Changing Face-Space: Perceptual Narrowing in the Development of Face Recognition

Peter James Hills

School of Psychology

PhD Thesis submitted to Cardiff University, 2007

UMI Number: U584234

UMI
Dissertation Publishing
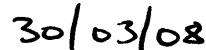
ProQuest

Declaration

This work has not previously been accepted in substance for any degree and is not being currently submitted in candidature for any degree. [1]

This thesis is a result of my own investigations, except where otherwise stated. Other sources are acknowledged by giving explicit references. A bibliography is appended.

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan and for the title and summary to be made available to outside organisations.

P. Hills

30/03/08

Peter James Hills   Date

---

[1] It should be noted that the following have been submitted for publication in peer-reviewed journals: Experiment 5 is in press at *Cognition* as Hills, Lewis, and Honey, 2007; Experiment 6 has been published in *Quarterly Journal of Experimental Psychology* as Hills and Lewis, 2006; Experiments 7 to 9 have been submitted as Hills and Lewis, 2007; Experiments 13 to 15 have been submitted as Hills, Elward, and, Lewis 2007; Experiments 16 and 17 have been submitted as part of Hills, Lewis, and Elward, 2007; Experiments 20 to 22 are in press at *European Journal of Cognitive Psychology* as Hills and Lewis, 2007; Experiments 23 and 24 have been published in *Perceptual and Motor Skills* as Hills and Lewis, 2007; Experiments 25 to 29 have been resubmitted to *Journal of Experimental Psychology: Human Perception and Performance* as Hills and Lewis, 2007.

Thesis Overview

The face-space metaphor for the encoding and storage of faces in memory has received a great deal of theoretical and empirical support. This metaphor is applied to the development of face recognition here. Three versions of the immature face-space are presented: the constant-face-space presumes that the number of dimensions of the immature face-space is the same as the adult face-space and all that changes is the distribution of faces within face-space; the expanding-face-space presumes that as more faces are encountered, new dimensions are added to the face-space to discriminate between highly similar faces; the shrinking-face-space suggests that the immature face-space contains many more dimensions than the adult face-space and, through processes akin to perceptual narrowing, most become dormant, leaving a default set of dimensions. These dormant dimensions are still contained within the face-space and can be activated under certain circumstances leading to a more flexible and dynamic face-space in adulthood.

Thirty-two experiments were conducted that aimed to discriminate between these three models of the immature face-space. The experiments presented in Chapter 2 explored the recognition and discrimination of: upright faces compared with inverted faces; own- and other- race, gender, and age faces; and faces with an "unnatural" facial configuration in adults and children aged 5- to 15-years-old. Chapter 3 used recognition tests and manipulated participants' attentional focus to explore what happens to unattended dimensions of face-space. Chapter 4 used adaptation procedures to recalibrate the dimensions of face-space over short and longer timeframes. In Chapter 5, the lower-level perceptual nature of dimensions of face-space is explored indirectly. In Chapter 6, the data reported are incorporated into a developmental model of shrinking-face-space, leading to an adult flexible-face-space.

.

Contents Page

8

List of tables

List of Figures

List of Abbreviations

ANOVA – Analysis of Variance
ERP – Event Related Potential
FDAE – Face Distortion After Effect
FIAE – Face Identity After Effect
FRU – Face Recognition Unit
HSD – Honestly Significant Difference
IAC – Interactive Activation and Competition model
IT – Inferotemporal
MSE – Mean Square Error
NRU –Name Recognition Unit
ORB – Own-Race Bias
PCA – Principal Component Analysis
PET – Positron Emission Tomography
PIN – Person Identity Nodes
SDT – Signal Detection Theory
SIU – Semantic Information Unit
VIQ – Visual Imagery Questionnaire
VRU – Voice Recognition Unit
WRU – Word Recognition Unit

Changing Face-Space: Perceptual Narrowing in the Development of Face Recognition

Chapter 1: Literature Review

Ellis (1986) suggests that face recognition takes its place at the acme of all human abilities surpassing all other human abilities and unrivalled by computers. When face recognition is modelled, researchers tend to indicate that a human adult can recognise 1,000 or 5,000 faces (Lewis, 2001) with some not having been seen in over 40 years (Bahrick, Bahrick, and Wittlinger, 1975). This expert ability (e.g., Farah, Wilson, Drain, and Tanaka, 1998) encompasses a wide array of perceptual, attentional, and cognitive abilities. The face can give an indication of identity, mood, age, gender, and health in a single brief glance (e.g., Bruce, 1982). Several influential models of this impressive ability have been proposed to explain adult face recognition. These models; however, have not been extensively applied to children.

The present thesis aimed to explore one particular model of face recognition, the face-space metaphor (Valentine, 1991). The adult version of this model and how it accounts for various established effects in adult face recognition will be described in some detail first. Then this background chapter will highlight how the adult model of face-space fails to account for face recognition effects observed in children. Three possible models of immature face-space shall be put forward and some predictions based upon them will be made. Furthermore, how these developmental models relate to an adult form of face-space will be described.

1.1. Face-Space

In 1991, Valentine devised the concept of face-space to represent how faces may be stored within memory. In this model, all encountered faces are stored within a multidimensional space, whereby each dimension represents a physiognomic characteristic used to encode faces. The dimensions of face-space are largely undefined but could include absolute featural information such as eye colour (e.g., Ellis, Deregowski, & Shepherd, 1975; Shepherd & Deregowski, 1981), absolute relational information such as distance between the eyes (e.g., Rhodes, 1988), and relative relational information such as the relative distance between the eyes compared with the length of the philtrum (e.g., Pedelty, Levine, & Shevell, 1985). Relative relational information as possible dimensions of face-space contradicts an implied assumption that dimensions of face-space are always positive and mutual exclusive values (Valentine, 1991). However, evidence for this mutually exclusivity is not extensive. Dimensions could also relate to grouping features such as secondary sexual characteristics which would group according to gender (e.g., Byatt & Rhodes, 1998; Rhodes, Hickford, & Jeffery, 2000). A principal component analysis (PCA) analysis has been used to show how gender might be a dimension of face-space by creating a set of Eigenfaces that can be used to describe the dimensions of face-space (c.f., Hancock, Burton, & Bruce, 1996; Turk & Pentland, 1991; Valentin, Abdi, Edelman, & O'Toole, 1997).

Valentine (1991) suggests that Euclidean distances are the best representation for the scales along each dimension of face-space. This has been a largely uncontentious assumption of face-space (but see Craw, 1995 for an alternative). Since face-space is based on Euclidean distances and the dimensions of the space represent biological information, the distribution of data points along each dimension is normal. The normal

distribution assumption is also largely unchallenged. Nevertheless, if Eigenfaces correspond to the dimensions of face-space, then it is possible that some dimensions of face-space are dichotomous, such as gender. Within face-space, gender could be considered a scale if visual characteristics are the only ones taken into consideration. In this way, the dimension for gender would be more visually defined as a dimension representing, for example, the size of the Adam's apple, coupled with length of stubble, and other secondary sexual characteristics (c.f., Bruce, Burton, Hanna, Healey, et al., 1993; Burton, Bruce, & Dench, 1993). In this way, dichotomous variables can be treated as a scale – however, they will be bimodal and not normally distributed.

As the dimensions of face-space represent biological information, the distribution along each dimension will be normal (with the exceptions described above). Thus, the distribution of faces within face-space as a whole will also be normal. There will be more faces clustered around the centre of face-space and these are the more typical faces (Valentine, 1991). More distinctive faces are stored further away from the centre in a region that has fewer faces in it. This description is not necessarily valid, when the typicality paradox is considered. The typicality paradox is whereby, when a series of faces are rated, there are few that are typical and few that are distinctive (Vokey and Read, 1992). Instead there is a normal distribution of ratings around the midpoint of the scale (Burton & Vokey, 1998). Thus, this suggests that there is not a cluster of normal faces around the middle of the face-space as Valentine (1991) suggests, at least in a one-dimensional face-space.

In addition to normality, biologically determined dimensions indicate that each dimension will have a different range of values along the scale. These differences are due to the fact that population variations for each facial feature differ from other features. For example, the distance between the eye line and the hair line in North American White

faces is typically 67.1 mm with a standard deviation of 7.1 mm (Wilson, Ferrara, and Yo, 1992); the height difference between the two eyes has a typical mean difference of 2.8 mm and a standard deviation of 1.1 mm (Hreczko and Farkas, 1994). People are covertly aware of these population deviation differences, stored as implicit knowledge according to Valentine (1991). Valentine did not explain whether this implicit knowledge is stored the same way as the explicit face memory or whether the dimensions of face-space act as the implicit knowledge. The latter seems most probable since the dimensions of face-space are set up to be most appropriate for the majority of faces encountered (Valentine & Endo, 1992). This implicit knowledge indicates that each dimension of face-space has a scale that accurately represents the physiognomy of typically encountered faces (McKone, Aitkin, & Edwards, 2005).

Valentine (1991) suggested two possible versions of face-space. One was based the idea that, from all the faces encountered, a prototype is extracted. Thus, each face is stored as a vector deviation from this prototype. There is only one prototype stored within face-space according to Valentine. However, Ellis (1981) has indicated that there may be more than one prototype within the face recognition system (e.g., a male and a female prototype). A recognition decision is based on a comparison process between the vector of the new face from the prototype and the nearest two vectors from the prototype. The second suggestion offered by Valentine is an exemplar based account whereby each face is stored as an exemplar. Similarity between faces is based on a monotonic function of distance between exemplars and recognition decisions are based on the distance to the nearest two exemplars. These two versions of face-space are depicted in two-dimensional form in Figure 1.1.
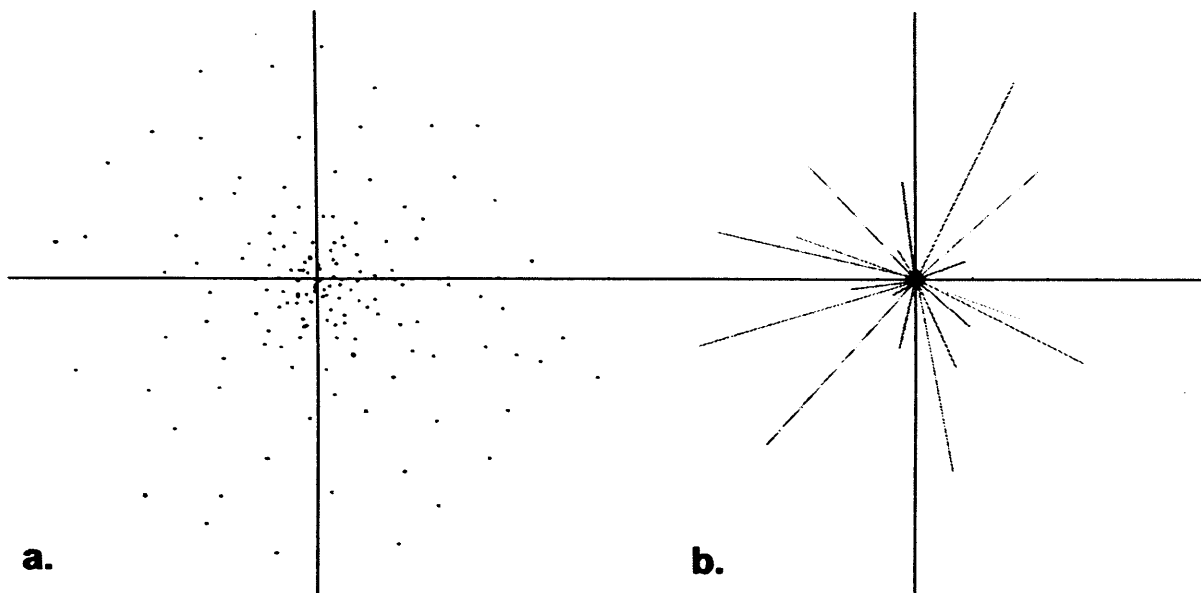
Figure 1.1. Two-dimensional representations of Valentine's (1991) original versions of face-space. a. The exemplar account - each face is represented by a point. b. The prototype account - each face is represented by a vector.

The exemplar face-space model was extended by Lewis (2004) and mathematically modelled using learning algorithms. Lewis suggested that the "strength" of each exemplar may not be equal and some are stored with insufficient strength for recognition. Ease of identification is based on activation strength and exemplar strength is directly related to experience with the face. Thus, a familiar face will be recognised more accurately. When a probe stimulus (a face) is presented, it is compared to stored exemplars. The probability of the probe stimulus being perceived as different to a particular stored exemplar is proportional to the difference in activation of the two stimuli (c.f., Thurstone, 1927). The final response is directly related to the strength of the original activation. Lewis implemented two assumptions into his face-space-r: the assumption of necessity, whereby the number of dimensions is not less than that required to recognise all faces; and the assumption of sufficiency, whereby there are not more dimensions of face-space than are required to recognise all faces. Reducing the available information

with which to process a face reduces the number of dimensions available to store the face. Lewis also indicated that new dimensions can be incorporated to distinguish highly familiar faces.

Combining the necessity and sufficiency assumptions leads to an implicit assumption of face-space-r that is all faces will have a value on all dimensions and that all dimensions are used equally. The sufficiency assumption states that there are no redundant dimensions and the necessity assumption indicates that all faces must be recognised by the dimensions within face-space. Together these assumptions lead to a state of affairs whereby removing a dimension of face-space will lead to a violation of the necessity assumption, or adding a dimension will violate the sufficicency assumption. Lewis (2004) puts it thus, "none of the dimensions are redundant... without violating the necessity assumption" (pp. 45). During all of his implementation, Lewis always used all dimensions within the model equally. Thus, all dimensions will be used and that all faces will be represented on all dimensions and that – by implication from the modelling – all dimensions are used equally.

Having described the theoretical infrastructure, a more detailed discussion of the perceptual and cognitive processes that occur when a recognition decision is made shall be made, in order to fully appreciate the face recognition system and the subtleties of face-space. To make a recognition decision, the face has to be seen in the visual surrounding and classified as a face. The amount of detailed information that a face contains must then be filtered, such that only the useful properties are attended to. This percept must then be compared to faces stored within memory and finally a decision has to be made. Each of these stages shall be discussed in detail.

To perceive a face, one must separate the face from the visual scene. The visual scene is made up of light of different wavelengths and intensities. Patches of light next to one another have different contrasts, spatial frequencies, contour boundaries, and so on that differentiate between objects and backgrounds. From the incoming light, the human visual system is able to detect a face and extract its identity. Multiple brain regions are involved in this process.

The human visual system is hierarchically organised, such that the lower visual areas process a smaller visual field and simpler stimuli than higher visual areas (Ito, Tamura, Fujita, and K. Tanaka, 1995). Cells in the primary visual cortex (Area V1, the striate cortex) have distinct responsive patterns, whereby cells will selectively respond to, for example, a particular edge (simple cells), a moving edge (complex cells), or moving edges of a particular shape (hypercomplex cells) (see de Valois, Yund, and Helpler, 1982). Low level areas of the visual cortex process selected classes of stimuli: Area V2 contains cells that respond to local geometric shapes; Area V3 is considered to process form and local movement (Zeki, 1993); Area V4 processes colour (e.g., Zeki, 1973, 1977) and contains cells with relatively large receptive fields that respond to combinations and relative positions of stimuli (e.g., Gallant, Connor, Rakshit, J. Lewis, and van Essen, 1996; Hanazawa & Komatsu, 2001; Pasupathy & Connor, 1999, 2001). Cells in these regions link to the inferotemporal (IT) cortex and on to the superior temporal cortex such that object perception can occur.

Cells in the monkey superior temporal cortex have been found that selectively respond to faces (C. Bruce, Desimone, and Gross, 1981; Perrett and Mistlin, 1987). At this high-level, neural tuning curves have been reported for variations in global face

attributions such as: hair and facial outline (M. Young and Yamane, 1992), expression (Hasselmo, Rolls, and Baylis, 1989), and three-dimensional representations (G. Wang, K. Tanaka, and Tanifugi, 1996). In humans, functional imaging studies have indicated the fusiform gyrus (a small region within the superior temporal cortex) to be one key locus for face processing (Aguirre, Zarahn, and D'Esposito, 1999; Gauthier, Tarr, Moylan, Skudlarski, Gore, and Anderson, 2000; Haxby, Hoffman, and Gobbini, 2000; Kanwisher, McDermott, and Chun, 1997; Kanwisher, Stanley, and Harris, 1999; McCarthy, Puce, Gore, and Allison, 1997). It is active when processing faces but not when processing objects. Moreover, the fusiform gyrus is significantly more active when processing upright faces than inverted faces (Gauthier, Tarr, Anderson, Skudlarski, and Gore, 1998; Haxby, Ungerleider, Clark, Schouten, Hoffman, and Martin, 1999; Kanwisher, Tong, and Nakayama, 1998; Rossion and Gauthier, 2002; Yovel and Kanwisher, 2004). This indicates that face processing takes place in the fusiform gyrus amongst other areas. However, this region is also active for the visual processing of cars and birds for experts in these classes of objects (Gauthier, Anderson, Skudlarski, and Gore, 2000). Indeed, people can be trained to recognise "Greebles" (novel objects that are complex and require fine detail processing) and this causes the fusiform gyrus to be active (Gauthier, Williams, Tarr, and Tanaka, 1998). As such, the fusiform gyrus appears to process any visual object that the human has expertise in processing even though some researchers maintain that the fusiform gyrus is specialised for faces. Recent evidence has suggested that the fusiform gyrus may actually be multi-modal and can be activated by auditory input (Kung, 2007).

Having established that there is a brain region that is associated with the expert processing of faces, it is important to have an understanding of what aspects of the visual stimulation are processed here. The human visual system filters out 90% of the incoming

light (c.f., Moran and Desimone, 1985; Spitzer, Desimone, and Moran, 1988) through a processing of attentional filtering. This filtering occurs throughout the visual cortex, such that not the entire facial image is being processed by the fusiform gyrus. To have some idea of what is being processed, some perceptual empirical studies concerning spatial frequencies shall be described.

There has been a great deal of evidence suggesting the importance of certain spatial frequency channels in face recognition. A multitude of methods have been used to determine the critical spatial frequencies involved in the identification of faces, including selectively removing spatial frequency channels until identification is near impossible. These studies have indicated that spatial frequencies between 8 and 16 cycles per face are most critical in recognition of faces (Bachmann, 1991; Costen, Parker, and Craw, 1994, 1996; Fiorentini, Maffei, and Sandini, 1983; Ginsburg, 1978, 1986; Gold, Bennett, and Sekuler, 1999; Konorski and Petersik, 2003; Näsänen, 1999; Parker and Costen, 1999; Tieger and Ganz, 1979), though some suggest spatial frequencies of 25 cycles per face are crucial to recognition of faces (Hayes, Morrone, and Burr, 1986).

The involvement of spatial frequency channels in face recognition is not as simple as that, however. Bruce (1988, see also Goffaux, Hault, Michel, Vuonq, and Rossion, 2005; Schyns and Oliva, 1999) suggests that low spatial frequencies carry configural information and high spatial frequencies carry featural information (for a discussion of configural and featural information, see section 1.2.1.). Moreover, the spatial frequency channels involved in the identification of faces are higher than those used in gender classification (Schyns, Bonnar, and Gosselin, 2002, see also Schyns and Oliva, 1999). This difference is borne out in Event Related Potential (ERP) studies also, whereby the face selective ERP N170 responds according to the spatial frequencies required by the task, in that for identification it reacts to high spatial frequencies, but for gender

classification it reacts to low spatial frequencies (Goffaux, Gauthier, and Rossion, 2003). Indeed, tasks involving holistic face perception involve lower spatial frequencies than tasks involving featural face perception (Goffaux and Rossion, 2006). Furthermore, high-frequency spatial information appears to be used at the early stages of face perception (Halit, de Haan, Schyns, and Johnson, 2006), in contrast to low-frequency spatial information, which is used at latter stages of face perception.

Having suggested the importance of certain spatial frequency channels, it must be noted that, absolute spatial frequencies are less important than the differences in spatial frequencies seen in the face at learning from those seen in the face at test. If the spatial frequencies in the test faces do not match those learnt by a single octave, recognition performance is reduced by 20% (Liu, Collin, Rainville, and Chaudhuri, 2000). This matching of spatial frequencies at learning and at test is more important for faces than objects and is similar for upright and inverted faces (Collin, Liu, Troje, McMullen, and Chaudhuri, 2004). Indeed, recognition of faces missing the critical spatial frequency band is possible at the same level of accuracy as those containing them provided that the faces are learnt and tested with the same spatial frequencies (Konorski and Petersik, 2003).

From the neurological and perceptual research summary presented here, some idea of the neurological locus of face-space and what its dimensions are can be made. The visual system filters out a great deal of irrelevant information and this differs according to the task at hand. This visual filtering is similar to the idea that face-space only contains dimensions that are most diagnostic for the most frequently encountered faces. It is, also, possible to conceive the dimensions of face-space as a collection of spatial frequencies if one considers Fourier analysis (c.f., Graham, 1989; Weisstein, 1980). Irrespective of this, it is important to understand how spatial frequency analysis impacts on face recognition and how these may relate to the dimensions of face-space.

22

1.1.2. The Number of Dimensions of Face-Space

Several attempts have been made to estimate the number of dimensions of face recognition and in particular face-space using different computational methods. Using factor analysis of participants' descriptions of facial features, Shepherd, Ellis, and Davies (1977) found that 11 factors accounted for approximately 60% of the total variance between the 100 faces they examined. Turk & Pentland (1991) used PCA to assess dimensions of recognition and estimated that there were 50 dimensions useful for recognition.

Burton and Vokey (1998) assessed the dimensions of face-space more specifically than the dimensions of face recognition using the typicality paradox, as described earlier. Recall that when faces are rated for how typical or distinctive they are, few faces are rated as typical or distinctive (a normal distribution) (Vokey and Read, 1992), yet face-space predicts that there will be more typical faces (a skewed distribution). This paradox is true for lower numbers of dimensions but not for higher numbers of dimensions. Burton and Vokey (1998) resolved the paradox by mathematically showing that as the number of dimensions increases the skew decreases. Normality is approximately reached between 10 and 200 dimensions if all the dimensions are used equally.

More recently, Lewis (2004) used a different method for assessing the dimensions of face-space. He modelled face recognition employing the necessity and sufficiency assumptions described earlier. Lewis used a computational model, adjusting the number of dimensions until errors were made. This gave an estimate of between 15 and 22 dimensions of face-space. Thus, these converging lines of evidence suggest there are between 10 and 200 dimensions of face-space that may be used equally (Lewis, 2004;

Vokey and Read, 1992) or not used equally (Valentine and Endo, 1992). The next section looks at how face-space is used to make recognition decisions.

