel, G. and N. Kanwisher (2004a). Face perception: Domain specific, not process specific. *Neuron* 44(5): 889-898.

Yovel, G., & Kanwisher, N. (2004b). Face perception engages a domain-specific system for processing both configural and part-based information about faces [abstract]. *Journal of Vision*, 4(8), 133a.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology* 9 (Supp.): 1-27.

Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science* 10(6): 224-228.

Zizak, D. M. and. Reber A. S (2004). Implicit preferences: The role(s) of familiarity in the structural mere exposure effect. *Consciousness and Cognition* 13(2): 336-36.