1.1.3. Recognition Decisions

The basic recognition decision will be based on the magnitude of the response to the probe exemplar (Lewis, 2004). If the comparison process between the new face and the stored face indicates that two faces share the same values, within an error range, then the decision will be made that the face is a target (a known face). Error ranges around each exemplar ensure that some recognition errors occur (Young, Hay, and A. Ellis, 1985). When a face is encountered, and a recognition decision is made, this can either be: "I know the face" or "I do not know the face". This response can either be correct or incorrect leading to a 2 by 2 matrix of possible outcomes of a simple recognition decision. These possibilities are: a correct recognition (*hit* in Signal Detection Theory – SDT – terms, e.g., Swets, 1966); a correct no recognition (*correct rejection*); falsely recognising someone (*false alarm*), or a failure to recognise someone (*miss*). Since there are two types of error that are possible, the perceiver can make a cognitive decision to make one type more than the other by shifting their response criterion (e.g., Balakrishnan, 1998).

Response criterion is defined as the criterion set to make a recognition decision. Participants with a conservative response criterion will only say they have seen a face before if they are absolutely sure. Participants with a liberal response criterion will say they have seen a face before even when they are not sure. Experimental instructions and participant individual differences are likely to play a part in response criterion (c.f., Balakrishnan, 1999).

24

Faces not only convey information for which a recognition decision needs to be made – they also contain visual signals regarding affective and semantic information. Some of this is based purely on visual signals, such as a smile indicating positive affect. Others, however, are due to stored knowledge interacting with the visual input. As such, it is important to explore how semantic information is linked to visual cues in the face recognition system.

It is well established that the conditions under which events are encoded are directly related to the success with which they are subsequently recalled (for a review see Coin and Tiberghein, 1997). Encoding conditions have been shown to exert a greater influence on successful recognition than retrieval conditions (c.f., Bruce, 1998). Winograd (1981) assessed face recognition after participants had made one of nine judgements during the initial encoding of a face. Participants had been instructed to rate faces on one attribute that could either pertain to their physical characteristics (big nose, straight hair, and heavy), various abstract traits (intelligent, anxious and friendly), or occupations (actor, businessman, and teacher). Recognition accuracy was poorer when physical judgements were made about the faces (e.g., big nose or straight hair) than when abstract judgements were made (e.g., intelligent or teacher; see also, e.g., Light, Kayra-Stuart, and Hollander, 1979; Mueller, Carlomusto, and Goldstein, 1978).

There are several potential explanations for the different effects of making physical as opposed to abstract judgements about a novel face on the accuracy of subsequent face recognition. For example, the Levels of Processing framework (Craik and Lockhart, 1972) assumes that the more deeply an item is processed, the better it is recalled. Within this framework, making a physical judgement about a face during encoding should result in poorer subsequent performance than making an abstract judgement. However, other research has shown that selecting the most distinctive feature

in a face during encoding (a surface judgement) results in similar recognition accuracy to making an abstract judgement during encoding (Daw and Parkin, 1981; Deffenbacher, Leu, and Brown, 1981). This finding suggests that, at least for the case of faces, the influence of abstract judgements on recognition might not simply reflect the depth with which they are processed. In contrast to the depth of processing analysis, Courtois and Mueller (1979) argued that the critical factor determining face recognition accuracy was the *number* of facial features assessed during encoding. According to this view, making either a distinctive-feature judgement or an abstract judgement results in a greater number of features being processed during encoding and it is this fact that supports greater recognition accuracy. However, there is no direct evidence that making abstract judgements about a face either requires or results in a greater number of features being processed than does making a gender classification (Kerr and Winograd, 1982). In fact, the evidence concerning the reaction times to make these differing judgements is inconsistent: Bloom and Mudd (1991) found that "honesty" judgements took longer than gender classification, whereas Daw and Parkin (1981) found the opposite pattern of results using remarkably similar methods.

It is clear that making an abstract judgement about a face somehow increases the readiness with which it is encoded. This kind of influence has also been demonstrated by Shepherd, Ellis, McMurran, and Davies (1978) using a quite different method of assessment. In their study, Shepherd et al. (1978) asked participants to create a Photofit from a picture they had seen previously. Half of their participants were told during presentation of the picture that the face belonged to a convicted murderer; whereas the other half were told that the face was that of a lifeboat captain. The constructed Photofits were then rated for traits such as intelligence, humorousness, sociability, and attractiveness by another set of participants, who were unaware of the stereotype labels

that had originally been provided. The Photofit of the 'lifeboat captain' was rated as more attractive and more intelligent than the Photofit of the 'criminal'. These results can be taken to suggest that occupation labels presented during encoding altered the way a face is encoded and/or subsequently retrieved (in this instance for reconstruction).

Other research has revealed an influence of the provision of semantic information during encoding that appeared to be much more selective. It depended on whether or not the presented face was congruent with the stereotype of the occupation label that accompanied it. Thus, Klatzky, Martin, and Kane (1982a, b), using a recognition paradigm, provided evidence that occupation labels given by the experimenter influences subsequent face recognition. Klatzky et al. collected a set of faces that were reliably classified as having one of 11 occupations (e.g., accountant, hairdresser, actor, and musician; c.f., Bull and Green, 1980); there were 11 sets of 4 such pictures. The participants were presented with a subset of the 44 faces. The faces were accompanied during encoding with an occupation label that was either stereotypically congruent or incongruent with that face. Participants given the congruent occupation label recognised more faces in the subsequent recognition test than did those who had received an incongruent label. Because the occupation labels were also presented at test, this effect was evident in participants' propensity to say "yes" both to occupation-congruent targets and to occupation-congruent distractors (false hits). An analysis of d' (SDT measure of stimulus dsicriminability) demonstrated that sensitivity was higher for occupation-label congruent faces, even accounting for the increase in false hits. However, in a recognition test where the ratio of targets to distractors is greater than one to one, as it was in Klatzky et al. (1982a), response bias is likely to be increased (c.f., Balakrishnan, 1998). It is, therefore, possible that the effect reported by Klatzky et al. reflects response bias rather than accuracy.

The results described here indicate that semantic information may interact with visual input. As face-space is primarily a metaphor for the visual representations of stored faces, it fails to account for semantic effects in face processing. For this the Interactive Activation and Competition (IAC) model (e.g., Burton, 1994) can be considered.

1.1.4. Beyond Face-Space: Interactive Activation and Competition

The IAC model is based upon neural network architecture (c.f., Grossberg, 1978; McClelland, 1981), whereby it comprises processing units organised into hierarchically arranged pools. Each unit is connected to other units within the same pool by inhibitory links which are all of equal strength (Burton and Bruce, 1993), but less strong than excitatory links between units that cross pools (Burton, et al., 1999). Each unit takes on an activation level between fixed maximum and minimum values which are updated with each processing cycle. The model is stabilised by a global decay function which drives all units to a resting level of activation. All the links are bi-directional (Burton, Bruce, and Johnston, 1990). Activation of units can be external or internal. External input is through processing something in the real world, whereas internal input is activation spreading from other units. The activation level of a particular unit is proportional to the product of the strength of activation of the connected units and the connection strengths. The ease of recognition is based upon the number of processing cycles required to reach the threshold value (Burton et al., 1990).

The IAC contains a number of pools that represent the different aspects of the face recognition system. At the lowest level is the visual input. Though generally underspecified, the perception and coding of a face is done here. Following the visual input level is the face recognition pool, whereby nodes representing all the faces within

memory are stored. There is a Face Recognition Unit (FRU) for every face that is stored within memory, and each is purely visual. It may be an average for all the many views of a particular person, or it may contain all possible views encountered of that person. Burton, et al. (1990) suggest that the FRUs store the visual structural descriptions of faces which allow views of one known face to be distinguished from views of other faces, known or otherwise. Some researchers have indicated that the FRU pool is best represented by the face-space (e.g., Valentine, Chiroro, and Dixon, 1995). Connected to each FRU, through Person Identity Nodes (PINs), is person-specific information in terms of many Semantic Information Nodes (SIUs). These nodes provide all the semantic information concerning that identity. Each aspect represents a different characteristic of someone's identity connected to the visual face store. A basic architecture is shown in Figure 1.2.

The original IAC has been developed and combined with a perceptual front-end based upon PCA (Burton, Bruce, & Hancock, 1999). The PCA mechanism is connected to the FRU pool. The PCA analysis is based upon shape-free transformations of faces, whereby the faces used for subsequent processing all have the same shape. It creates a series of Eigenfaces that are compared to known FRUs. Each face has a value on the Eigenfaces that give its FRU a PCA signature. This computational model incorporating a perceptual front-end is an advancement on the original IAC (Burton, et al., 1990). The weighting between nodes between pools within the original IAC does not necessarily have to be equal, nor do the connexions between the FRUs and the PCA front-end. The PCA front-end may operate in a similar manner to face-space given that an Eigenface appears to correspond with a dimension within the face-space.
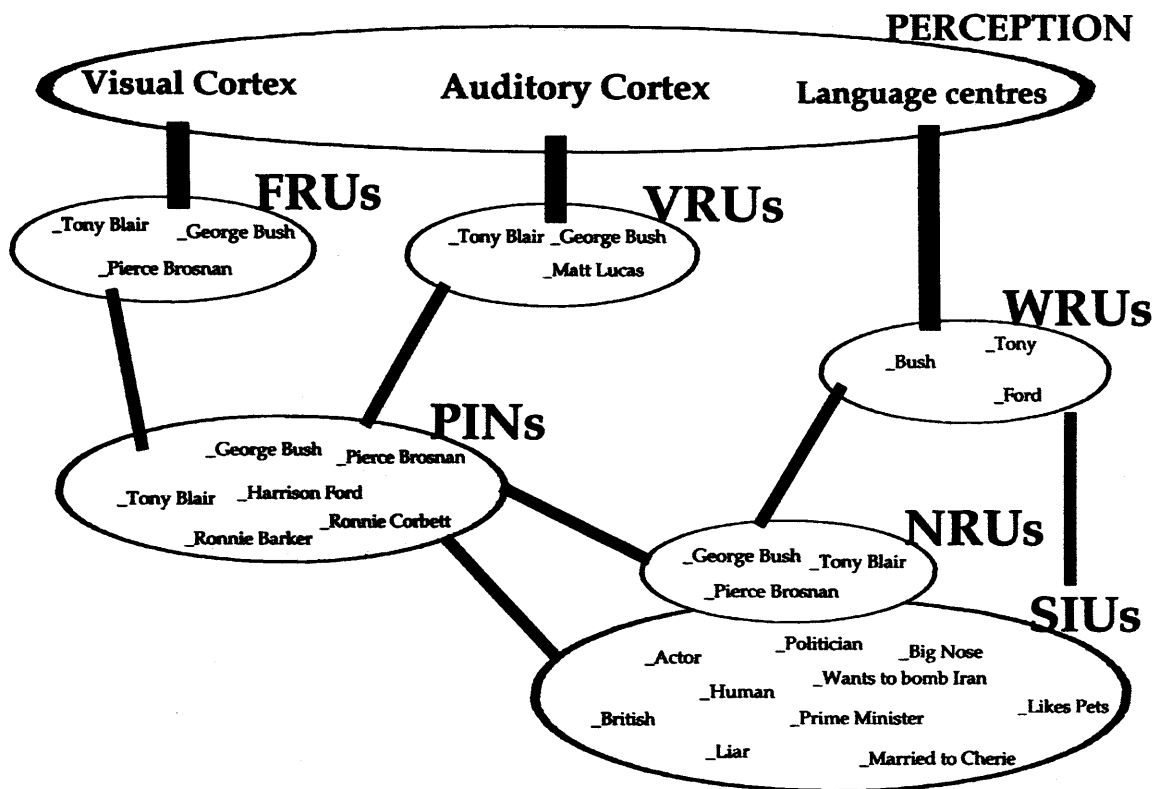
Figure 1.2. A schematic representation of the IAC based upon Burton and Bruce (1993). The perceptual front-end of the IAC links to the FRUs, VRUs and Word Recognition Units, which link to the PINs (through the NRUs for WRUs). The PINs are also connected to the NRUs and the SIUs. All the links shown are bi-directional.

Face-space, incorporating the IAC, can account for a vast majority of the basic effects in face recognition. It offers a neat account for how faces are detected and encoded, and this is supported by neurological evidence. It explains the comparison process with other faces, and how decisions can be reached based the magnitude of activation within the model. Nevertheless, there are many unanswered questions about face-space. In particular, the number and types of dimensions of face-space require further definition. There is also debate as to whether all the dimensions are used to recognise all faces as Lewis (2004) implies, or whether a subset can be used under certain circumstances

(Valentine and Endo, 1992). Either is plausible. To make an informed suggestion, empirical work into the main effects of face recognition shall now be presented.

1.2. Key effects in face recognition

Having discussed the basic structure of adult face-space, it is important to explain how this model accounts for various established effects in face recognition that will be tested in subsequent experimental chapters. The effects of interest are the inversion effect, own-group biases, and adaptation.

1.2.1. The Inversion Effect as Shift in Processing Style

The inversion effect is where inverted stimuli are harder to recognise than their upright counterparts. Faces are more affected by inversion than other classes of stimuli, such as aeroplanes, houses, and visual scenes (e.g., Hochberg & Galper, 1967; Leder & Carbon, 2005; Yin, 1969, 1970). The inversion effect is so robust that it does not interact with familiarity of the face, whether the familiarity is experimentally manipulated (Scapinello & Yarmey, 1970) or based on celebrity (Yarmey, 1971). The effect inversion has on the encoding of faces has been theorised many times. In the face recognition literature, there is a distinction between featural and configural encoding (see Pellicano, Rhodes, & Peters, 2006; Rhodes, Brake, & Atkinson, 1993; Tanaka & Sengco, 1997). Featural encoding is usually defined as piecemeal processing of individual features in isolation (e.g., Bartlett and Searcy, 1996; Tanaka, Kay, Grinnell, Stansfield, & Szechter, 1998). Configural encoding is defined as processing the relations between the features in a whole (e.g., Thompson, 1980; Young, Hellawell, & Hay, 1987). Holistic encoding may

be considered as either configural encoding or as both featural and configural coding combined (Farah, et al., 1998; Searcy and Bartlett, 1996). Face recognition is based primarily on configural encoding (e.g., Bruce, Doyle, Dench, & Burton, 1991). All faces share the same basic arrangement, which Rhodes (1993) terms first-order relational information. Second-order relational information, according to Rhodes (1993) is the small idiosyncrasies in the relationships between features, so long as they maintain the basic facial configuration. These definitions have not been consistently agreed upon within face recognition and are largely unspecified. For reviews, see Maurer, Le Grand, and Mondloch (2002) and Rossion and Gauthier (2002).

According to many researchers, inverting a face will disrupt configural information (Freire, Lee, & Symons, 2000). This is certainly true of first-order relational information, whereby an inverted face has a mouth above the eyes. However, the second-order relational and featural information is not affected by inversion. Lewis and Glenister (2002) define configural encoding as that "which is disrupted by inversion" (p. 8) and featural encoding as that "which can occur for inverted faces" (p. 8).

Although inverting a face causes it to be subsequently recognised more poorly than an upright face, an effect of inversion is not always found in face-matching tasks (e.g., Bruyer and Velge, 1981; Valentine, 1986). However, inversion does disrupt how attention is distributed over a face (Davies and Hoffman, 2002). Using a test of cue saliency, Endo (1982; 1986) demonstrated that changes to the least salient facial features (nose, mouth, and chin; see Ellis, Deregowski, and Shepherd, 1975) were disproportionately harder to identify than changes to the most salient facial features (hair and eyes, see Shepherd and Deregowski, 1981). The pattern of performance was in the same direction for upright and inverted faces (see also Endo, 1983). If the changes made

to the face are cued, they are easier to detect, irrespective of the orientation of the face (Barton, Deepak, and Malik, 2003).

Face recognition is characterised by a distinct attentional process. In general, when presented with a face, a distinct scan-path as measured by an eye-tracker is observed (e.g., Cook, 1978). This scan-path is associated with a first glance to the eyes (e.g., Stacey, Walker, and Underwood, 2005), followed by many prolonged fixations around the eyes and forehead and few shorter fixations to the mouth (e.g., Althoff and Cohen, 1999). This distinct face-scanning pattern is especially evident for novel faces, and less so for highly familiar faces. Inverted faces, by contrast, have a more random scan path starting at the mouth (Barton, Radcliffe, Cherkasova, Edelman, and Intriligator, 2006). Thus, the scan-path for an inverted face has an incorrect first glance and fixates on less diagnostic features.

Another method for potentially examining the type of coding used to process a face is with verbalisation. First studied by Schooler and Engstler-Schooler (1990), verbal overshadowing is where describing a face in the interval between witnessing an event (in their case a mock crime) and making a recognition decision (a line up) disrupts recognition. Those participants who have verbalised a face are less accurate in their recognition of the face than those who did not verbalise. Schooler and Engstler-Schooler (1990) suggested that the perceptual memory of the face has been overshadowed by a verbal memory. This verbal memory causes a mismatch between perceptual and verbal knowledge and is known as the modality mismatch hypothesis (Schooler, Fiore, & Brandimonte, 1997). Alternatively, a transfer-inappropriate processing shift explanation has been put forward as described below.

When people verbalise faces, it is often the case that they describe features and use a feature-matching strategy (Dodson, Johnson, & Schooler, 1997). Since face

recognition is primarily a configural process as described above, verbalisation causes faces to be processed featurally. This is a non-optimal process (c.f., Roediger, 1990) and does not match the encoding conditions of the original face. This mismatch reduces recognition performance.

Fallshore and Schooler (1995) indicate that the verbal overshadowing effect will only occur for tasks involving a high degree of perceptual expertise, which explains why no verbal overshadowing effect was found for car memory (Brown & Lloyd-Jones, 2003). Fallshore and Schooler also noted that verbalising an inverted face causes it to be recognised more accurately. Since an inverted face is best processed featurally, the verbalisation has increased the amount of featural information to be used. This would, therefore, be a transfer-*appropriate* shift. Nevertheless, the verbal overshadowing effect has proven difficult to replicate and several authors have failed to find an effect of verbal overshadowing (e.g., Lyle & Johnson, 2004; Memon & Bartlett, 2002). A meta-analysis conducted by Meissner and Brigham (2001a), revealed the effect size of verbal overshadowing was small, around $r = .12$.

An apparently similar effect was reported by Macrae and H. Lewis (2000). These authors found that during the interval between watching a video of a crime and a line-up, processing the global figure in a Navon letter improved accuracy in the line-up. A Navon letter (Navon, 1977) is a global figure made up of local features. Whereas control participants (who read a magazine) had an accuracy of 63% in the line-up, participants asked to identify the global letters for ten minutes had an accuracy of 83%. Furthermore, participants asked to identify the local letters for ten minutes had reduced performance, at an accuracy level of 30%. Processing the local features of a Navon letter depreciates recognition accuracy by 33% in this study.

This pattern of results has been subsequently replicated by Perfect (2003) but not by Brand (2005), Lawson (2006), or Ryan, Aulenbach, Becker, Johnson, Christman, Murray, and Bergmaier (2006). Indeed, Brand (2005) indicates that the effect size of this finding is small, roughly r = .10. Nevertheless, Macrae and H. Lewis (2001) and Perfect (2003) have hypothesised that processing the local features of a Navon stimulus causes local processing to be transferred to the subsequent face recognition task. This is, thus, a transfer-inappropriate processing shift.

There are two key effects regarding the way Navon stimuli are processed. The first is that the global shape is identified quicker than the local shapes (Navon, 1977). This is the global precedence effect. There is also an interference effect, whereby processing the global component of a Navon figure causes the response time for a subsequent local feature to be significantly greater. These effects are malleable, since changes in the size (Kinchla and Wolfe, 1979) and spatial frequency (Badcock, Whitworth, Badcock, and Lovegrove, 1990; Hughes, Nozawa, and Ketterle, 1990; La Gasse, 1993; Lamb and Yund, 1993; 1996a, b; Robertson, 1996; Sierra-Vazquez, Serrano-Pedraza, and Luna, 2006) of the Navon stimuli can remove or reverse the global precedence effect. It must be noted that not all changes to Navon stimuli affect the global precedence effect - colour, polarity, or contrast appear not to have an effect on the global precedence effect.

The processing of local and global information has been identified in neuroimaging studies. On average, the left hemisphere processes featural information and the right hemisphere processes configural information (Delis, Robertson, and Efron, 1986; Lamb, Robertson, and Knight, 1989; 1990; Robertson, Lamb, and Knight, 1988; Robertson, Lamb, and Zaidel, 1993). This correlates with research on high and low spatial frequency processing, which suggests the left and right hemisphere dominate

processing for each respectively (Christman, Kitterle, and Niebauer, 1997; Kitterle and

Selig, 1991). A direct test of this correlation was performed by Han, Yund, and Woods

(2003) in which ERPs were recorded for the processing of local and global targets in the

left and right visual fields that were either high- or low-pass filtered. The ERP Nd190

usually observed when processing global targets was eliminated when the targets had

their low spatial frequencies removed. This suggests a direct link between spatial

frequency channels and the type of processing of Navon stimuli. Similar results have

been observed in fMRI and PET studies (Peyrin, Chokron, Guyader, Gout, Moret, and

Marendaz, 2006), whereby activation for low spatial frequencies and global targets is in

the right lateral areas of the occipital cortex and activation for high spatial frequencies

and local targets is in the left medial areas of the occipital cortex. Boeschoten, Kemner,

Kenemans, and van Engeland (2005) demonstrated that the removal of high spatial

frequencies in Navon stimuli decreased identification performance of local but not global

targets. Conversely, removal of low spatial frequencies decreased identification

performance for global but not local targets.

Shulman, Sullivan, Gish, and Sakoda (1986) demonstrated that adaptation to low

spatial frequency sine-wave gratings prior to identifying global figures in Navon stimuli

made such identification significantly more difficult. Furthermore, following the

processing of the global component of a Navon stimulus, high spatial frequency sine-

wave gratings were more easily detected than low spatial frequency sine-wave gratings

(Shulman and Wilson, 1987).

The inversion effect may well be caused by a disruption to configural encoding

when a face is inverted. It is not the only method to disrupt configural processing,

causing similar detriments in face recognition performance. The research presented here

indicates the importance of spatial frequency information in configural and featural

processing, however, spatial frequency correlates with type of processing is not necessarily a one-to-one relationship.

1.2.2. The Own-Group Biases

The own-race bias (ORB) is characterised by the ability to recognise faces of one's own culture far more readily than those of another culture (e.g., Brigham & Malpass, 1985; Brigham & Williamson, 1979; Chance & Goldstein, 1986; Chance, Goldstein, & McBride, 1975; Leippe, 1995; Shepherd, Deregowski, and Ellis, 1974). In a meta-analysis, Meissner and Brigham (2001b) reported that people are 2.23 times more likely to successfully recognise own-race faces than other-race faces. The ORB is not due to one race of face being more homogenous than another race, as borne out by pretests involving own-race participants (see e.g., Goldstein and Chance, 1976). Moreover, there is usually no association between racial attitudes and the magnitude of the ORB (e.g., Galper, 1973; Malpass, 1981; Seeleman, 1940, but see Brigham and Barkowitz, 1978; Brigham and Ready, 1985; Walker and Hewstone, 2006).

Several results indicate that contact with members of another race reduces the ORB. For example, Brigham, Maass, Snyder, and Spaulding (1982) found that self-reported amount of experience with other-race faces correlates with level of ORB, though this is not always found (Brigham and Barkowitz, 1978; Luce, 1974a, b). Black participants (see Anthony, Cooper, and Mullen, 1992; Hosch and Platz, 1984; Malpass and Kravitz, 1969) and Japanese participants (Valentine and Endo, 1992) show a lower ORB than White participants (but see Bothwell, Brigham, and Malpass, 1989, for an alternative result). A possible explanation is that Black and Asian participants have more contact with White people in their everyday lives (on television, and at work), whereas

White participants can avoid Black people quite easily due to some degree of racial segregation still present in many Western countries (c.f., Valentine and Endo, 1992). Indeed, the magnitude of the ORB has been found to be smaller in White children brought up in desegregated neighbourhoods than White children brought up in segregated areas in the US (Cross, Cross, and Daly, 1971). The same results were not found for Black children. Moreover, children in integrated schools showed a lower ORB than those in segregated schools (Feinman and Entwhistle, 1976).

To test the idea that the ORB is dependent on contact with faces of another race, several authors have provided training paradigms in experimental settings to reduce the ORB. Malpass, Lavigueur, and Weldon (1973) reduced White participants' ORB levels in a visual training paradigm. Their participants were given feedback in each trial of a four-alternative forced-choice recognition test. Training improved recognition scores for both own and other-race faces, especially if the participants were given shock feedback. However, Devine and Malpass (1985) found that they could not reduce the ORB by encouraging participants to process other-race faces more deeply. Moreover, the ORB is not reduced by providing reward for accurate recognition (Barkowitz and Brigham, 1982). These results indicate that experimental training and exposure can reduce the ORB in some cases, but only to a small degree (Slone, Brigham, and Meissner, 2000).

Malpass (1974) has indicated that rather than simple exposure to other-race faces, the type of individuating processing engaged with the other-race people will result in smaller ORB levels. Indeed, a training experiment was conducted by Lavrakas, Buri, and Mayzner (1976), in which White participants were given a concept learning visual training task. The training faces were made from an Identi-Kit and the two concepts participants had to learn were either light eyes or a conjunction of dark eyes and thick lips. Participants were given feedback for their judgement. Lavrakas, et al. also tested

whether the type of contact and field-dependence/independence of the participants influenced the magnitude of their ORB. Their results indicated that those White people who had Black friends had a small ORB than those who simply knew more Black participants. Moreover, field-independent White participants were better able to recognise Black faces than field-dependent White participants. Finally, the training on specific features did reduce the ORB, possibly due to removing the deleterious effect of remembering only skin-colour (Lavrakas et al.).

These results are consistent with a recent theoretical account of the ORB. Levin (2000) advocates a feature-selection process in the recognition of faces. By this, participants extract the most diagnostic visual feature to aid in recognition. As such, processing an own-race face, the most diagnostic visual feature could be any number of individuating information. However, processing an other-race face involves a race feature as the most diagnostic according to Levin (2000). As such, an other-race face is processed according to a prototype and therefore not individuated, thus, making it more difficult to recognise.

This explanation is related to an explanation of the ORB based on the face-space metaphor. Other-race faces do not vary consistently along the same dimensions as own-race faces, due to the physiognomic differences between Black and White faces (e.g., Shepherd & Deregowski, 1981). Thus, other-race faces are grouped close together, but farther away from the centre of face-space. This makes them more difficult to distinguish from each other, but easy to differentiate from own-race faces (Valentine & Endo, 1992). Nevertheless, it has not been explored directly whether the dimensions of face-space for own-race faces are different from other-race faces. Though physiognomic differences between Black and White faces exist, few studies have looked at whether Black and White people use different physiognomic characteristics to process faces.

Ellis, et al. (1975) looked at the frequencies with which Black and White participants referred to various facial features when describing faces. Their results clearly showed a different pattern of features examined by Black participants and White participants. Face outline, eye size, eyebrows, chin, and ears were significantly more frequently used by Black participants as descriptors, while White participants used hair colour, hair texture, and eye colour more than other features and Black participants. Their data also shows that Black subjects used lips, mouth, and nose more frequently than White subjects especially when describing Black faces rather than White faces. Furthermore, Black participants use a more varied description, using more features than White participants. This may represent an important differentiation between the face recognition abilities of Black and White participants or physiognomic differences between Black and White faces.

Shepherd and Deregowski (1981) looked at how facial descriptions differ across Black and White faces using INDSCAL analyses. Internal facial features, with round fat faces, thick lips, and broad noses are related to Black physiognomy. This is contrasted with features associated with hair, which are related to White physiognomy. Hair features explain roughly 85% of the variance in the INDSCAL solution for White faces, but only 35% for Black faces. The internal lower facial features explain roughly 75% of the variance in the INDSCAL solution for Black faces, but the same features only explain 35% of the variance for White faces. Nevertheless, using a homogenous set of faces, both Black and White participants will use the same features to discriminate between them.

Several theories exist to explain the ORB. Levin (2000) suggests that own-race faces are processed in an individuating manner, whereas other-race faces are processed according to race. Sporer (2001) indicates an in-group/out-group model of face recognition, whereby people individuate faces of their in-group and not the out-group

faces. Valentine (1991) suggests that the dimensions of face-space are not diagnostic for other-race faces, and thus other-race faces are stored further from the centre of face-space in a tightly packed cluster. These theories are not necessarily mutually exclusive.

1.2.3. Adaptation

Adaptation is where the perceptual system is altered following constant stimulation of a particular characteristic (Blakemore, Nachmias, and Sutton, 1970). It is possible to adapt to specific spatial frequencies, causing these to become harder to discriminate following adaptation (e.g., Menees, 1998). In practice, this means the visual cortex can become adapted to, for example, high resolution images, whereby the after effect makes it easier to identify objects in low resolution images. An adaptation of three minutes will only affect a very narrow band of spatial frequencies (Webster & Mollon, 1999), such as causing visual fatigue for text on a computer screen (Lunn & Banks, 1986). Adaptation to combinations of basic stimuli can also occur; for example, adaptation to coloured sine-waves of high contrast and high spatial frequency, causes an after-image that is an 'opposite' blurred colour sine-wave (Webster, & Mollon, 1995).

An example of an adaptation experiment is that of Webster and Miyahara (1997). In a test of detection thresholds of spatial frequencies in natural images, these authors adapted participants to a particular spatial frequency for five minutes. Subsequently, identification of spatial frequencies was tested in a sequential test, with a further 6 seconds adaptation between each test stimulus. Their results were that the threshold for detecting similar spatial frequencies to the adaptor is much higher following adaptation.

Adaptation to spatial frequencies is dependent on contrast (Heinrich & Bach, 2002; Snowdon & Hammett, 1996). Georgeson and Harris (1984) have noted that

41

contrast threshold elevation functions are neither straight nor parallel at different spatial frequencies. In other words, the increase in contrast required to detect a sine-wave after adaptation is dependent on the spatial frequency of the adaptor. A high-contrast high-spatial-frequency adaptor causes detection thresholds to be raised more than an adaptor that is of the same spatial frequency but of low contrast. Adaptation can also cause an improvement in discrimination abilities under certain circumstances (Abbonizio, Langley, & Clifford, 2002; K. de Valois, 1977). Participants' discrimination of high contrast sine-wave gratings was tested at various levels of contrast and spatial frequency. Testing at higher contrasts (80%) discrimination improvement was found, but not at lower contrasts. Moreover, there is a temporal delay of 500 ms in spatial frequency adaptation (Wilson & Humanski, 1993).

Face distortion after effects (FDAEs) have been reported whereby a face distorted in one direction (e.g., compressed) will cause post-adaptation faces to appear distorted in the opposite direction (e.g., expanded; Rhodes, Jeffery, Watson, Clifford, and Nakayama, 2003; Rhodes and Jeffery, 2006). Webster and MacLin (1999) created a series of stimuli of faces that were distorted from the norm in a Gaussian fashion in vector format. The resulting set of faces was presented to participants using a nulling-match procedure, whereby participants had to adjust a distorted face to appear normal. After inspecting an adaptation face for five minutes and for eight seconds between the test images, participants had to adjust the distorted face such that it would appear normal. The adjustments the participants made were distorted in the opposite direction to the adaptation stimuli the participants had seen. The results were replicated in a normality rating procedure.

Beyond these adaptation effects, Webster and MacLin also noted that it was impossible to adapt to a normal (undistorted face). Moreover, after effects transferred

across faces and even to the perceivers' own faces. The after effects occurred for upright faces and for inverted faces, but only if the orientation of the adaptation face was matched with the orientation of the test faces. The perceived distortions (the after effects) disappeared quickly over time.

These adaptation effects are quite clear. Adaptation to a face distortion makes a test face appear distorted in the opposite direction. The fact that the distortion transfers across faces is interesting, but a comparison with non-face stimuli is required before drawing firm conclusions. The after effects may also simply be a distortion after effect and not special to faces.

The FDAE is partially size-tolerant since it transfers from an adaptor of one size to test stimuli of a different size (Zhao and Chubb, 2001). The assumption was that, since face-selective neurons in the superior temporal sulcus (Perrett, Rolls, & Caan, 1979) are size invariant (Grill-Spector, Kushnir, Edelman, Avidan, Itzchak, and Malach, 1999; Perrett, Rolls, & Caan, 1982; Rolls and Baylis, 1986), the adaptation effects should occur irrespective of the difference in size at adaptation and at test. In fact, Zhao and Chubb found significant adaptation effects even when the size difference between the adaptor and the test stimulus was altered by a factor of four, although the magnitude of the adaptation was significantly smaller when the test face and the adaptor did not match in size. Thus, part of the effect is size-invariant, but part is size-dependent.

FDAEs also transfer across parts of the retina (Hurlbert, 2001; Anderson and Wilson, 2005) and partially across viewpoints (Jeffery, Rhodes, and Bussey, 2006). In a review of how FDAEs transfer across changes in spatial frequency, contrast polarity, colour, and size, Yamashita, Hardy, K. De Valois, and Webster (2005), noted that the magnitude of adaptation was dependent of visual similarities between the adaptor and the test stimuli. Size and colour differences between the adaptor and the test stimuli reduced

the magnitude of adaptation significantly less than spatial frequency and contrast differences between the adaptor and test stimuli.

Not all kinds of distortions can cause the FDAE, however. Robbins, McKone, and Edwards (2007) demonstrated that adult participants' oddness ratings[2] were significantly affected when they were adapted to a "natural" facial configuration but were not affected by adaptation to "unnatural" facial configurations. In their study, participants were adapted to a stimulus with eyes shifted up but still aligned ("natural" facial configuration) or to a stimulus with one eye shifted up and the other eye shifted down ("unnatural" facial configuration). The effect of adaptation transferred from the adaptor identity to other faces, indicating a shift along the dimensions of face-space used to encode faces.

Similar to FDAEs are face identity after effects (FIAEs), whereby the perceived identity of a face is altered after adaptation to a particular identity. Leopold, O'Toole, Vetter, and Blanz (2001) conducted an elaborate study into FIAEs. In their study, a collection of 200 faces was morphed together to produce a prototype face. This was assumed to be the centre of the face-space (see Valentine, 1991). Some of the faces were selected to be used as targets. Due to the morphing process, each face identity could be measured in terms of Euclidean distances from the prototype face. Thus, a series of faces were created ranging from the prototype face to the face identity, each differing in identity 'strength'. Identification thresholds (the required identity strength to perceive the face identity) were taken before and after adaptation to an anti-face identity (opposite from the face-identity in terms of Euclidean geometry). Post adaptation to the anti-face, the identification threshold was lowered by 12.5% suggesting it was easier to perceive the identity following adaptation.

---

[2] Oddness ratings were of the form "How much does this face look like Bob?"

Leopold et al. (2001) instructed their participants to examine the adaptation face rather than fixate on it. Thus, FIAEs are translation-invariant and are not based on low-level after effects. The FIAE reported by Leopold et al. (2001) was maintained with delays up to 2400 ms, but died away quickly after the presentation of the test faces. Furthermore, the magnitudes of the adaptation effects were similar for upright and inverted faces, provided that the adaptation and test faces were in the same orientation.

The magnitude of FIAE is typically measured in terms of difference in identification thresholds pre- to post-adaptation and is dependent on the presentation duration of both the adaptor and the test stimuli (Leopold, Rhodes, Müller, and Jeffery, 2005). Stronger FIAEs were observed when the adaptation stimuli were presented for longer durations (16000 ms) than shorter durations (1000 ms). Moreover, FIAEs were significantly stronger when the test stimuli were presented for shorter durations (100 ms) than longer durations (1600 ms).

Another facet of the FIAE is that it transfers across viewpoints at least in some participants (Jiang, Blanz, and O'Toole, 2006). Their participants were adapted to a face image in one pose and tested on images in the same or different poses. Their results indicated that although significant adaptation occurred when the faces are in a different pose, the magnitude is significantly less than when the images are in the same pose. This study certainly indicates that this adaptation is not based solely on the visual similarity between adaptation and test.

The results of Jiang et al. (2006) are at least partially similar to those of Benton, Jannings, and Chatting (2006) regarding the FDAE. Benton, et al. found that the magnitude of FDAEs were greater for faces tested in the same pose as the adaptor face than for faces tested in a different pose to the adaptor face. These authors strongly suggest that the FDAE is viewpoint dependent. However, their conclusion is unwarranted

for two reasons. Firstly, the magnitude of the FDAE was significantly greater than zero for faces tested in a different pose to the adaptor face. This suggests that, at least in part, the FDAE is viewpoint independent. Secondly, there were substantial individual differences in their data, whereby one participant showed the same magnitude of the FDAE for faces tested in the same and different pose to the adaptor face, and one participant showed no adaptation to faces tested in a different pose to the adaptor face. These individual differences were not discussed by these authors.

According to Moradi, Koch, and Shimojo (2005), the FIAE requires conscious perception, since it is significantly reduced if the adapting face is invisible to participants. Moradi et al. tested the effect different types of suppression had on the magnitude of the FIAE. Though the FIAE transferred from the adapted retina to the unadapted retina, it disappears when the participants are attending to the eye that is not adapted. For example, one eye is being adapted to a face identity, while the other eye is presented with a pattern of moving random dots. Participants who attended to the moving pattern often ignored the face and failed to show the FIAE.

Moradi, et al. (2005) also tested whether imagination can cause the FIAE. Six participants were trained to associate names with the antifaces of Leopold, et al. (2001). This training lasted 300 trials. They were then asked to imagine one of the faces and were asked how clear their mental image was. Participants reported that their visualisation was vivid and yet demonstrated no FIAE even after prolonged visual imagery. This observation is apparently at odds with the fact that there is activation in the fusiform gyrus (the "face specific" area of the inferotemporal cortex, IT, Kanwisher, et al., 1997) during mental imagery of a face (Kreiman, Koch, and Fried, 2000; O'Craven and Kanwisher, 2000) or with the absence of awareness (Marois, Yi, and Chun, 2004; Moutoussis and Zeki, 2002). Since there are significant individual differences in the

ability to mentally visualise (e.g., Amedi, Malach, and Pascual-Leone, 2005; Bywaters, Andrade, and Turpin, 2004; Issac and Marks, 1994; McKelvie, 1994), it is possible that the small sample (6 participants) tested by Moradi et al. was affected by a single participant unable to accurately visualise faces. Training participants to associate a name with a two-dimensional digitised face that is only ever seen in one pose and in one picture is unlikely to lead to an accurate face representation (c.f., Burton, et al., 1999; Jiang, Blanz, and O'Toole, 2007). Moreover, the assessment of visualisation clarity was not based on previous mental imagery work (e.g., Marks, 1973).

The studies described thus far have tested FIAEs in unfamiliar faces. The representation of familiar and unfamiliar faces may well be different (Ryu and Chaudhuri, 2006) and involve different face recognition systems (Megreya and Burton, 2006). Indeed, the magnitude of the after-effects interacts with familiarity with the face (Fang, Blanz, & O'Toole, 2007). Fang et al. (2007) tested identity adaptation across viewpoints using familiar and unfamiliar faces. Adaptation effects were greater with familiar faces and would transfer across viewpoints more readily for familiar faces than unfamiliar faces. Ryu and Chaudhuri (2006) also report adaptation of unfamiliar faces that crosses viewpoints (see also, Pourtois, Schwartz, Seghier, Lazeyras, & Vuilleumier, 2005). In terms of the FDAE, Carbon and Leder (2005) have shown that the FDAE in a familiar face (Princess Dianna) does reliably transfer to novel viewpoints. This result is in contrast to the data reported by Benton, et al. (2006). In addition to using familiar rather than unfamiliar faces, Carbon and Leder (2005) used more participants than Benton et al. and all their participants were naïve to the purpose of the experiment. Carbon and Leder (2005) have also shown that the FIAE for familiar faces lasts a long time (they tested up to 80 minutes after adaptation), much longer than the 2400 ms suggested by Leopold, et

al. (2001) and may even result in permanent changes to the stored representation of the face in memory (see Carbon and Leder, 2006).

Jiang, et al. (2007) specifically tested the degree of familiarity that participants have with a face and the magnitude of the FIAE in within- and between-viewpoint adaptation. They trained 90 participants on a set of 16 faces to varying degrees of familiarity. In the low-familiarity condition, the participants only saw the face in a frontal view twice. In the high-familiarity conditions, participants saw the faces eight times each in frontal views only, rotated views only or half frontal and half rotated views. An extreme familiar condition was also included where participants saw the faces 16 times in frontal views and 16 times in rotated views. This highly controlled study demonstrated that, although the magnitude of adaptation was greater for within-viewpoint adaptation, there was still significant adaptation for between-viewpoint adaptation. Moreover, greater FIAEs were observed when the participant had higher familiarity with the face. Indeed, the difference in magnitude of adaptation between the same- and different-viewpoint conditions was smallest for the extremely familiar condition. This suggests that familiar and unfamiliar faces are not only represented differently, but also that the results of Benton et al. were misleading due to the use of only frontal view presentations and largely unfamiliar faces.

1.2.3.1. Explanations of Adaptation

Clifford, Wenderoth, and Spehar (2000) put forward a set of principals to account for the cortical coding of sensory information based upon lateral inhibitions and excitations between orientation-tuned cortical neurons. The axioms of this model are based upon the idea that the sensory system maps environmental attributes onto patterns of fixed neuronal responses. The mapping is dynamic in such a way that it optimises the fixed

neuronal responses, and is based on the statistical stationarity in the structure of the environment. The mapping is changed when the structure of the environment is altered.

Implicit within this model is that neuronal populations have tuning curves for particular stimuli characteristics. The width of tuning curves represents the population's response bandwidths, whereas the peak represents the preferred stimulus property. The perceived response is given by the weighted vector average of the units responding to the stimulus. Distribution-shift theory posits that the adaptor affects the way the stimulus is mapped onto the responses of the population of neurons (c.f., Mather, 1980), resulting in a population vector that no longer corresponds to, for example, the motion of the stimulus. Thus, at a single cell level, this may look like neuronal fatigue, but at a population of cells, this is adaptive gain control (Clifford and E. Weston 2005). Since responding to the constant presence of a stimulus is wasteful on resources, such a model allows for adaptive gain suppression as a functional method for reducing the need for constant firing.

In terms of face-space, one can consider what the neuronal representations of particular dimensions are. The neural analysis provided by Robbins, et al. (2007) is shown in Figure 1.3 and will be applied here. Traditionally, two-pool models have been considered when explaining apparently opponent-processing (e.g., Over, 1971; Regan and Hamstra, 1992; Rogers and Graham, 1985). These suggest that at one end of dimension there is a neural population for the extreme of particular feature and the other end of that same dimension, there is a neural population for the opposite extreme of the same feature. The amount that each pool is firing determines the size of the feature seen. This premise has received support from physiological data (e.g., Maddess, McCourt, Blakeslee, and Cunningham, 1988; Movshon and Lennie, 1979). For an alternative multi-channel model see Hubel and Wiesel (1974).
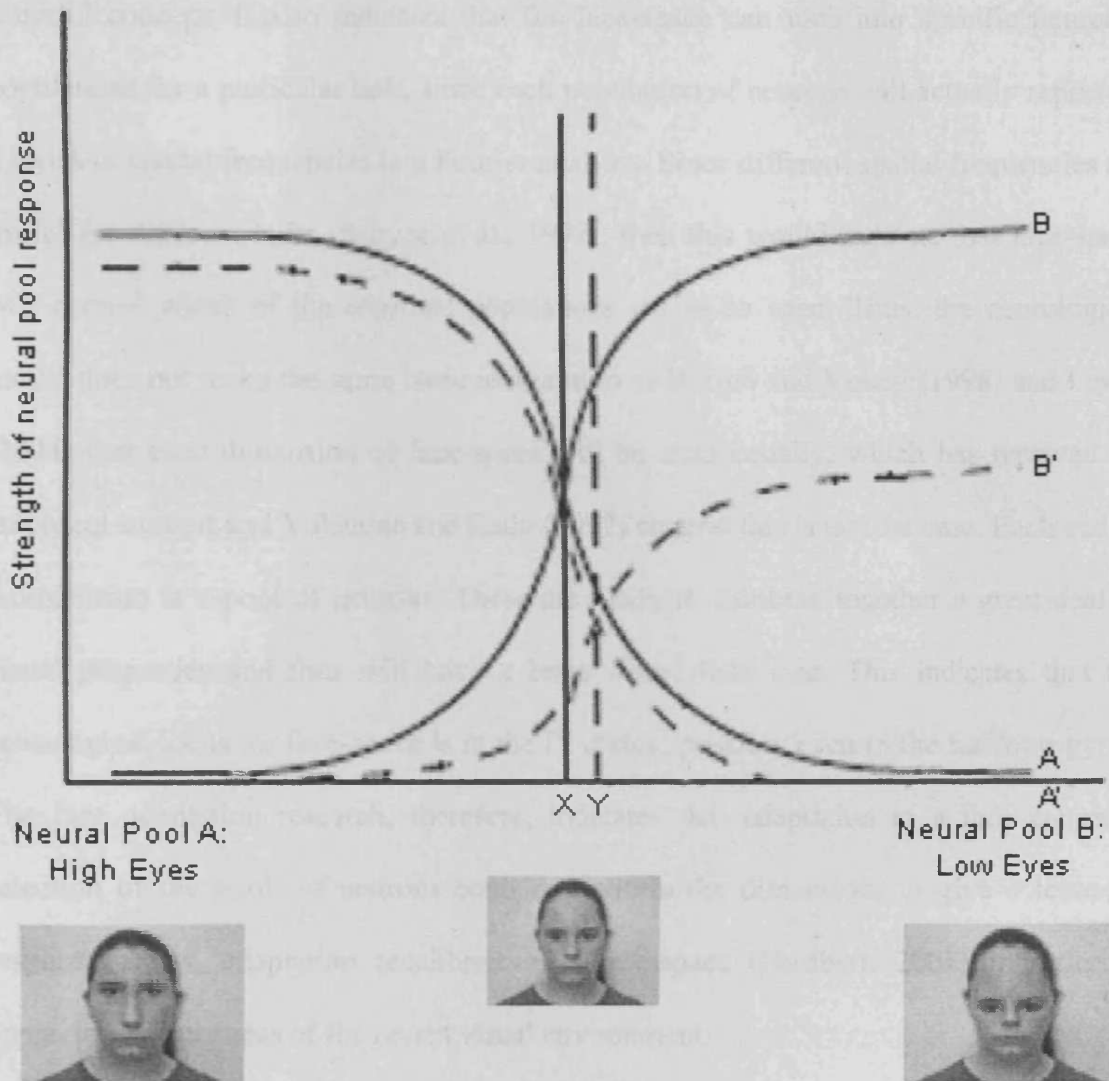
Figure 1.3. A model for the neural response to FDAEs. Strength of the neural response for neural pool A is represented by line A pre-adaptation and line A' post-adaptation to a face with low eyes. Strength of the neural response for neural pool B is represented by line B pre-adaptation and B' post-adaptation to a face with low eyes. The neural response in pool B is lower following adaptation to an image mainly represented by this pool, shifting the perceived midpoint from point X to point Y. Adapted from Robbins, et al. (2007).

This neurologically based model for how dimensions may express themselves may prove a useful concept. It also indicates that the face-space can tune into specific neuronal populations for a particular task, since each population of neurons will actually represent a series of spatial frequencies in a Fourier analysis. Since different spatial frequencies are useful for different tasks (Schyns et al., 1992), then this would indicate that face-space will control which of the neuronal populations are to be used. Thus, the neurological model does not make the same basic assumption as Burton and Vokey (1998) and Lewis (2004) that each dimension of face-space will be used equally, which has received no empirical support and Valentine and Endo (1992) suggest this is not the case. Each end of a dimension is a pool of neurons. These are likely to combine together a great deal of visual properties and thus will have a large visual field size. This indicates that the neurological locus for face-space is in the IT cortex, possibly even in the fusiform gyrus. The face adaptation research, therefore, indicates that adaptation to a face causes a selection of the pools of neurons contained within the dimensions to give a lessened response. Thus, adaptation recalibrates the face-space (Hurlbert, 2001) in order to appreciate the nuances of the recent visual environment.

1.3. Face recognition in children

Significant attempts have been made to understand how face recognition skills develop. These have led to varying and complex research findings that, at times, appear contradictory. Research indicates both quantitative and qualitative changes in face recognition abilities. In this section, research showing the developmental trend of face recognition shall be presented. The following subsections more closely examine the inversion effect, the ORB, and finally perceptual development in children.

Face recognition abilities improve through childhood (Adams-Price, 1992; Blaney and Winograd, 1978; Ellis and Flin, 1990; Goldstein and Chance, 1964; List, 1986; Searcy, Bartlett, and Memon, 2000; Searcy, Bartlett, Memon, and Swanson, 2001). Feinman and Entwisle (1976) demonstrated that this improvement in face recognition abilities follows a linear pattern reaching an asymptote at the age of 12. Ellis, Shepherd, and A. Bruce (1973), however, found that 12-year-old children performed worse on a face recognition task than 17-year-old children. These differences may be due to the stimuli set used by these authors. Feinman and Entwisle (1976) tested children's recognition of 40 children's faces, whereas Ellis et al. (1973) tested children's recognition of 60 adult faces. Since an own-age bias has been reported in face recognition, whereby participants are better at recognising faces of their own age than other ages (Anastasi and M. Rhodes, 2005, 2006; Anastasi, Hodges, and M. Rhodes, 2006; Chung, 1997; Wright and Stroud, 2002), it is probable that the Feinman and Entwisle (1976) stimuli benefited children whereas the Ellis et al. (1983) stimuli benefited adults. Thus, there is a problem in determining absolute levels of performance in face recognition tasks for children at different ages, especially since the magnitude of the own-age bias is different for adults and children (e.g., Anastasi and M. Rhodes, 2005; Fulton and Bartlett, 1991). Furthermore, many studies only test two or three age groups and may fail to fully explore the nature of the development of face recognition.

Carey (1978) reported a more extensive study of face recognition in 6-, 8-, 10-, 12-, 14-, and 16-year old children. She demonstrated that there is a gradual improvement in face recognition that is apparently linear up to the age of 10. Performance on her face recognition task is lower at ages 12 and 14, before returning to levels of performance obtained by 10-year-old children. Flin (1980, 1985) replicated these findings and extended them, measuring hits, false alarms, and d'. Flin's data are summarised in Figure

1.4. Flin's data are indicative of a dip in face recognition performance at age 11 or 12. This dip is mainly due to a change in false alarms, with a sudden decrease in false alarms at age 10.6 and increase at age 11.6. Indeed, if a linear function is fitted to Flin's data, the d' measure shows a blip at age 11.6 with a peak at age 10.6.
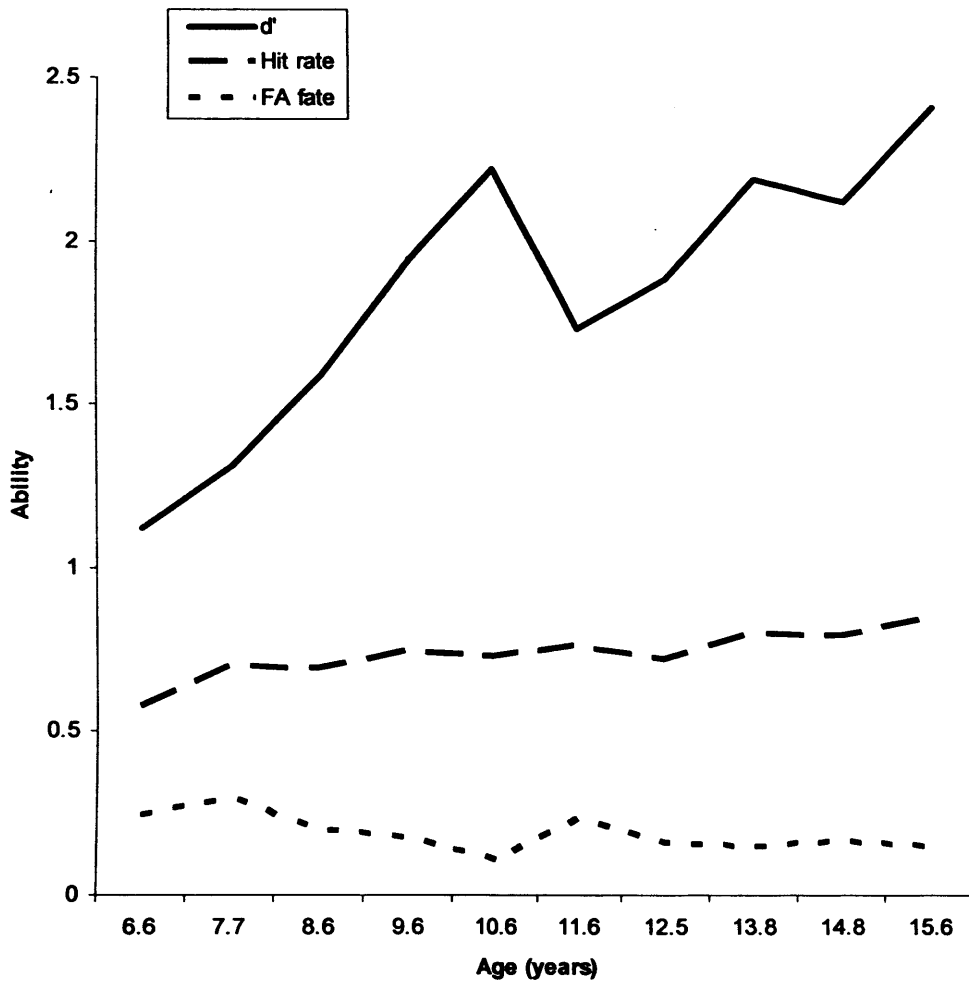
Figure 1.4. Mean ability score is split by age, with three measures of ability presented: hits, false alarms, and d', adaptated from Flin (1980).

This apparent blip in face recognition roughly at age 12 has not been fully explored and is difficult to explain. This blip in performance is not unique to faces (Flin, 1983), and thus may represent a general cognitive deficit at this age. Explanations for this may be

maturational or involve a change in cognitive style (e.g., Carey, 1981; Flin, 1985; Flin and Dziurawiec, 1989). Such changes may be qualitative. A demonstration of qualitative changes has been shown in the matching tasks.

Matching tasks have also been used to examine children's face recognition abilities. Saltz and Sigel (1967) compared adults' and six-year-old children's abilities to match two photographs that were simultaneously presented. Children were limited in their ability to match faces if they differed in terms of pose, expression, or paraphernalia (see also, Ellis, 1992a, b; Freire and Lee, 2001). Diamond and Carey (1977) provided some of the first data for research on the errors made by children when recognising faces. Using a matching paradigm, their participants were instructed to identify which two faces were the same as each other. The crucial manipulation was whether the target and distracter faces had the same expression and/or paraphernalia. Recognition accuracy improved between the ages of 6 to 10, with little difference in performance for older age participants. The younger children were susceptible to making errors involving paraphernalia. This was the case irrespective of whether the matching task was one based on simultaneous presentation or short-term memory. This was not the case for familiar faces, which were not associated with any paraphernalia based errors. These studies indicate that children are not adequately encoding faces, since face matching tasks involving simultaneous presentation do not place demands on memory, just on perception (Carey and Diamond, 1994), in addition to having poorer memory capacity (Carey, Diamond, and Woods, 1980).

1.3.1. The Inversion Effect as Shift in Processing Style

Six- and eight-year-old children do not show the inversion effect when tested in a matching paradigm (Carey and Diamond, 1977; Joseph, Gathers, Liu, Corbly, Whitaker, and Bhatt, 2006) or a recognition paradigm (Goldstein, 1975). Ten-year-old children do show the adult inversion effect, however. Indeed, whereas in adults, there is greater activation in the left superior occipital lobe when processing upright faces than inverted faces, in children under 8-years of age this advantage is not observed (Joseph, et al., 2006). Schwarzer (2000) demonstrated, using a categorisation task, that seven-year-old children processed face stimuli and non-face stimuli in a similar analytical manner. Moreover, they did not demonstrate a significant inversion effect, whereas adults showed a significant inversion effect, and some 10-year-old children demonstrated an inversion effect. However, these studies, and others that fail to show an inversion effect in children (e.g., Hay and Cox, 2000), are primarily matching or classification tasks. In recognition tests, and when floor and ceiling effects are controlled for, the inversion effect is apparent by the age of three (Carey, 1981), five (Fagan, 1979; Flin, 1983), or seven (Young and Bion, 1981, 1982). At eight-years-old, children show the adult level of the inversion effect and their improvement in face recognition skills is continuing in parallel (Itier and Taylor, 2004). The results regarding the inversion effect in children are thus quite mixed. Nevertheless, an age by orientation interaction is often found in recognition tests, indicating that the *magnitude* of the inversion effect is dependent on age, where it increases with age (Carey, 1981; Carey and Diamond, 1994; Flin, 1983; Goldstein and Chance, 1964).

Due to the differences in the magnitude of the inversion effect in children from that in adults, some authors have suggested that children rely more on featural rather than configural coding (e.g., Hay and Cox, 2000). However, using the parts and wholes test (Tanaka and Farah, 1993) in which participants are required to recall individual features

that were either presented in isolation or as part of a whole face, Pellicano and Rhodes (2003) found evidence that children as young as four-years old use configural processing. These authors found that children as young as four show the adult pattern of results, whereby they can recall features learnt as part of a whole face better than those presented in isolation.

Carey and Diamond (1994) investigated the presence of the composite effect in children. The composite effect is where the top half of one face is combined with the bottom half of another face, creating a new face. The two halves are more difficult to identify upright than they are inverted (Young, et al., 1987). This effect is shown in children aged 6 and 10 with the same magnitude as adults (Carey and Diamond, 1994). The only age effect observed in this study was that 6-year-old children were generally slower than the older children and adults in identifying both the upright and inverted composites.

Thus, the evidence is mixed as to whether younger children show an inversion effect or not. Furthermore, if the inversion effect is based on a disruption of configural information in the same way as the parts and wholes task and the composite face task, then it is even more unclear why these effects give different results in children. It is likely that children do show an inversion effect, but its magnitude is less than that of adults.

1.3.2. The Own-Group Biases

Chance, Turner, and Goldstein (1982) tested children from 6-years old to adulthood in a recognition test using own-race and other-race faces. Their results are equivocal: the youngest children were able to recognise both own- and other-race faces equally well, whereas children above the age of 13 were better at recognising own-race faces.

Moreover, until the age of 13 years, boys and girls performed equally well at this recognition task. Following this age, however, girls began to outperform boys.

The results of Chance, et al. (1982) have not always been replicated. Feinman and Entwisle (1976) demonstrated a significant ORB children age 6 years and over. The key differences between these two experiments is that the Chance et al. study only tested White children on White and East Asian faces, whereas Feinman and Entwisle tested White and African-American children on White and African faces. Moreover, Chance, et al. examined the effect using the highly sensitive d' measure of accuracy, whereas Feinman and Entwisle only measured hit rates as their dependent measure.

In a particularly poor attempt to disentangle these results, Pezdek, Blandon-Gitlin, and Moore (2003) tested the ORB in 6- and 9-year old children and adults using own-race and other-race participants and faces. Using the non-parametric measure[3] A_g to test, these researchers discovered the ORB was present at all ages that they tested. Whereas there was a general improvement in face recognition abilities with age, there was no race of face by race of participant interaction. However, their results are contaminated by two problems. The first is the near chance levels of other-race recognition for the youngest age group. This is indicative of a floor effect. Moreover, the magnitude of the ORB in African-American participants is far greater than that of the White participants. This finding is highly unusual and the opposite is usually found, where African-American participants show a much smaller ORB (e.g., Anthony, et al., 1992; Feinman and Entwisle, 1976; Malpass and Kravitz, 1969; see also Valentine and Endo, 1992, testing White and East Asian participants and faces).

Another commonly observed effect in adults is the distinctiveness effect. This effect is where distinctive faces are easier to learn, recognise, and make familiarity

---

[3] The logic for the use of a non-parametric statistical analysis was not offered by these authors, when the convention is to use the parametric d'.

decisions about (see e.g., Ellis, Shepherd, Gibling, and Shepherd, 1988; Valentine and Bruce, 1986). Ellis (1992a, b) has demonstrated that in a recognition paradigm, 6-year old children failed to show this effect. Whereas adults can discount distinctive distractors more easily than typical distractors in a face recognition test, 13-year old children are unable to discount distinctive distractors (Barlett, Hurry, and Thorley, 1994), possibly explaining the observed developmental "blip" in performance.

Johnston and Ellis (1995) used a standard old/new recognition paradigm to show that there is an effect of age on response latency (younger children are slower in responding than older children) and on d' accuracy, testing 5-, 7-, 9-, 11-, 13-, and 20-year old participants. A significant age by distinctiveness interaction was also reported, whereby older participants recognised distinctive faces more accurately and quickly than typical faces. Younger participants did not show this effect. Adult performance levels were observed for 7-year-old children for accuracy and 9-year-old children for response latency.

Johnston and Ellis (1995) also tested the same age children using a face classification task. Adults are significantly better at classifying a typical face as a face than a distinctive face. Younger children showed the standard distinctiveness disadvantage in face classification for response time and accuracy, though younger children were less accurate and slower than adults for both distinctive and typical faces. The improvement was gradual and consistent over all of the ages tested. Unusually, for accuracy there was not a distinctiveness advantage for 9- and 11-year old children. No explanation was offered for this.

Like the inversion effect, the data on the ORB in children is mixed. Some studies report an ORB in children, whereas others do not. Related distinctiveness effects are also not observed in children, suggesting differences in facial memory. Before embarking on a

theoretical analysis of these findings, perceptual development research will be summarised.

1.3.3. Perceptual and Attentional Development

Developmental psychologists have explored cognitive and perceptual development in great detail. It is worthwhile exploring how the development of face recognition fits in with such developmental trends. By far the most influential model of child development is that of Piaget (e.g., 1952). This model divides development into discrete stages. The age at which Carey (1978) and Flin (1980) found a blip in overall face recognition performance is roughly the transition from concrete operational to formal operational thought in Piagetian terms. Concrete operational thought is characterised by a need to physically see an object, whereas formal operational thought can be more abstract. Whereas concrete operational thought involves simple flexible schemas, formal operational thought involves the use of many more complex schemas (c.f., K. Nelson, 1981; 1996). Schema theory has been used to explain differences between adults' and children's face recognition abilities, whereby the face schema develops over time (Goldstein and Chance, 1980) and is used by adults and not children.

Perception researchers have indicated that there is a developmental progression in the use of perceptual schemata or scripts. This development is evident in a change in visual search and attention. Younger children are less systematic (Vurpillot, 1968) and take longer to decide where to move their eyes than adults (Cohen, 1981). Attentional mechanisms and, in particular, attentional shifting abilities develop slowly (e.g., Akhtar and Enns, 1989; Brodeur and Enns, 1997; Enns and Brodeur, 1989) not reaching adult levels until after 11 years of age (Pearson and Lane, 1991). Since shifting attention is

more difficult for children than adults, it is likely that a perceptual script has developed to process the most frequently encountered faces. Indeed, eye-tracking data indicates that the attention paid to faces follows a set script, starting at the eyes and hair line for adult White participants (Althoff and N. Cohen, 1999; Cook, 1978; Stacey, Walker, and Underwood, 2005). This eye movement script may well relate to the development of focusing attention on the more diagnostic visual properties of a face, since eye movements are affected by task demands, meaning, context, and expectations (Antes and Penland, 1981; Findlay, 1981; Stark and S. Ellis, 1981).

Another attentional processing related to the neurological changes that occur throughout development is that of perceptual narrowing. This is the process whereby the perceptual system develops in such a way that it "learns" to filter out unwanted stimuli (Macchi-Cassia, Kuefner, Westerlund, and Nelson, 2006; Pascalis, de Haan, and Nelson, 2001). For example, infants can discriminate native and non-native speech sounds whereas children and adults cannot (Nelson, 2001) and infants can discriminate between pairs of monkey faces, whereas children and adults cannot (Pascalis, et al., 2001). Neurologically, this perceptual narrowing phenomenon is related to neuronal pruning (Johnson and Vecera, 1996), whereby the number of axons in the visual cortex decreases with age. Perceptual narrowing, chunking (Miller, 1956), and schemas reduce attentional load. This reduction is because many unwanted and non-diagnostic stimulus features can be ignored. Without this narrowing, all stimulus features will be processed and this may overload the attentional resources.

Carey and Diamond (1994) offer an explanation for the time course for the maturation of face recognition. They noted that the inversion effect is observed for perceptual experts of other categories of objects made up of features that bind together to form a meaningful whole. Dog experts who receive 10 years of training in judging dogs

for shows demonstrate an inversion effect for dogs' faces of a similar magnitude to that of faces (Diamond and Carey, 1986). As such, Carey and Diamond posit that perceptual expertise takes 10 years to develop. However, perceptual organisation can occur within a matter of hours following congenital blindness (Maurer, T. Lewis, Brent, and A. Levin, 1999), suggesting such prolonged experience may not be required for the observed face recognition effects to occur.

The developmental research described here has been largely unsystematic and atheoretical. In many cases, the stimuli used to test children's abilities are unsuitable, leading to floor or ceiling effects. This will hide particular important effects. Moreover, the face recognition research has largely not been tied to general cognitive and perceptual development. Using the research conducted on face recognition and keeping in mind developmental research, models of the immature face-space can be devised.

1.4. Developmental Models of Face-Space

Very few attempts have been made to understand how face-space may exist in very young infants (but see Johnston and Ellis, 1995). Face-space in adulthood is characterised by many dimensions (between 10 and 200, see section 1.1.3.) and more faces stored around the centre of the space (Valentine, 1991, but see Burton and Vokey, 1998). The distribution is uneven, with a denser region around the centre, and sparse distribution at the outer regions of the space. This offers several alternatives for how the immature face-space may differ from the adult face-space. Three suggestions shall be described in detail here and the predictions they make regarding the development of face recognition. The first model is static version, whereby the dimensions of the immature face-space are the

same as the adult face-space. The second model is the expanding face-space, whereby the number of dimensions in the immature face-space is fewer than that in the adult face-space. The third model is a novel model and is the shrinking face-space, whereby the number of dimensions in the immature face-space is greater than in the adult face-space.

1.4.1. Constant-face-space

Johnston and Ellis (1995) suggested that the immature face-space may be based upon the same dimensions as the adult face-space. They termed this the uniform model of the immature face-space. To begin with, the face-space is empty. As faces are encountered, so the face-space begins to fill up. Experience of faces dictates the distribution of the faces stored in face-space.

This model of face recognition predicts that the dimensions of face-space are predetermined (either genetically or through some other means, e.g., CONSPEC, see Morton & Johnson, 1991) and are thus suitable for recognition of the most frequently encountered faces. Alternatively, the effects observed in adulthood (the ORB and distinctiveness effects) are caused only by the changes in the distribution of the stored faces. Due to this, in young children, the distance between each stored exemplar is roughly equivalent (c.f., Johnston and Ellis, 1995). Thus, young children will not show an ORB or distinctiveness effects. These effects develop with experience leading to a more crowded face-space around the centre of the space.

1.4.2. Expanding-Face-Space

Johnston and Ellis (1995) suggest that their evidence on the lack of distinctiveness effects in children's face recognition is indicative of their face-space having fewer and "less appropriate" dimensions than adults (pp. 463). Both Lewis (2004) and Valentine (1991) indicated that new dimensions may be added to face-space during development especially when highly typical faces are encountered. Consider the situation early on in face-space: Two faces are encountered and will appear very similar. If they cannot be automatically discriminated then the face recognition system must learn to distinguish between them. Perceptual learning (e.g., McLaren, 1997; Mundy, Honey, and Dwyer, 2007) indicates that when two stimuli are presented, participants will attend to the differences between the stimuli and inhibit the similarities. In a classic phrasing, two stimuli are presented, AX and BX, where X is the common elements and the unique elements are A and B (c.f., Hall, 1980, 1991; Honey, 1990). Presenting AX and BX together, participants will attend to A and B and inhibit X as X is not useful in distinguishing between the two features. In terms of face recognition, X may be first-order relational information (two eyes above a nose above a mouth). This information is not beneficial for discriminating between two faces. Expanding-face-space indicates that when AX and BX are presented, the dimension that distinguishes between A and B will be added to the face-space unless the face can be discriminated using existing dimensions. This process will continue until most faces are recognisable.

This basic explanation ignores a nuance within perceptual learning: That is the fact that perceptual learning is more effective if two stimuli are similar to each than if the two stimuli to be discriminated are very different (e.g., Mundy, Dwyer, and Honey, 2006). Due to this, dimensions will only be added when two faces that are similar are

encountered provided that the existing dimensions cannot be used to distinguish between the faces. Since the majority of faces encountered are typical – and, by definition average – dimensions will only be added to distinguish between typical faces. Thus, experience of faces dictates what dimensions will be added to face-space.

There are problems with this model, however. Johnston and Ellis (1995) suggest that the child's face-space contains less appropriate dimensions than that of the adult. If this is the case, and only new dimensions are added, the less appropriate dimensions still remain. To have non-optimal dimensions in face processing is unusual considering the expert nature of the human face recognition system.

Schema theory can also be used to describe this situation. The schema used by children to recognise faces may be a more vague and imprecise. Thus, children's face recognition will be poorer than that of adults. With development and experience, the schema becomes more appropriate and valid for recognising the most frequently encountered faces.

This model, therefore, predicts that children will not show standard distinctive effects in face recognition that adults do, since the face-space does not contain the dimensions to inform what is a distinctive and what is a typical face. However, expanding-face-space does predict that children will show an ORB since the dimensions in face-space will be added to help differentiate the most frequently encountered faces. Other-race faces are usually not the most frequently encountered and thus the face-space will not have the dimensions to encode other-race faces appropriately.

Younger infants are better able to discriminate between pairs of monkey faces than older infants are (Lewkowicz and Ghazanfar, 2006; Pascalis, et al., 2002; Scott, Shannon, and Nelson, 2005, 2006). Moreover, a new born will be able to discriminate between phonemes of virtually every language, but by 9-months, the infants can only discriminate native speech sounds (Cheour, Ceponiene, et al., 1998; Kuhl, Williams, Lacerda, Stevens, et al., 1992). These effects are known as perceptual narrowing (Nelson, 2001), whereby experience narrows the ranges of stimuli that will elicit a response (stimulus discrimination in animal learning, Pavlov, 1927; Sheffield, 1965). During an experimental session, participants generally improve on many perceptual tasks. Such practice effects are due to participants orientating their perceptual attention to the most relevant characteristics (Saarinen & Levi, 1995). This is the crux of perceptual narrowing.

A related effect of expertise in memory comes from research into chess (e.g., DeGroot, 1965; 1966). When planning a move in chess, expert players consider fewer moves than novice players (DeGroot and Gobet, 1996; Gobet and Simon, 1996a, b). The information processing of expert chess players is characterised by laziness (Charness, 1981; Saariluoma, 1994; Wagner and Scurrah, 1971) and efficiency (Simon and Barenfeld, 1969; Simon and Gilmartin, 1973). Expert chess players can remember the positions of pieces if they occur in a position likely during a real game at a far greater accuracy level than that of novice chess players. However, novice chess players can remember the positions of chess pieces better than experts if the pieces are in random positions not likely during a game of chess (Chase and Simon, 1973a, b). Thus, memory is affected by expertise such that the more frequently encountered stimuli are better

remembered. Thus, there is some perceptual or memory narrowing associated with expertise.

It is possible that face-space develops in a similar manner. Experience narrows the number of dimensions used to recognise faces. As this develops, the perceiver learns to focus on particular aspects of faces. Thus, the dimensions become more appropriate with development. This inherently makes face-space a more flexible system, whereby dimensions can be added, discarded, changed, or ignored. Since there is an idea that nothing is permanently lost from long-term memory without brain damage (c.f., Loftus and Loftus, 1980; Penfield, 1969; but see Eysenck, 1993), then it is quite likely that the dimensions of face-space that experience has taught are less relevant are not discarded, but left unused or changed. In this way, it may be possible to reactivate those dimensions that have not been used for a while. Using an analogy of the dimensions of face-space as types of volcano, there are: active dimensions that are used to process all faces currently; dormant dimensions that have once been used to process faces but are no longer used; and extinct dimensions that have been used but are no longer useful and thus are discarded or changed into an active one[4].

Consistent with the idea that with development, the face-space becomes less flexible is the idea of the face schema. Goldstein and Chance (1976) have indicated that adults use a schema to compare other faces to. This does not exist in children, and thus the immature face-space may well compare to more features than would be present in the adult face schema.

Shrinking-face-space makes several novel predictions about face recognition effects that are likely to be observed in children. Primarily, it suggests that children will be able to recognise and process faces that adults will not be able to process due to the

---

[4] The term "shrinking" is slightly misleading given that dimensions can be added. However, overall, the number of dimensions in adult face-space is smaller than in the child's face-space.

fact the immature face-space contains may more dimensions than the adult face-space. Thus, children are likely to show a much reduced inversion effect and ORB. Furthermore, children will be more able to learn "unnatural" facial configurations. One reason why children have a poorer general face recognition ability is due to attentional demands being taken up by the face-space trying to encode along too many dimensions.

Conclusions

Face-space as a metaphor for how faces are stored in memory is a highly successful model of face recognition (Valentine et al., 1995) in that it explains many effects of face recognition succinctly and is backed up by much evidence. How it accounts for the process of recognising a face has been presented here, followed by how it may account for several key effects in face recognition. However, the adult model of face-space is unable to account for children's face recognition abilities. For this, one must turn to one of the three models of immature face-space presented here: constant-face-space, expanding-face-space; shrinking-face-space.

Four experimental chapters will be presented that aimed to test these three models and explore the nature of face-space in adults further, based upon these three models. Chapter 2 presents four experiments that tested the face recognition abilities of children and adults using the inversion effect, the own-group biases, and adaptation. These experiments tested whether children were relatively better able to process less frequently encountered faces than adults, possibly indicating an immature face-space containing more dimensions than adults.

Eight experiments in Chapter 3 explored what happens to dimensions in adults that are no longer used by default. Regions of face-space were primed to allow for better

encoding. Perceptual learning tasks were employed to train participants to use different dimensions. Cueing paradigms were also employed to guide participants' attention to particular dimensions. These experiments reveal the flexibility of the adult face-space in terms of how dimensions are used.

In Chapter 4, adaptation procedures were employed to recalibrate dimensions of face-space. Recent experience was used to adjust the perception of novel faces. This experience was not only restricted to visual, but auditory, semantic, and imagined experience was also tested. Furthermore, the time-scale of adaptation was assessed.

In the final experimental chapter, Chapter 5, low-level spatial frequency aspects of the dimensions of face-space were explored. The effect Navon stimuli have on the recognition of faces was tested in line-up tasks and recognition paradigms. Furthermore, different types of Navon stimuli were used to affect the perceptual experience of faces. Finally, different types of bandpass filtered faces were used to assess the effect Navon processing had on face processing.

These experimental chapters tested between the three developmental models and explored what happens to dimensions of face-space in adulthood. The flexibility of the adult face-space was assessed using a variety of methods in terms of the use of dimensions and the distribution of faces stored along each dimension. The resulting findings were incorporated into a developmental model of shrinking-face-space and the adult flexible-face-space.

Chapter 2 – The Dimensionality of Face-Space in Children

Three models of the immature face-space were presented in Chapter 1 based upon the literature exploring the types of face processing, the own-group biases, and attentional development. Four experiments in Chapter 2 tested between these models. Experiment 1 tested the inversion effect using a recognition paradigm in children from the age of 5 years to adults. Experiment 2 tested children from the age of 5 years to adulthood in the recognition of own- and other- race and age faces. Experiment 3 tested the magnitude of adaptation to "natural" and "unnatural" facial configurations in children and in adults. Experiment 4 explored the ability of children and adults to match "natural" and "unnatural" facial configurations.

2.1. Experiment 1 - The Inversion Effect Developmentally

Virtually all faces share the same first-order relational information, such that two eyes are level and side-by-side and above the nose. For face identification, ignoring this first-order relational information is an economical process. Classifying faces as faces is likely, however, to depend on first-order relational information. Expert processing is more economical, involving more chunking (e.g., Newell and Simon, 1987) and schemas (e.g., Goldstein and Chance, 1976). Expert face processors may be able to take this into consideration and ignore first-order relational information in face recognition tasks. Inexpert face processing, however, may take into account this first-order relational information in recognition tests as well as classification tasks. Experiment 1 tested this

using the inversion effect in children. Inverted faces violate this first-order relational information, and are disproportionately harder to recognise than inverted objects (e.g., Yin, 1969). This effect occurs reliably in adults (see Section 1.2.1.), but has not been tested on children in a valid test throughout development (see Section 1.4.1.).

Experiment 1 aimed to test the inversion effect in children (from 5-years-old) compared to adults, in a standard recognition paradigm. Three possible developmental tracks are possible: The magnitude of the inversion of the inversion effect may increase with age as a product of experience; The magnitude of the inversion effect may be constant throughout development if the effect is not based upon expertise; Finally, developmental maturation in Piagetian terms would suggest that the inversion effect appears at a particular critical age (e.g., the change from concrete to abstract reasoning, Piaget, 1956). These hypotheses are termed the increasing inversion effect, the constant inversion effect, and the "all or none". It must be observed that with increasing age, face recognition abilities will improve. This will cause the absolute decrement due to inversion to appear to be increasing in the constant inversion hypothesis. As such, the constant inversion effect hypothesis predicts a relatively constant inversion effect rather than an absolute constant inversion effect. Floor and ceiling effects were avoided by choosing a number of stimuli that was manageable for the youngest children tested and will be sufficient to cause errors in the adult participants (c.f., Ellis, 1992a).

2.1.1. Method

2.1.1.1. Participants

Participants were: 13 five-year-old (six male); 19 six-year-old (nine male); 18 seven-year-old (eight male); eight eight-year-old (four male); 15 nine-year-old (nine male); 12

10-year-old (five male); 42 11-year-old (23 male); 45 12-year-old (26 male); 50 13-year-old (23 male); 22 14-year-old (10 male); and 10 15-year-old (four male) children recruited from a sample who returned consent forms to their school[5] and 22 adults (six male; mode age = 18) who volunteered for this study. These were pooled into larger groups (aged 5- and 6-years, 7- to 9-years, 10- to 12-years, 13- to 15-years, and adult). Uneven group size was a product of the recruitment process. All participants were ethnically White and self-reported that they had normal or corrected vision. All of the children were considered typically developing. The adults were recruited from the participant panel at Cardiff University and participated in this Experiment for money.

2.1.1.2. Materials

Thirty-two faces randomly selected from the Minear and Park (2004) face database were used in this experiment. Two images of each face were used, one presented during the learning phase and one presented during the test phase to ensure face recognition rather than picture recognition. The faces were of males and females aged between 20 and 22, all with similar hairstyles, positioned in a frontal view in neutral or happy expressions (this was randomised). All extraneous paraphernalia (e.g., glasses, beards, or ear-rings) were masked. All the faces had been pretested in other research to be in the middle of attractiveness and distinctiveness scales. The faces were presented 100 mm by 110 mm dimensions in 72 dpi resolution and these were counterbalanced, such that each image appeared the same number of times in each phase of the experiment and in both orientations. These were presented using Superlab Pro 2™ Research Software using a Toshiba Tecra M4™ Tablet PC.

---

[5] I would like to thank the schools and teachers for their invaluable support in conducting this research.

2.1.1.3. Design

A cross-sectional design was employed, in which age of participants (5 levels) was one independent variable (IV) and a second IV was orientation of the face (upright or inverted). Participants saw both upright and inverted faces. This created a 2 by 5 mixed design. The faces were counterbalanced such that each face was a target as often as it was a distracter and upright the same number of times as inverted. Each participant saw the same number of upright and inverted faces. Faces were presented in a random order. Recognition accuracy was measured using the SDT measure, d'. This combines hits and false alarms and is independent of response bias (Swets, 1966; but see Balakrishnan, 1998). This measure typically ranges from 0 to 4, and 0 is recognition at chance levels.

2.1.1.4. Procedure

A standard old/new recognition paradigm was employed. Participants were tested individually. Participants responded verbally and the experimenter entered the responses into a standard computer keyboard. Participants sat 50 cm from the computer screen. The Experimenter was blind to the contents of the screen, since it was turned away from him. Thus, the Experimenter could not influence the participants' responses. The set-up is shown in Figure 2.1.

The subsequent experimental procedure involved three consecutive phases: learning, distraction, and test. The learning phase involved showing the participants half of the set of faces (N = 16). Participants were instructed to rate each face for how attractive they thought the face was using a 1-9 Likert-type scale, where 1 was ugly and 9 was beautiful (c.f., Light, Hollander, and Kayra-Stuart, 1981). If a participant did not understand the scale, it was explained to them using alternative synonyms. The

presentation of each face was response terminated. Half of these faces were presented upright and half presented inverted.
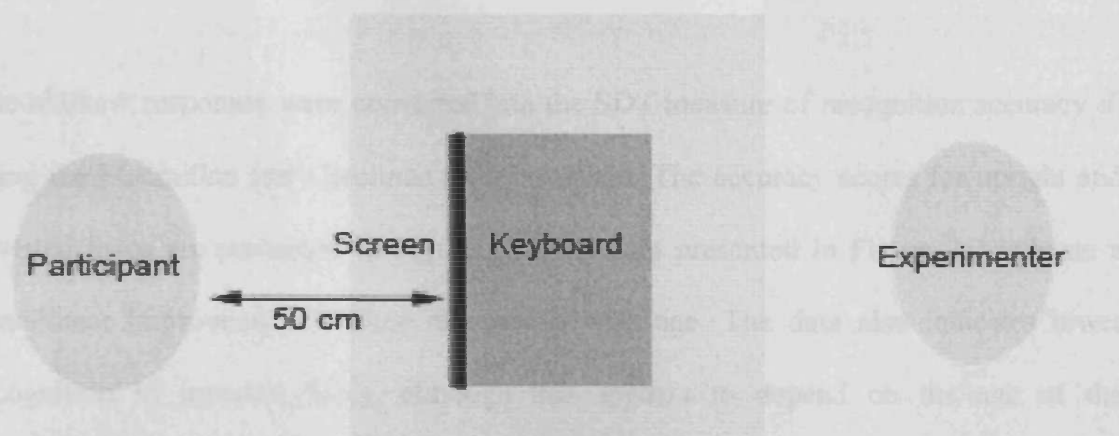
Figure 2.1. The set-up of experiments 1 to 4. The participants faced the computer screen and gave their responses verbally. The experimenter controlled the keypad, but could not see the screen.

Immediately after this presentation, participants were given some control questions. These were: "What is your first name?" "What is your surname?" "What is your gender?" "How old are you?" "What school do you go to?" "What school year group are you in?" "Where were you born?". If the participant did not understand the question it was explained using simpler synonyms, or the Experimenter was able to complete the information from school records. These questions took no longer than 60 seconds to administer.

Following this, the participants were given the test phase. In this, the participants saw all 16 target faces and 16 distractor faces sequentially and had to make an old/new recognition judgement to each face. Each presentation was response terminated. Half the faces were upright and half the faces were inverted. Orientation of the target faces was matched at learning and at test. The participants were told to be as quick and as accurate as possible. Once this phase was completed, the participants were thanked and debriefed.

2.1.2. Results

The old/new responses were converted into the SDT measure of recognition accuracy d' using the Macmillan and Creelman (2005) method. The accuracy scores for upright and inverted faces are presented in Figure 2.2. The data presented in Figure 2.2 indicate a near linear improvement in face recognition with age. The data also indicates lower recognition in inverted faces, although this appears to depend on the age of the participant.
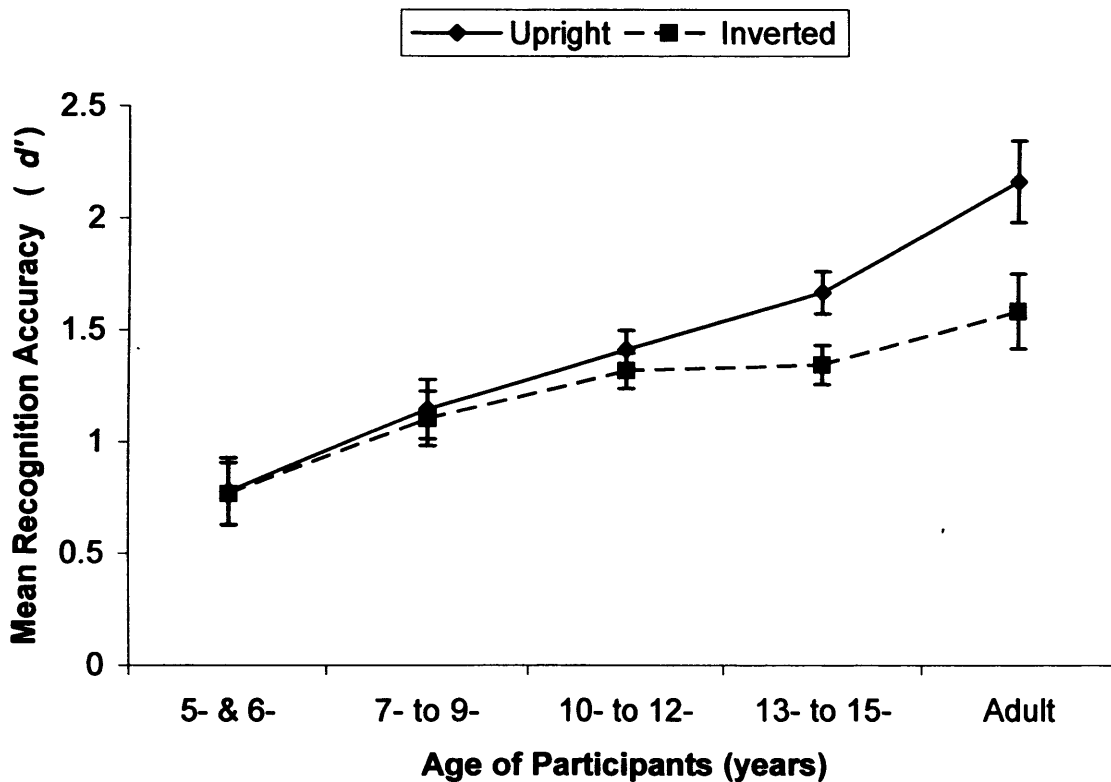
Figure 2.2. Mean recognition accuracy (d') for upright and inverted faces, split by age. Higher numbers indicate better accuracy, zero is chance. Error bars represent standard error.

A 2 by 5 mixed factorial ANOVA was conducted on the data summarised in Figure 2.2. with the factors: inversion (upright or inverted) and age of participants (5 to 15 and adult). This analysis revealed a significant effect of age, $F(4, 171) = 10.756$, MSE = $0.966, p < .05$, whereby older participants had significantly greater accuracy than younger participants. Tukey HSD *post hoc* tests were used here and for all subsequent experiments unless otherwise stated. These revealed that adult participants had greater recognition accuracy than 5- & 6-year-old children (mean difference = $1.101, p < .05$), 7- to 9-year-old children (mean difference = $0.748, p < .05$), 10- to 12-year-old children (mean difference = $0.508, p < .05$), and 13- to 15-year-old children (mean difference = $0.367, p < .05$). Moreover, 5- & 6-year-old children had a poorer accuracy than: 7- to 9-year-old children (mean difference = $0.353, p < .05$), 10- to 12-year-old children (mean difference = $0.592, p < .05$), and 13- to 15-year-old children (mean difference = $0.734, p < .05$). In addition, 7- to 9-year-old children had a poorer recognition accuracy score than 13- to 15-year-old children (mean difference = $0.381, p < .05$). No other pairwise comparisons were significant.

There was also a main effect of orientation of the test faces, $F(1, 271) = 11.810$, MSE = $0.376, p < .05$. This revealed itself through greater accuracy for upright faces than inverted faces (mean difference = $.209$). Crucially, the interaction between age-of-participant and orientation-of-face was also significant, $F(4, 271) = 2.539$, MSE = 0.376, $p < .05$. This interaction was explored using simple effects. This revealed that the inversion effect was significant for adults (mean difference = $0.579, p < .05$) and for 13- to 15-year-old children (mean difference = $0.323, p < .05$). No other simple effects were significant.

2.1.3. Discussion

The present experiment found a significant interaction between age of participants and orientation of faces. This was revealed such that children under the age of 13 did not show a significant inversion effect. At first glance, the magnitude of the inversion effect does appear to be increasing gradually. This suggests that there is an increasing inversion effect with development rather than a constant inversion. However, the present data do not permit a rejection of the "all or none" hypothesis.

These results are not consistent with the majority of studies that did show an inversion effect in children (Young and Bion, 1981; 1982). Some authors indicate that floor and ceiling effects hide the inversion effect in children (Carey and Diamond, 1994). The present study controlled for floor and ceiling effects by presenting sufficiently small number of stimuli to allow the youngest children to adequately complete the task and also a sufficiently large number of stimuli to permit errors in the oldest participants (c.f., Ellis, 1992a). Examining the means presented in Figure 2.2., it is clear that the observed values of d' are within the accepted range for this statistic. Floor effects would be shown if the values of d' were not significantly different from zero and ceiling effects would be demonstrated with a d' of 4. Neither ceiling nor floor effects affected the present study.

2.2. Experiment 2 – The Own-Group Biases Developmentally

Having established that the inversion effect is not as robust in children as it is in adults, Experiment 2 was conducted to explore the ORB in children. A standard recognition paradigm was employed similar to that used in Experiment 1, testing children of the same

age range. Own- and other-race and own- and other-age faces were presented to participants. These biases could thus be compared.

2.2.1. Method

2.2.1.1. Participants

Participants were: 10 five-year-old (four male); 14 six-year-old (nine male); 14 seven-year-old (six male); 11 eight-year-old (eight male); 10 nine-year-old (five male); 16 10-year-old (12 male); 13 11-year-old (seven male); 12 12-year-old (four male); 11 13-year-old (five male); and nine 14-year-old (three male) children recruited from a sample who returned consent forms to their school and 13 adults (two male; mode age = 18) who volunteered for this study. These were pooled into larger groups (aged 5- and 6-years, 7- to 9-years, 10- to 12-years, 13- & 14-years, and adult). Uneven group size was a product of the recruitment process. All participants were ethnically White and self-reported that they had normal or corrected vision. All of the children were considered typically developing.

2.2.1.2. Materials

Forty-eight faces from the Minear and Park (2004) face database were used in this Experiment. These were: 16 Black 20- to 22-year-old male and female faces; 16 White 20- to 22-year-old male and female faces; and 16 White 60- to 80-year-old male and female faces. In addition, 16 White 8-year-old faces were used from those used by Bruce, Campbell, Doherty-Sneddon, Import, Langton, McAuley, and Wright (2000)[6]. The images were cropped and adjusted such that they were the same dimensions and

---

[6] I would like to thank Vicki Bruce and Debbie Riby for allowing us to use these stimuli and providing them to use so promptly.

resolution as the first database. All faces in this Experiment were presented in greyscale using the same software and hardware as in Experiment 1.

2.2.1.3. Design and Procedure

The design for Experiment 2 was similar to that of Experiment 1. Participants saw all four types of faces, thus creating a 4 by 5 mixed factorial design with the factors: type of face and age of participant. All other aspects of the design and procedure were identical to that used in Experiment 1.

2.2.2. Results

Mean recognition accuracy (d'), for each type of face split by age, is presented in Figure 2.3. This figure shows that face recognition abilities improve with age, and at younger ages, participants show no recognition advantage for faces of their own-race. Each bias was analysed separately in three separate parallel 2 by 5 mixed factorial ANOVAs with the factors: type-of-face and age-of-participant. Only faces crucial to the calculation of the bias independently of other biases were included in each analysis.

2.2.2.1. The own-race-bias

The first ANOVA compared the types of face: White (20- to 22-year-old own-gender) and Black (20- to 22-year-old own-gender). This revealed a main effect of age, $F(4, 128)$ = 9.639, MSE = 0.669, $p < .05$, whereby older participants were more accurate than younger participants as explored by Tukey HSD *post hoc* tests. Adults had higher recognition accuracy than: 5- & 6-year-old children (mean difference = 1.135, $p < .05$); 7- to 9-year-old children (mean difference = 0.754, $p < .05$); and 10- to 12-year-old

children (mean difference = 0.748, $p < .05$). In addition, the youngest age group tested had a lower accuracy than 7- to 9-year-old children (mean difference = 0.380, $p < .05$), 10- to 12-year-old children (mean difference = 0.387, $p < .05$), and 13- & 14-year-old children (mean difference = 0.735, $p < .05$). There was also an effect of type-of-face, $F(1, 128) = 8.261$, MSE = 0.394, $p < .05$, whereby White faces were recognised more accurately than Black faces (mean difference = 0.240). The interaction between type-of-face and age-of-participant was approaching significance, $F(4, 128) = 1.992$, MSE = 0.394, $p < .10$. Simple effects were used to explore this interaction and revealed that the ORB was only significant for adults (mean difference = 0.640, $p < .05$) and for 13- & 14-year-old children (mean difference = 0.431, $p < .05$). No other simple effects were significant.
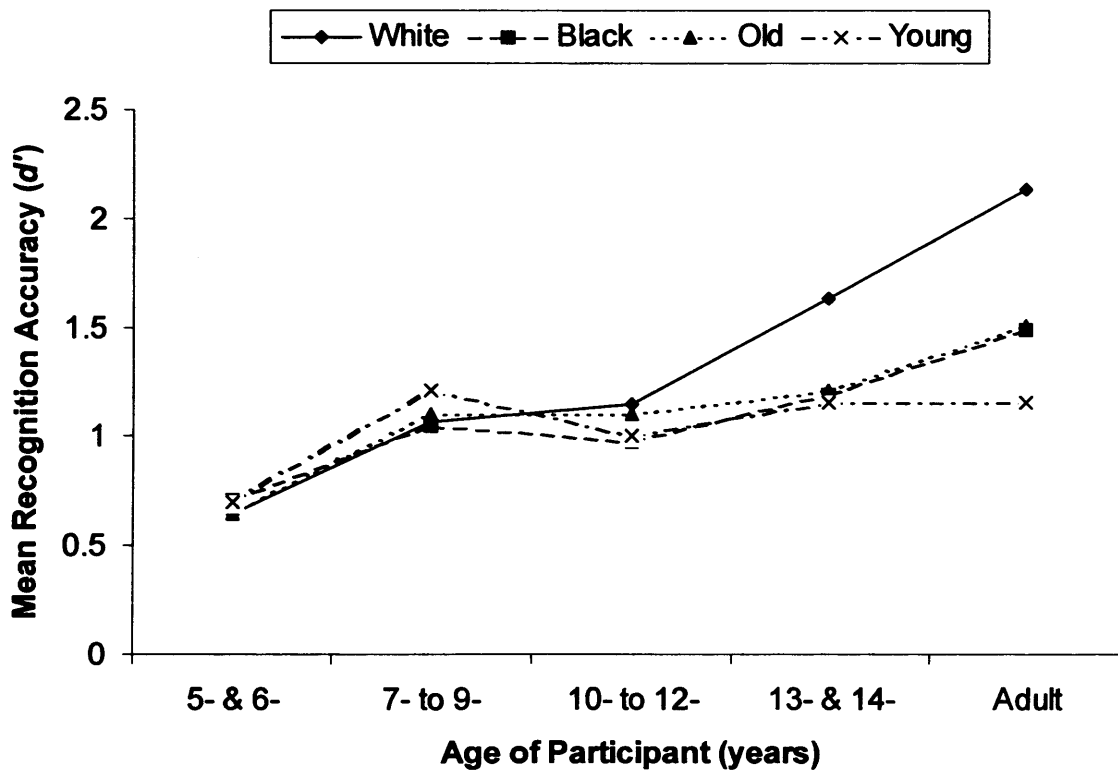
Figure 2.3. Mean recognition accuracy (d') for the four types of faces, split by age. Unless otherwise stated, the faces are White.

Due to Feinman and Entwisle (1976) demonstrating an ORB in children when analysing hits only, an analysis on the hit rate was also conducted. This revealed a significant effect of age-of-participants on hits, $F(4, 128) = 4.969$, MSE = 0.054, $p < .05$. Tukey HSD tests revealed that 5- and 6-year-old children had a lower hit rate than 13- and 14-year-old children (mean difference = .174, $p < .05$) and adults (mean difference = .220, $p < .05$). There was a significant main effect of type-of-face, $F(1, 128) = 4.784$, MSE = 0.025, $p < .05$, where White faces received a higher hit rate than Black faces (mean difference = .046). There was also a significant interaction, $F(4, 128) = 2.967$, MSE = 0.025, $p < .05$. Simple effects revealed that adults had a higher hit rate for own-race faces than other-race faces (mean difference = 0.171, $p < .05$). No other simple effects were significant.

2.2.2.2. The own-age bias

The analysis comparing the 60- to 70-year-old White own-gender faces with 22-year-old White own-gender faces also revealed a significant main effect of age, $F(4, 128) = 13.394$, MSE = 0.505, $p < .05$. Tukey HSD comparisons revealed that adults had a greater accuracy than: 5- & 6-year-old children (mean difference = 1.172, $p < .05$); 7- to 9-year-old children (mean difference = 0.736, $p < .05$); 10- to 12-year-old children (mean difference = 0.698, $p < .05$); and 13- & 14-year-old children (mean difference = 0.400, $p < .05$). The youngest age group had a lower mean recognition accuracy than: 7- to 9-year-old children (mean difference = 0.436, $p < .05$); 10- to 12-year-old children (mean difference = 0.474, $p < .05$); and 13- & 14-year-old children (mean difference = 0.772, $p < .05$). The analysis also revealed an effect of type of face, whereby 22-year-old faces were better recognised than 70-year-old faces, $F(1, 128) = 6.708$, MSE = 0.370, $p < .05$ (mean difference = 0.210). The interaction between type-of-face and age-of-participant was marginal, $F(4, 128) = 2.157$, MSE = 0.370, $p > .08$. Simple effects revealed that

adults demonstrated a significant own-age bias (mean difference = 0.622, $p < .05$) and that 13- & 14-year-old children demonstrated a marginal own-age bias (mean difference = 0.413, $p < .08$).

2.2.2.3. The own-age bias

The third analysis compared the accuracy of recognition for 22-year-old White own-gender faces with 8-year-old White own-gender faces. This revealed a main effect of age of participant, $F(4, 128) = 8.598$, MSE = 0.579, $p < .05$. Tukey HSD comparisons revealed that adults had greater recognition accuracy than: 5- & 6-year-old children (mean difference = 0.972, $p < .05$); 7- to 9-year-old children (mean difference = 0.506, $p < .05$); and 10- to 12-year-old children (mean difference = 0.566, $p < .05$). The youngest age group tested had a lower recognition accuracy performance than: 7- to 9-year-old children (mean difference = 0.466, $p < .05$); 10- to 12-year-old children (mean difference = 0.406, $p < .05$); and 13- & 14-year-old children (mean difference = 0.721, $p < .05$). The analysis also revealed a significant effect of type of face, $F(1, 128) = 9.637$, MSE = 0.457, $p < .05$, whereby 22-year-old faces were better recognised than 8-year-old faces (mean difference = 0.279). There was also a significant interaction between type-of-face and age-of-participant, $F(4, 128) = 4.041$, MSE = 0.457, $p < .05$. This interaction was due to 8-year-old children being better at recognising 8-year-old faces than 22-year-old faces though not significantly (mean difference = 0.140, $p < .20$), whereas adults were better at recognising 22-year-old faces than 8-year-old faces (mean difference = 0.976, $p < .05$) as were 13- and 14-year-old children (mean difference = 0.470, $p < .05$).

2.2.3. Discussion

The most reliable finding from Experiment 2 was that face recognition performance improves with age. The ORB was not observed until age 13 in the present experiment as demonstrated by the simple effects analysis.

The experiments reported thus far indicate that children do not show the same detriment in face recognition that adults do to certain subgroups of faces. In this way, the face recognition system of children appears to be more flexible than that of the adults. Nevertheless, adults are more accurate than children in the recognition of all types of faces, indicating a general cognitive improvement with age. The face recognition system of children will encode subtypes of faces less frequently encountered whereas the adult face recognition system is set up primarily for those most frequently encountered faces. This proposal will be explored in Experiment 3.

2.3. Experiment 3 – Adaptation to "Unnatural" Facial Configurations in Children and Adults

Experiments 1 and 2 indicated that the immature face recognition system is characterised by the ability to recognise inverted and other-race faces at the same levels as upright and own-race faces. Inverted faces are an "unnatural" class of faces as they violate first-order relational information. Thus, it can be hypothesised that children are better able to process other "unnatural" facial configurations than adults.

Faces are typically made up of the same configuration of features, in which there are population variations in the make-up of these faces. Adults appear to be unconsciously aware of these population variations, as explained in Section 1.2.3. Indeed,

82

Robbins et al. (2007) demonstrated that it was not possible to adapt adult participants to an "unnatural" facial configuration. In their study, participants were adapted to a stimulus with eyes shifted up but still aligned (a "natural" facial category) or to a stimulus with one eye shifted up and the other eye shifted down (an "unnatural" facial category). The effect of adaptation transferred from the adaptor identity to other faces, indicating a shift in the dimensions of face-space used to encode faces (Rhodes, et al., 2004).

An explanation for this may be that adults cannot be adapted to an "unnatural" facial category. Since eyes are aligned in virtually all faces, one eye shifted up is unlikely to be processed as a face. If children have a more flexible face recognition system, then they may be adapted to this "unnatural" face configuration as well as the more "natural" configuration. Moreover, since children have experienced fewer faces than adults they may well be more likely to show adaptation effects. To look into this, a standard face distortion adaptation procedure (c.f., Fox and Barton, 2007) shall be tested on adults and children. Due to the length of testing, a between-subjects design was employed, whereby each participant was either adapted to a "natural" or "unnatural" facial configuration. As such, only two ages groups were tested, adults and children. Participants above the age of 12 were considered adults and participants below the age of 12 were considered children due to the blip in performance found in the previous experiments occurring at the age of 11. A third type of adaptor was employed, whereby both eyes were shifted in the same direction but by different amounts. Thus, three adaptation conditions were tested: asymmetric (one eye up and one eye down); symmetric (both eyes shifted together); and mixed (both eyes shifted in the same direction but by different amounts).

2.3.1. Method

2.3.11. Participants

The participants were 34 children, of which six participated in the mixed condition (mean age = 10 years)[7], 13 participated in the symmetric condition (mean age = 9.2 years), and 15 participated in the asymmetric condition (mean age = 9 years), and 50 adults, of which 15 participated in the mixed condition (mean age = 16 years), 15 participated in the symmetric condition (mean age = 16.2 years), and 20 participated in the asymmetric condition (mean age = 15.96 years). All participants were White, considered typically developing by their school, and had normal or corrected vision.

2.3.1.2. Materials

Two faces from the Minear and Park (2004) were used as the stimuli. These were full frontal views in neutral expressions. These were distorted in that both eyes were shifted using a Java™ computer program written for the purposes of this experiment[8]. The program allowed for each eye to be moved up or down individually by any number of pixels the user required. In addition, it stretched the area around the eyes such there were no box or contrast marks. Three types of face stimuli were created. One set had one eye shifted up and the other eye shifted down by the same amount (type asymmetric). A second had both eyes shifted up or down (type symmetric). The final had both eyes shifted in the same direction first, then one eye shifted down by 8 pixels (type mixed). Twenty seven faces were created in each set, with the extreme distortion being +26 pixels or -26 pixels. Each picture had a resolution of 96 dpi, and was 639 pixels wide by 490 pixels high. The series of faces for asymmetric and symmetric and their respective

---

[7] It must be noted that 6 participants in this condition is too few to be completely confident of the subsequent data.

[8] This computer program was written by Andrew M Holland and is available from the author.

distortion values (measured in pixels) is shown in Figure 2.4. All stimuli were presented to participants using Superlab Pro 2^{TM} research software on a Toshiba Tecra^{TM} M4 Tablet PC.

Figure 2.4. *Top panel*. The asymmetric series of distorted faces ranging from left eye up 26 pixels – right eye down 26 pixels to left eye down 26 pixels – right eye up 26 pixels in equal steps of -2 and +2 pixels. The undistorted face is the farthest right on the second line. *Bottom panel*. The symmetric series of distorted faces ranging from both eyes moved 26 pixels up to 26 pixels down in equal steps of 2 pixels. The undistorted face is the farthest right on the second line.

2.3.1.3. Design

A between-subjects design was employed, whereby participants were randomly allocated to be presented with the asymmetric series, the symmetric series, or the mixed series. The second IV was participant age (adult or child). The dependent variable was magnitude of adaptation, measured by the change in oddness ratings of the series of faces post-adaptation from pre-adaptation. The order of the faces presented was randomised. The facial identity of the adaptor was counterbalanced, such that both faces appeared as the adaptor in each condition an approximately equal number of times. Participants were either adapted to the extreme positive or extreme negative stimulus in the series and this was counterbalanced across participants.

2.3.1.4. Procedure

Experimentation took place in school libraries and was set up as shown in Figure 2.1. Participants were first introduced to the undistorted image of one of the identities and were given a name and a short story about the face. The face remained on screen for the duration of this. The story lasted no longer than one minute. Following this, the participants were presented with each of the faces of both facial identities in the series they had been allocated to participate in. The participants were instructed to say how odd they thought the face looked on a 1 to 9 scale, whereby 9 was very odd and 1 was normal. The participants gave their response verbally, and the experimenter entered the number into the computer keyboard. Each presentation was response terminated. Participants only saw faces in the direction of the distortion they were to be adapted to and the undistorted face. Moreover, participants saw faces from the identity they were to be adapted to and the identity they were not to be adapted to. Thus, each participant saw 28 faces.

Once each of the 28 faces had been presented, the participants were asked the same filler questions as in Experiment 1. Immediately following this, participants were presented with the adaptor. The adaptor was the most extreme distortion from the series they had previously rated. The participants were instructed to examine the face for the duration it was on screen. The adaptor remained on screen for 60 seconds.

Immediately following the adaptor, the participants were presented with the same 28 faces they had rated and asked to rate how odd each one was. Intermixed between each test face was the adaptation face that the participants also rated. The test faces remained on screen until the participants made a verbal response, which the experimenter typed on the keyboard. The intermixed adaptation faces remained on screen for four seconds each (c.f., Rhodes, et al., 2003). Once all 28 faces had been viewed, the participants were thanked.

2.3.2. Results

Differences in oddness ratings before and after adaptation were computed according to the formula:

$$\text{Adapt} = \text{Odd}_{\text{before}} - \text{Odd}_{\text{after}} \qquad [2.1]$$

whereby the magnitude of adaptation was the oddness rating before adaptation minus the oddness rating after adaptation. A low number indicates little adaptation, whereas a higher number indicates greater adaptation. These data are presented in Figure 2.5.
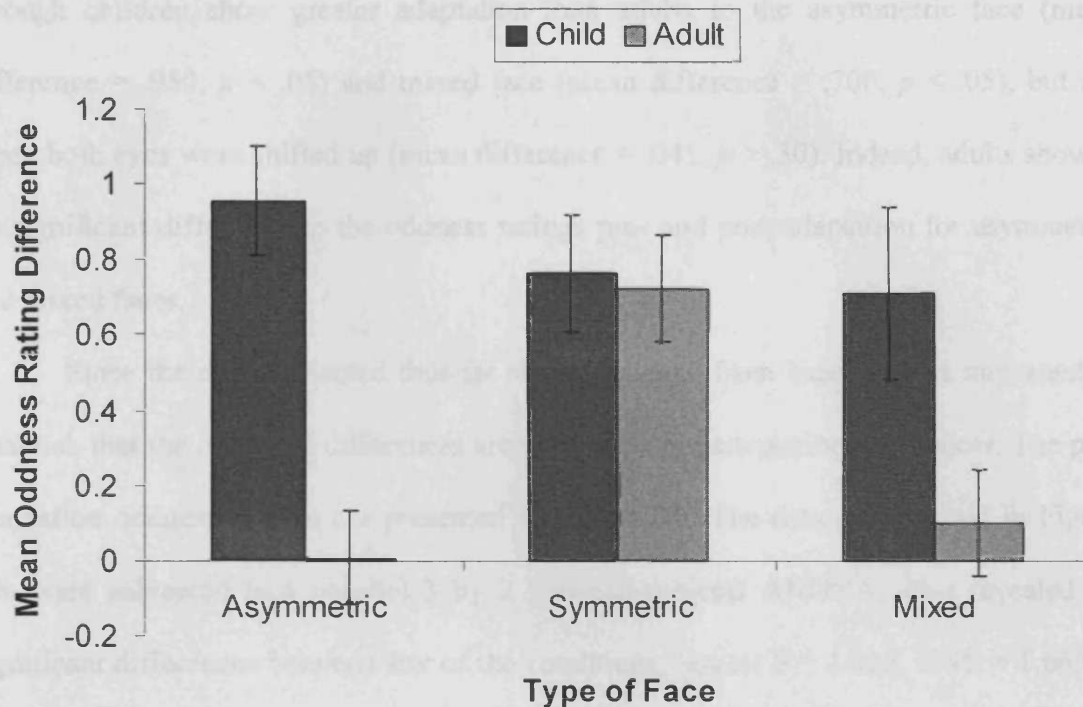
Figure 2.5. Mean difference in oddness rating pre- and post-adaptation for the three face types and participant age. Zero indicates the oddness ratings were the same pre- and post-adaptation. Error bars represent standard error of the mean.

The data summarised in Figure 2.5 were entered into a 3 by 2 between-subjects ANOVA with the factors: type-of-stimulus and age-of-participant. This revealed a significant effect of age-of-participants, F (1, 78) = 16.835, MSE = 0.309, $p < .05$, whereby children were more affected by adaptation than adults (mean difference = 0.534). The main effect of type-of-face was approaching significance, $F(2, 78) = 2.530$, MSE = 0.309, $p < .10$, whereby more participants were adapted to faces that had both eyes shifted than single eye (mean difference = .262, $p < .10$) and mixed (mean difference = .338, $p < .05$) after testing with Tukey HSD *post hoc* tests. These main effects are qualified by a significant interaction between type-of-face and age-of-participant, $F(2, 78) = 5.170$, MSE = 0.309, $p < .05$. This interaction revealed itself

through children show greater adaptation than adults to the asymmetric face (mean difference = .950, p < .05) and mixed face (mean difference = .700, p < .05), but not when both eyes were shifted up (mean difference = .041, p > .30). Indeed, adults showed no significant difference in the oddness ratings pre- and post-adaptation for asymmetric and mixed faces.

Since the data presented thus far show a change from baseline it is important to establish that the observed differences are not due to pre-adaptation differences. The pre-adaptation oddness ratings are presented in Figure 2.6. The data summarised in Figure 2.6. were subjected to a parallel 3 by 2 between-subjects ANOVA. This revealed no significant differences between any of the conditions, largest F = 1.922, MSE = 1.665, p > .17. This indicates that the effects demonstrated in Figure 2.5. are not due to differences in the baseline oddness ratings.
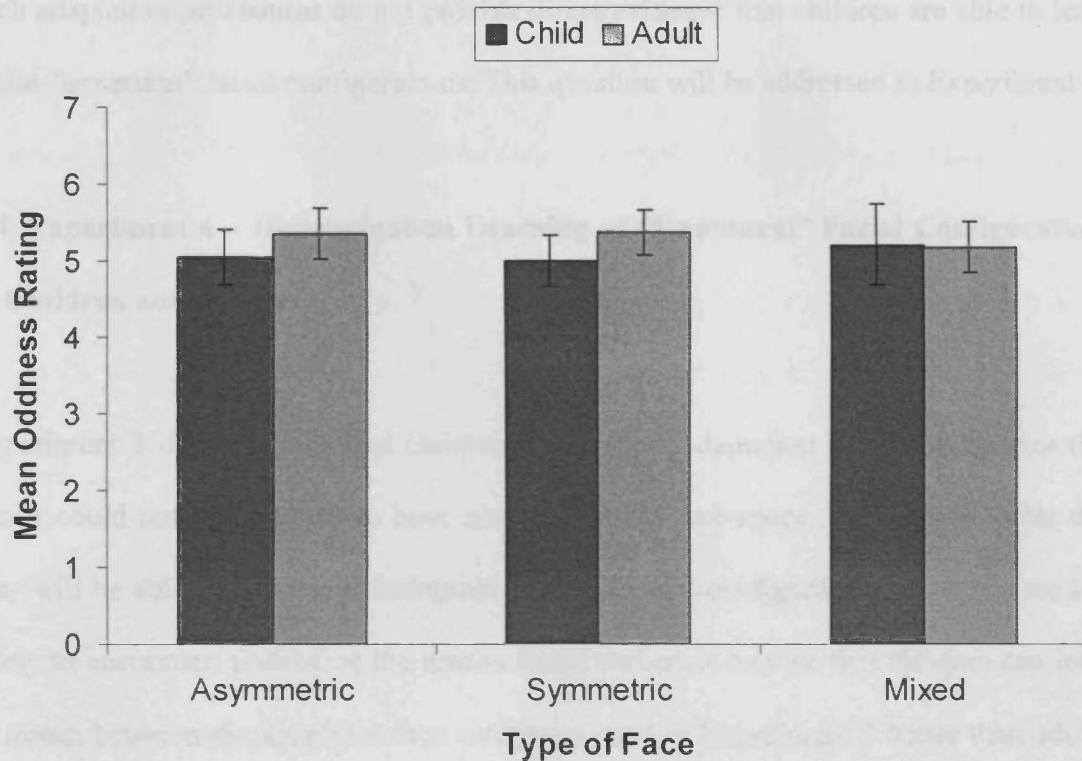
Figure 2.6. Mean oddness rating pre-adaptation for the three face types and participant age. Higher numbers indicate an odder-looking face. Error bars represent standard error of the mean.

2.3.3. Discussion

Children had their oddness ratings reduced following adaptation to a face that had one eye moved in one direction and the other eye moved in the opposite direction (an "unnatural" configuration), both eyes moved by the same amount (a "natural" configuration), or both eyes moved by a different amount (an "unnatural" configuration), whereas adults could only be adapted if both eyes were shifted by the same amount. The results for adults are consistent with those reported by Robbins et al. (2007). The present study shows a clear difference in the adaptability of adults and children. Nevertheless,

such adaptation procedures do not provide direct evidence that children are able to learn about "unnatural" facial configurations. This question will be addressed in Experiment 4.

2.4. Experiment 4 – Discrimination Learning of "Unnatural" Facial Configurations in Children and Adults

Experiment 3 demonstrated that children could show adaptation to face categories that adults could not. If children do have a more flexible face-space, then it is possible that they will be able to learn and distinguish between facial configurations that they are less likely to encounter. Extending the results found earlier, it may be that children can learn to match between the unnatural face categories used in Experiment 3 better than adults. Experiment 4 tested this with a matching paradigm using simultaneous presentation. The same asymmetric, symmetric and mixed types of faces used in Experiment 3 were used here. From the data observed in Experiment 3, it was predicted that children would be able to match all three types of faces equally well, whereas adults would be better at the symmetric faces than the asymmetric and mixed faces.

2.4.1. Method

2.4.1.1. Participants

Seventy-one children (mean age 8.2 years, range 5 to 11 years) and 48 adults (mean age 17 years, range 12 to 21 years) participated in this experiment. All were unpaid volunteers recruited from local educational establishments. All self-reported that they were White, typically developing, and had normal or corrected vision.

2.4.1.2. Materials

Thirty-three faces randomly selected from the Minear and Park (2004) face set were used in this experiment. These were of White males and females aged between 20 and 22 years and in frontal neutral pose. All thirty-three were cropped and the external features were masked. All the faces were then distorted using the programme described in Experiment 3. Eleven were created with the asymmetric arrangement, 11 with symmetric eyes shifted up, and 11 with the mixed arrangement.

A look-a-like was created for each of the 33 faces in which the eyes were shifted up by 6 pixels. Thus, a total of 66 faces were constructed (a target face and a look-a-like). All faces were presented with the dimensions 100 mm by 110 mm with 72 dpi resolution. They were presented to participants using Microsoft™ PowerPoint™.

2.4.1.3. Design

A 3 x 2 mixed design was employed, whereby each participant saw all three types of faces (asymmetric, symmetric, and mixed). The between-subjects independent variable was age, split by participants over 11 being adults and under 12 being children. The presentation order of the faces was randomised. The dependent variable was discrimination performance measured by proportion of correct responses.

2.4.1.4. Procedure

Testing took place in school classrooms, and the same set-up was used here as shown in Figure 2.1. Participants were presented with 33 trials in a random order. Each trial consisted of three faces; one positioned above the other two. Participants were told that one of the two matched the one above. Alternative synonyms were used if the participant did not understand. Participants were instructed to point to left or right face of the two as

the impostor or "the different one". The faces remained on screen until the participants responded. As there are only two options, chance rate was 0.5. An example of the trial screen is presented in Figure 2.7.

Figure 2.7. An example of a trial screen used in Experiment 4. The face above is a target face. The instruction to the participants was to identify which one of the bottom faces was different from the top face. In this case it is the left face.

2.4.2. Results

The results were in terms of proportion of correct decisions made and are presented in Figure 2.8. Children appear to be better able to discriminate asymmetric and mixed type

93

faces better than adults, whereas adults appear better at discriminating between symmetric type faces.
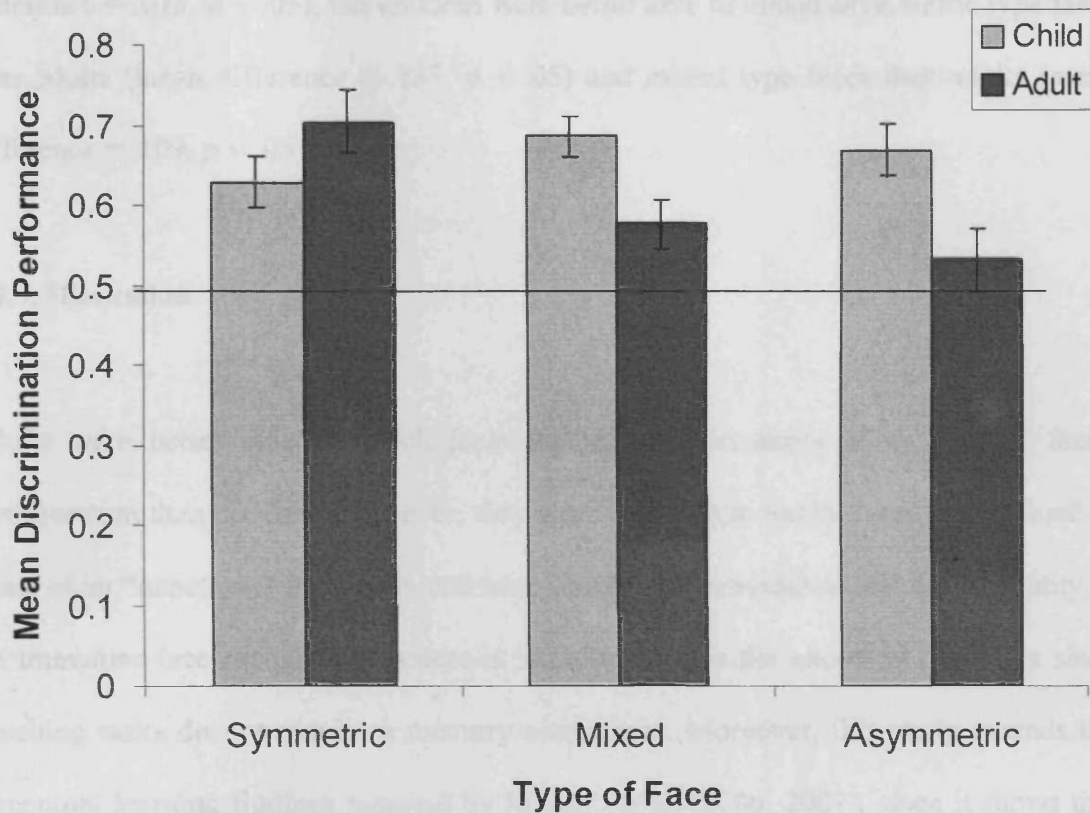
Figure 2.8. Mean performance on identification of the target face for the three types of faces, split by participant age. Chance levels would be 0.5, shown by the horizontal line. Error bars represent standard error.

The data summarised in Figure 2.8. were subjected to a 3 by 2 ANOVA with the factors: type-of-face and age-of-participant. This revealed a significant effect of age, $F(1, 117) = 4.369$, MSE $= 0.064$, $p < .05$, whereby children were better able to match faces than adults (mean difference $= 0.057$). There was also a marginal main effect of type of face, $F(2, 234) = 2.877$, MSE $= 0.045$, $p < .06$, whereby symmetric faces were better matched than asymmetric faces (mean difference $= 0.067$, $p < .10$). These main effects were qualified by a significant interaction between face-type and participant-age, $F(2, 234) =$

8.378, MSE = 0.045, $p < .05$. This interaction was explored using simple effects. These revealed that adults were better able to match symmetric type faces than children (mean difference = .075, $p < .05$), but children were better able to match asymmetric type faces than adults (mean difference = .137, $p < .05$) and mixed type faces than adults (mean difference = .109, $p < .05$).

2.4.3. Discussion

Adults were better able to match faces that differed in terms of a "natural" facial configuration than children. However, they were less able to match faces that differed in terms of an "unnatural" facial than children. This provides evidence that the flexibility of the immature face recognition system is likely to be due the encoding processes since matching tasks do not require a memory component. Moreover, this study extends the perceptual learning findings reported by Mundy, et al. (2006, 2007), since it shows that some features can be learnt more readily than others. When presented with two faces, differences in internal features are easier to discriminate *and* changes which maintain the first-order relational information, in adults but not necessarily children. This simple perceptual learning mechanism is useful in explaining adult face recognition abilities presented here, however does not explain why children can discriminate faces that adults cannot process. For this, one can look to the qualitative and quantative differences in adult and children's face recognition abilities.

The evidence thus far indicates that children have a flexible face recognition system in terms of the number of dimensions that becomes less flexible at the age of roughly 11. There are several possible types of developmental courses that may exist for face recognition. There may be some form of maturational development, whereby when a

child reaches the age of 11, a second, adult type face processing ability takes over from the primary, child type face processing (c.f., Carey and Diamond, 1977). Alternatively, the development may be entirely based on perceptual experience and thus the adult face recognition abilities develop slowly. There may also be a combination of the two.

At this point it is important to explore any alternative explanations for the findings in the present experiment. Since the mixed and asymmetrical facial configurations are "unnatural" in that they are not often experienced, they may be so odd as to be distracting to adults. Since children have experienced far fewer faces, the effect of oddness of the faces would be lower, and thus they would not be distracted by the oddness as much as adults. However, data from Experiment 3 can be used here to counter this argument. Oddness ratings for eyes asymmetrical, mixed, and symmetrical conditions are extremely similar for both adults and children. Adults and children do not appear to rate asymmetrical faces as odder than symmetrical distorted faces. Thus, this explanation is unlikely to be valid.

2.5. General Discussion

Four experiments here tested face recognition abilities in children and compared theirs to adults'. Experiment 1 indicates that the magnitude of the inversion effect increases with age, reaching adult levels at age 13. Experiment 2 found that adult levels of the own-race bias are also found at age 13. Experiment 3 demonstrated that children could be adapted to "unnatural" facial configurations, whereas adults could not be. Experiment 4 demonstrated that children could also discriminate between two faces that had "unnatural" facial configurations at a much higher degree of accuracy than adults, though this pattern was reversed for natural facial categories.

The data collected from the four experiments in the present Chapter advance our knowledge of face-space dramatically. The fact that children show a smaller inversion effect and smaller own-group biases is indicative of a less rigid face recognition system in children than in adults. The fact that children can be adapted to and learn about "unnatural" facial configurations is also supportive of this claim. These findings are largely incompatible with two of the immature face-space models described in Section 1.4.

Constant-face-space predicts all effects observed in adult face recognition will be observed in children. This does not happen as demonstrated in all the experiments observed here. One may assume that children are less accurate in their encoding of faces, and thus each face is encoded with larger error regions. Thus, the lack of inversion effects and the ORB may be due to floor effects in children's face processing abilities. This seems implausible for two reasons. Firstly, children were superior to adults in discriminating "unnatural" facial configurations in Experiment 4. If all faces were stored with greater error then children should not be able to perform at levels near to adult level on any face types, let alone perform better than adults for these two subgroups of faces (asymmetrical and mixed, see Experiment 3). Secondly, the data from Experiments 1 and 2 show that the face recognition abilities of children was significantly above chance, suggesting that floor effects did not play a role in the present experiment. Nevertheless, it may be suggested that the number of faces tested may not be sufficient to be sensitive enough to detect differences between upright and inverted faces in the youngest age groups. While this is possible, the number of stimuli were chosen based on previous data (Ellis, 1992a) and were not affected by floor effects. This suggests that the experiments were sensitive enough to detect differences in recognition should they have existed.

97

An assumption of expanding-face-space is that children use fewer dimensions to recognise faces than adults. New dimensions are added with development to aid in the recognition of similar faces. This version of face-space was also not supported for two reasons. Firstly, it would predict that the original dimensions are mostly suited for recognition of all faces, and the new ones that are added are only beneficial to the most frequently encountered faces (upright own-race faces). The second reason why an immature face-space with fewer dimensions than adults' is unlikely is that it was possible to adapt children to dimensions that it was not possible to adapt adults to. If adults retain all dimensions present in children then they should be able to learn "unnatural" facial configurations.

With two versions of the immature face-space largely discounted, the discussion leads inextricably to the validation of shrinking-face-space. Recall that shrinking-face-space suggests that there are more dimensions in the immature-face-space than the adult face-space. Due to these additional dimensions, children are able to recognise inverted faces (Experiment 1), other-group faces (Experiment 2), and "unnatural" facial configurations (Experiments 3 and 4) to the same degree as upright, own-race faces and "natural" facial configurations. As such, shrinking-face-space is able to account for the findings in this Chapter.

Conclusions

The results from these experiments are interpreted as being consistent with a model of immature face-space that contains many more dimensions than the adult face-space. These are indicative of greater flexibility of the child's face-space, as demonstrated by the ability of children to be adapted to and learn to discriminate faces based on

dimensions that adults cannot be adapted to or learn to discriminate. With development, the dimensions of face-space become more focused, such that only the most diagnostic ones are attended to by adulthood. Nevertheless, these dimensions may well still be contained within face-space. The subsequent chapters will explore what happens to the dimensions that are no longer attended to in adulthood, and whether these are extinct, dormant, or active.

Chapter 3 – Attentional Influences on Face-Recognition

Chapter 2 demonstrated that the immature face-space has more dimensions than the adult face-space. It was hypothesised that perceptual narrowing caused the least diagnostic dimensions to be discarded or ignored. Thus, the most diagnostic dimensions for the most frequently encountered faces become those used by default in the adult face-space. The less diagnostic dimensions may be deleted from face-space or may remain but are unattended to, and thus dormant. Eight experiments in this chapter explore what happens to dormant dimensions of face-space. Experiment 5 employed a priming paradigm to activate a region of the face-space useful for the encoding of stereotypical faces. Experiment 6 tested participants' ORB levels after a perceptual learning task involving training to other-race critical features. Experiments 7 to 9 used a cueing paradigm to shift attention when processing own- and other-race faces. Experiments 10 to 12 employed a similar cueing paradigm to shift attention when processing upright and inverted faces.

3.1. Experiment 8 – Stereotype Priming Activating Dormant Dimensions

Subgroups of faces can have their discriminability selectively improved by the accompaniment of a congruent stereotype label (Klatzky, et al., 1982a, b). This discriminability was confounded with a decrease in response criterion, whereby their participants were more likely to say "yes" in the recognition test to both stereotype-congruent targets (hits) and distractors (false alarms). Klatzky et al. used an uneven ratio of targets to distractors in the recognition phase of their experiment, which may have

100

caused the above pattern of results. Notwithstanding these problems, if one could establish that the presentation of semantic information affects the encoding of stereotype-congruent faces then it would indicate that semantic information can improve the encoding of particular faces. Experiments 5a and 5b aim to provide a demonstration of the stereotype priming effect and to examine its locus.

3.1.1. Phase 1: Development of Stimulus Set

Before the crucial experimentation can be conducted, a collection of face stimuli that could reliably be labelled according to various stereotypes were created.

3.1.1.1. Method

3.1.1.1.1. Participants

Four male and six female Cardiff University Psychology undergraduates participated as a partial fulfilment of a course requirement.

3.1.1.1.2. Materials

Two-hundred-and-fifty faces were created using Faces™ software package produced by Interquest™. This software package produces realistic pictures of faces from sets of feature pools[9]. These were presented to the participants using Microsoft™ Powerpoint™ software.

---

[9] Artificial faces were used as it was difficult to find a sufficient number of real faces that were reliably given occupation labels. However, in other studies that used real faces but only one stereotype label, the results were similar to those reported here (Hills, Lewis, & Honey, 2005).

3.1.1.1.3 Design and Procedure

Faces were presented sequentially and in a random order. Participants were instructed to write down on a separate sheet of paper which of 10 stereotype labels most accurately represented each face. The labels were actor, artist, banker, criminal, hairdresser, lawyer, philosopher, and politician. These labels were taken from Bull and Green (1980). Once the participant had made their judgement, they moved the Powerpoint™ presentation on one slide. Thus, each face had ten stereotype labels attached to it, one from each participant.

3.1.1.2. Results

All faces that had less than five of the same label given to them were discarded. Moreover, those which had five of one label and five of a separate label were removed (this occurred in six cases, all of which were judged to be either lawyers or criminals). The remaining 186 faces were used in the next phase. Only four stereotype labels were used consistently: actor, artist, banker, and criminal.

3.1.2. Phase 2: Development of Stimulus Set

The second phase was conducted in order to ensure both that the faces were reliably associated with a given label and that all categories of faces were of the same distinctiveness.

3.1.2.1. Method

3.1.2.1.1. Participants and Materials

Ten more Psychology undergraduates from Cardiff University took part in this study in partial fulfilment of a course requirement. Three were male and seven were female. All 186 faces created in Phase 1 were used in Phase 2.

3.1.2.1.2. Design and Procedure

Participants received 5 blocks of presentations. In each block the 186 faces were presented in random order and the participants were instructed to give each face a rating. In the first block the participants were asked to judge how much each face looked like the stereotype of an actor (using a 1 to 9 scale, where 9 was much like the stereotypical actor and 1 was unlike the stereotype). In the subsequent 3 blocks, participants made the same form of rating, but with respect to the labels artist, banker and criminal. In the final block, participants provided a distinctiveness rating of the form of a rating concerning how easy each face would be to spot in a crowd (on a 1 to 9 scale, where 1 is difficult to spot in a crowd (i.e., typical), and 9 is easy to spot in a crowd (distinctive); Light, et al., 1979). The order of the blocks was randomised, with the constraint that none of the participants received the same order.

3.1.2.2. Results

All faces that had a stereotype rating of greater than 4.5 were given that stereotype label. All faces that scored 4.5 on more than one stereotypical scale were removed from further consideration. This process left an uneven number of faces in each of the four stereotype

categories. Any face that scored more than 4 on distinctiveness was removed from the subsequent analysis. Finally, the 20 faces with the highest stereotype rating in each stereotype category were employed in Experiments 5a and 5b.

A one-way ANOVA was conducted on the strength of stereotype ratings for the 4 sets of 20 faces. This analysis revealed that strength of stereotype ratings did not differ across categories $F(3, 76) = 1.895$, MSE = 0.334, $p > .13$. The stereotype ratings were subjected to a regression to ensure that there was no co-linearity between the stereotype labels. This revealed that there were significant negative correlations between the stereotype ratings: bankers-actors, r (80) = -.21, $p < .05$; bankers-criminals, r (80) = -.365, $p < .05$; bankers-artists, r (80) = -.383, $p < .05$; criminals-artists, r (80) = -.322, $p < .05$; and a significant correlation between actors and artists, r (80) = .246, $p < .05$; and no significant correlation between actors and criminals, r (80) = -.074, $p > .51$. To ensure that one category of faces was no more distinctive than any other; the distinctiveness scores for the 4 sets of 20 faces were also subjected to a one-way ANOVA. This revealed that there were no significant differences in distinctiveness across the stereotype categories, $F(3, 76) = 0.620$, MSE = 0.287, $p > .60$. This process of stimulus set development thus produced 20 consistently rated stereotypical actors, artists, bankers, and criminals. An example of each is presented in Figure 3.1. There were no distinctiveness effects across categories and the stereotypes were mutually exclusive, except for artists and actors. These 80 faces were used in Experiments 5a and 5b.

3.1.3. Experiment 5a

Experiments 5a and 5b were identical with the exception that they used different pictures of faces (Experiment 5a: artists and criminals; Experiment 5b: actors and bankers) and

corresponding stereotype labels (Experiment 5a: artist/criminal/banker; Experiment 5b: actor/criminal/banker). In both experiments, participants viewed a subset of the stereotypical faces developed previously. In both experiments, the pictures were accompanied by a congruent (e.g., actor face plus label "actor") or an incongruent/irrelevant occupation label during training (e.g., actor face plus label "banker/criminal", respectively). During the old/new recognition test, the now familiar faces were presented together with other unfamiliar exemplars of the two stereotypes. The test faces were either accompanied by a congruent label (e.g., "artist" for the artist faces) or an incongruent/irrelevant label (e.g., "criminal/banker" for the artist faces). Test performance was assessed using signal detection measures; in particular, the stereotype priming effect was measured in terms of accuracy (d') and response bias (C).
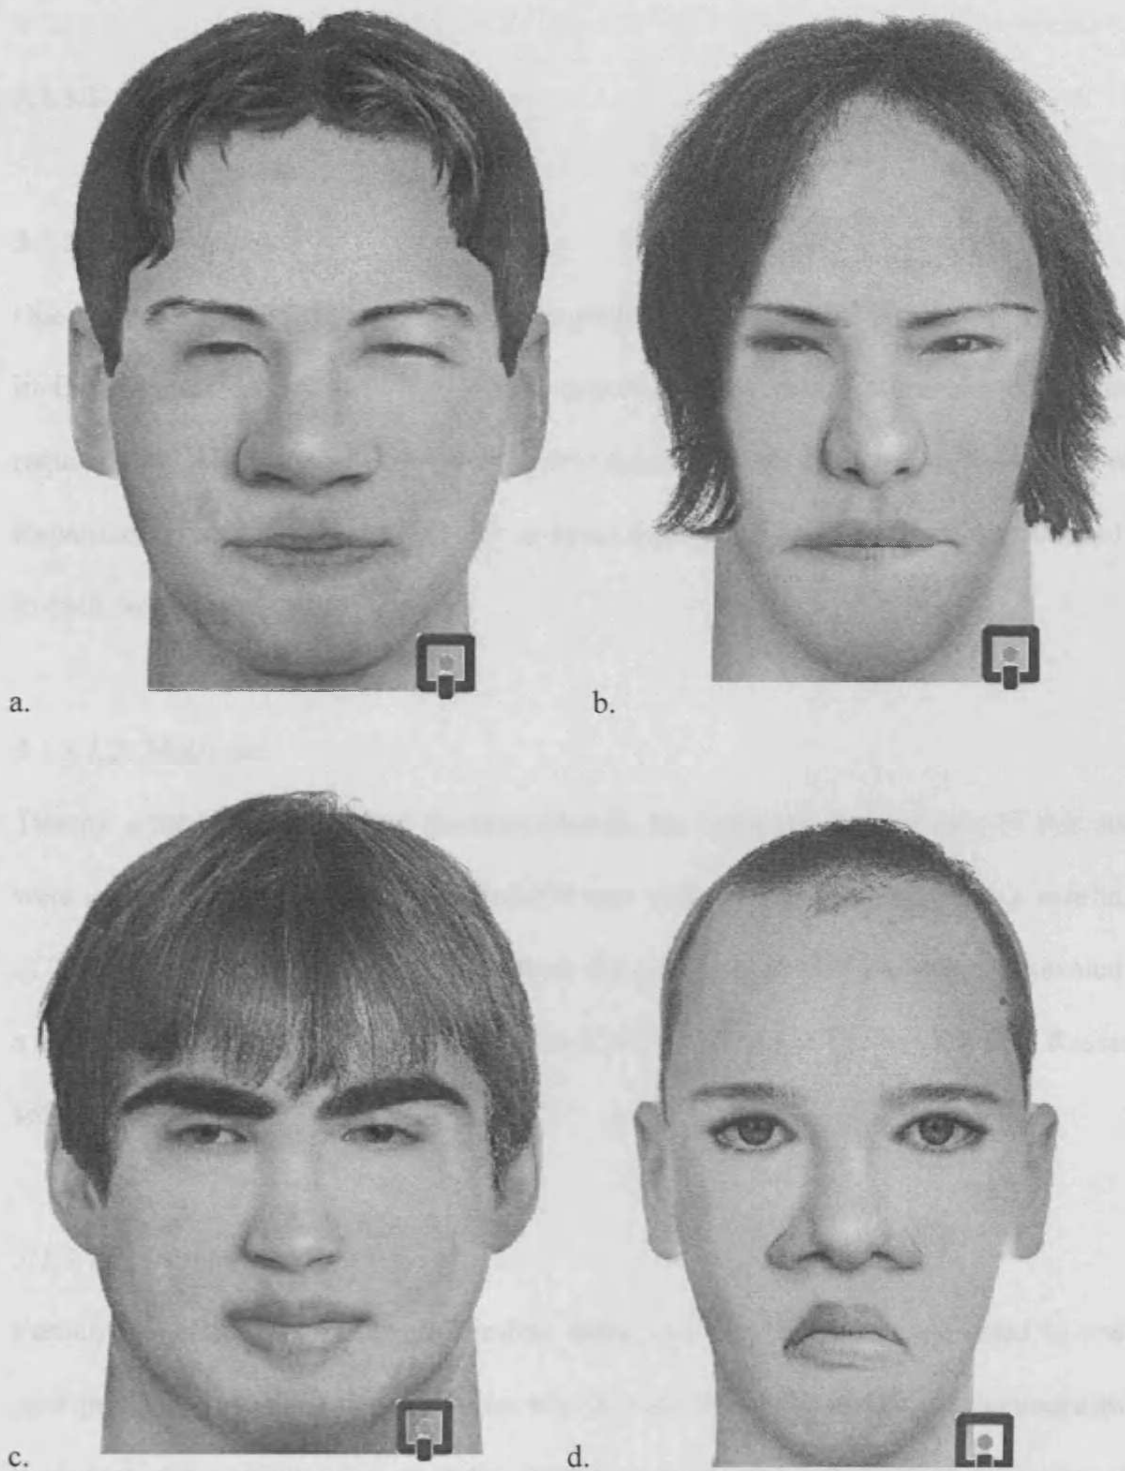
Figure 3.1. An example face from each stereotype category: a. An actor (distinctiveness score = 4.1, stereotype rating = 7); b. An artist (distinctiveness score = 4.6, stereotype rating = 6.7); c. A banker (distinctiveness score = 3.4, stereotype rating = 6); d. A criminal (distinctiveness score = 4.4, stereotype rating = 5.9).

3.1.3.1. Method

3.1.3.1.1. Participants

One hundred and eighty Psychology undergraduates from Cardiff University participated in Experiments 5a and 5b (90 in each experiment) as partial fulfilment of a course requirement. All had normal or corrected vision and were allocated to one of the nine Experimental conditions randomly but in equal numbers. This left equal group sizes of 10 in each condition.

3.1.3.1.2. Materials

Twenty artist and 20 criminal faces created in the stimulus development of this study were used. These faces were presented 130 mm wide by 130 mm high, with a resolution of 72 dpi. Participants sat 50 cm away from the computer screen. They were presented on a high definition colour monitor using an RM PC, running Superlab Pro 2™ Research software.

3.1.3.1.3. Design

Participants saw both artist and criminal faces, but were randomly allocated to one of nine priming conditions depending on whether the label was congruent, incongruent or irrelevant at encoding or whether the label was congruent, incongruent or irrelevant at test. Thus, a 2 (type of face; i.e., artist/criminal) x 3 (label at encoding; i.e., artist/criminal/banker) x 3 (label at test; i.e., artist/criminal/banker) design was employed. The irrelevant label was banker as it corresponded to none of the faces. The dependent variables were recognition accuracy, as measured using *d'*, and response bias, as

measured using C (C is propensity to respond with an "old" response, measured in standard deviation units, indicating the position of the decision criterion with respect to the neutral point where the signal and noise distributions cross over, $C = 0$). Participants were randomly allocated to one of the counterbalanced conditions, with the constraint that each condition had equal numbers of participants within it. There were two sets of artist faces (either designated as old or new), and these sets were divided into six subsets (accompanied by a congruent/incongruent/irrelevant label during encoding or test; ns=6/7).

3.1.3.1.4. Procedure

This experiment employed a standard recognition paradigm similar to that in Experiment 1. Prior to the learning phase, a set of instructions was given to the participants, presented on the computer screen. The instructions were:

> *You will see a set of faces of [stereotype label; e.g., "artists"]. Please rate*
>
> *these faces on how easy they would be to spot in a crowd.*

The participants were provided with the scale (from 1-9) on which to rate face distinctiveness. During the learning phase the participants saw the stereotype label for 250 ms, the label then disappeared and the face was presented for 3000 ms. Participants rated the distinctiveness of the face within this time. Immediately after the learning phase participants were given 25 anagrams to complete within 3 minutes. Finally, the participants were given an old/new recognition task consisting of all 40 faces. The instructions were:

> *You will see a set of faces of [stereotype label; e.g., artists]. Some of them you*
>
> *have seen before and some you haven't. Please indicate whether you've seen*
>
> *each face before by pressing the appropriate key (z=old, m=new).*

Participants were asked to complete this task as quickly and as accurately as they could. Finally all participants were thanked and debriefed.

3.1.3.2. Results and Discussion

The d' recognition accuracy scores are presented in Figure 3.2. In the overall ANOVA there was no effect of label congruence during the test or any interaction involving this factor (see below) and for presentational purposes Figure 3.2 collapses across this factor. Inspection of Figure 3.2 indicates that test accuracy was greater for the faces that were congruent with the label (e.g., "artist") than those that were either incongruent (e.g., "criminal") or irrelevant ("banker"). A three-way mixed-factorial ANOVA was conducted on the data summarised in Figure 3.2 (factors: type of face - artist or criminal; label at encoding - "artist/criminal/banker"; label at test - "artist/criminal/banker"). This analysis revealed that there was the significant main effect of label at encoding, $F(2, 81)$ = 4.413, MSE = 0.587, p < .05, but no other main effects, largest $F(2, 81)$ = 0.164, p > .84. However, the label at encoding by face type interaction was significant, $F(2, 81)$ = 12.092, MSE = 0.284, p < .001. Analysing this further, when artists were targets there was an effect of label $F(2, 87)$ = 9.748 p < .05. This was explored using Tukey *post hoc* tests and revealed that when the artist label presented with the artist face (congruent), recognition accuracy was significantly higher than when it was incongruent, criminal (mean difference = 0.541, p < .05), or irrelevant, banker (mean difference = 0.717, p < .05). When criminals were targets the effect of label was significant, $F(2, 87)$ = 4.385, p < .05. This was explored as above and revealed that when the criminal label was presented with the congruent face (criminals), recognition accuracy was significantly higher than when it was incongruent, artist (mean difference = 0.392, p < .05) or

irrelevant (mean difference = 0.458, $p < .05$. The effect size of presenting a congruent label at encoding was r = .38 (Cohen's d = 0.83). There were no other significant interactions, largest $F(4, 81) = 1.013$, $p > .40$.
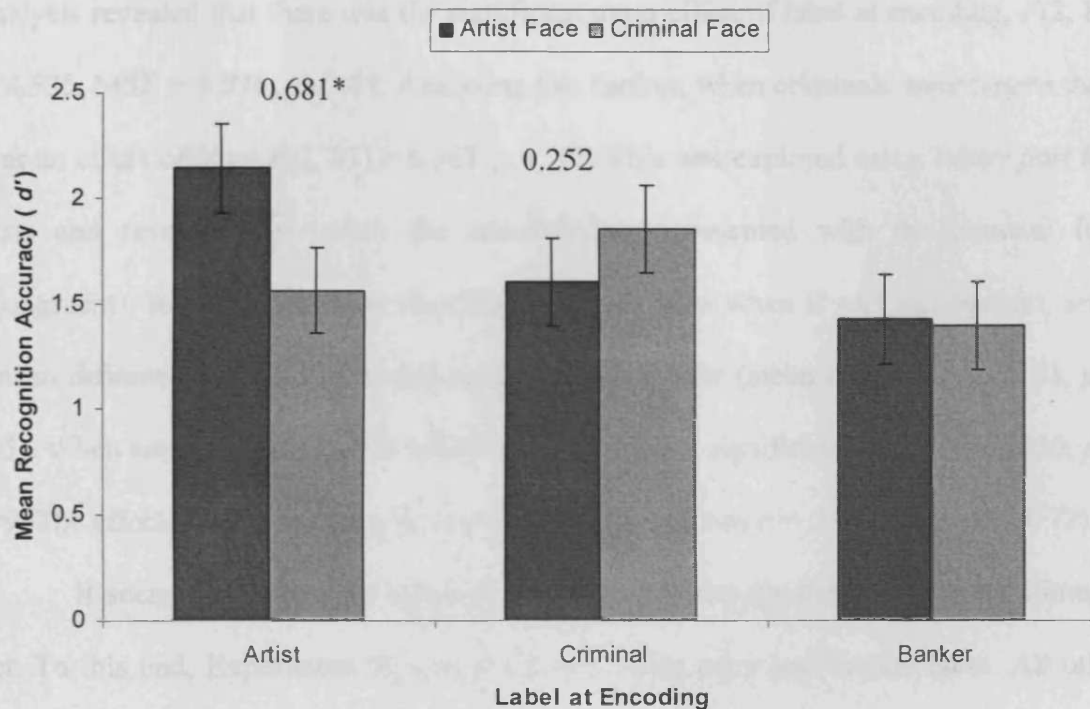
Figure 3.2. Mean recognition accuracy (d') for artist and criminal faces when the labels were artist, criminal, or banker (irrelevant) at encoding. Error bars represent standard error. Asterix represents significant differences between type of face for the particular label at encoding.

A parallel analysis was conducted on the response bias results summarised in Figure 3.3. Bias was assessed using C (Macmillan & Creelman, 2005). In the overall ANOVA there was no effect of label congruence during the test or any interaction involving this factor (see below) and for presentational purposes Figure 3.3 collapses across this factor. Inspection of Figure 3.3 indicates that response bias was more liberal for the faces that were congruent with the label (e.g., "artist") than those that were either incongruent (e.g.,

"criminal") or irrelevant ("banker"). A three-way mixed-factorial ANOVA was conducted on the data summarised in Figure 3.3 (factors: type of face - artist or criminal; label at encoding - "artist/criminal/banker"; label at test - "artist/criminal/banker"). This analysis revealed that there was the significant main effect of label at encoding, $F(2, 81)$ = 4.525, MSE = 0.076, $p < .05$. Analysing this further, when criminals were targets there was an effect of label $F(2, 87)$ = 6.563, $p < .05$. This was explored using Tukey *post hoc* tests and revealed that when the criminal label presented with the criminal face (congruent), response bias was significantly higher than when it was incongruent, artist (mean difference = 0.313, $p < .05$), or irrelevant, banker (mean difference = 0.181, $p < .05$). When artist were targets the effect of label was not significant, $F(2, 87)$ = 0.010, $p > .99$. The effect size for the congruent prime at encoding was $r = .34$ (Cohen's $d = 0.72$).

It seemed important to replicate this pattern of results using a different stimulus set. To this end, Experiment 5b was conducted using actor and banker faces. All other aspects of the method were identical to Experiment 5a except for the labels, where actor and banker were used as the two relevant labels and criminal was used as the irrelevant label.
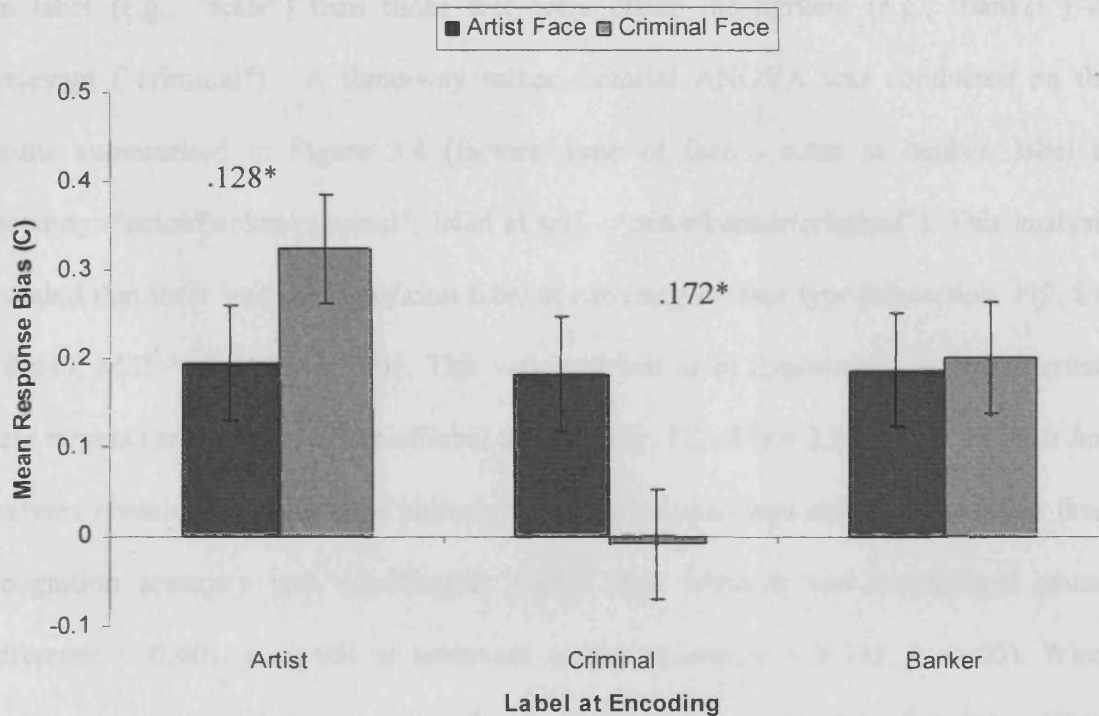
Figure 3.3. Mean response bias (*C*) for artist and criminal faces when the labels were artist, criminal, or banker (irrelevant) at encoding. Lower *C* represents a more liberal response bias (i.e., a tendency to respond with a hit or false alarm). Error bars represent standard error of the mean. Asterix represents significant differences between type of face for the particular label at encoding.

3.1.4. Experiment 5b

3.1.4.1. Results

The d' recognition accuracy scores are presented in Figure 3.4. This figure collapses across the congruence of the label at test because, as in Experiment 5a, the ANOVA revealed that this manipulation was again without influence (see below). Inspection of Figure 3.4 indicates that test accuracy was greater for the faces that were congruent with

the label (e.g., "actor") than those that were either incongruent (e.g., "banker") or irrelevant ("criminal"). A three-way mixed-factorial ANOVA was conducted on the results summarised in Figure 3.4 (factors: type of face - actor or banker; label at encoding: "actor/banker/criminal"; label at test - "actor/banker/criminal"). This analysis revealed that there was the significant label at encoding by face type interaction, $F(2, 81)$ = 8.040, MSE = 0.347, $p < .001$. This was explored as in Experiment 5a. When artists were targets there was an effect of label at encoding, $F(2, 87) = 3.531, p < .05$. Post hoc analyses revealed that when the stereotype label presented was congruent with the face, recognition accuracy was significantly higher than when it was incongruent (mean difference = 0.401, $p < .05$) or irrelevant (mean difference = 0.335, $p < .05$). When bankers were targets there was a significant effect of label at encoding, $F(2, 87) = 3.648$, $p < .05$. Post hoc analyses revealed that when the stereotype label presented was congruent with the face, recognition accuracy was significantly higher than when it was incongruent (mean difference = 0.317, $p < .05$) or irrelevant (mean difference = 0.317, $p < .05$). The effect size of presenting a congruent label at encoding was r = .28 (Cohen's d = 0.59). There were no other significant main effects or interactions, largest $F(4, 81)$ = 1.073, $p > .37$.

.511*
.350*
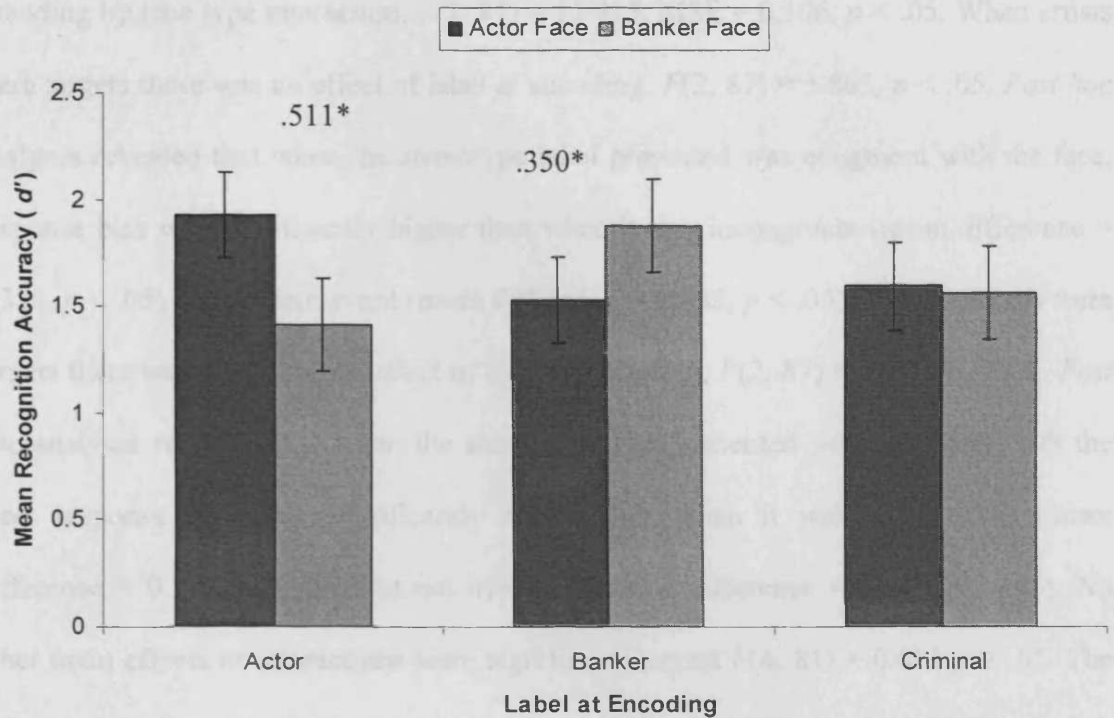
Figure 3.4. Mean recognition accuracy (*d'*) for actor and banker faces when the labels were actor, banker, or criminal (irrelevant) at encoding. Error bars represent standard error. Asterix represents significant differences between type of face for the particular label at encoding.

The *C* scores are presented in Figure 3.5. This figure collapses across the congruence of the label at test because, as in Experiment 5a, the ANOVA revealed that this manipulation was again without influence (see below). Inspection of Figure 3.5 indicates that response bias was more liberal for the faces that were congruent with the label (e.g., "actor") both than those that were incongruent (e.g., "banker") or irrelevant ("criminal") with respect to the face. A three-way mixed-factorial ANOVA was conducted on the results summarised in Figure 3.5 (factors: type of face - artist or criminal; label at encoding - "actor/banker/criminal"; label at test - "actor/banker/criminal"). This analysis revealed that there was the significant label at

encoding by face type interaction, $F(2, 81) = 15.915$, MSE $= 0.106$, $p < .05$. When artists were targets there was an effect of label at encoding, $F(2, 87) = 5.863$, $p < .05$. *Post hoc* analyses revealed that when the stereotype label presented was congruent with the face, response bias was significantly higher than when it was incongruent (mean difference $= 0.330$, $p < .05$) but not irrelevant (mean difference $= 0.182$, $p < .05$). When bankers were targets there was a significant effect of label at encoding, $F(2, 87) = 5.505$, $p < .05$. *Post hoc* analyses revealed that when the stereotype label presented was congruent with the face, response bias was significantly higher than when it was incongruent (mean difference $= 0.337$, $p < .05$) but not irrelevant (mean difference $= 0.137$, $p < .05$). No other main effects or interactions were significant, largest $F(4, 81) = 0.613$ $p > .65$. The effect size for the congruent prime at encoding was r $= .31$ (Cohen's $d = 0.66$).
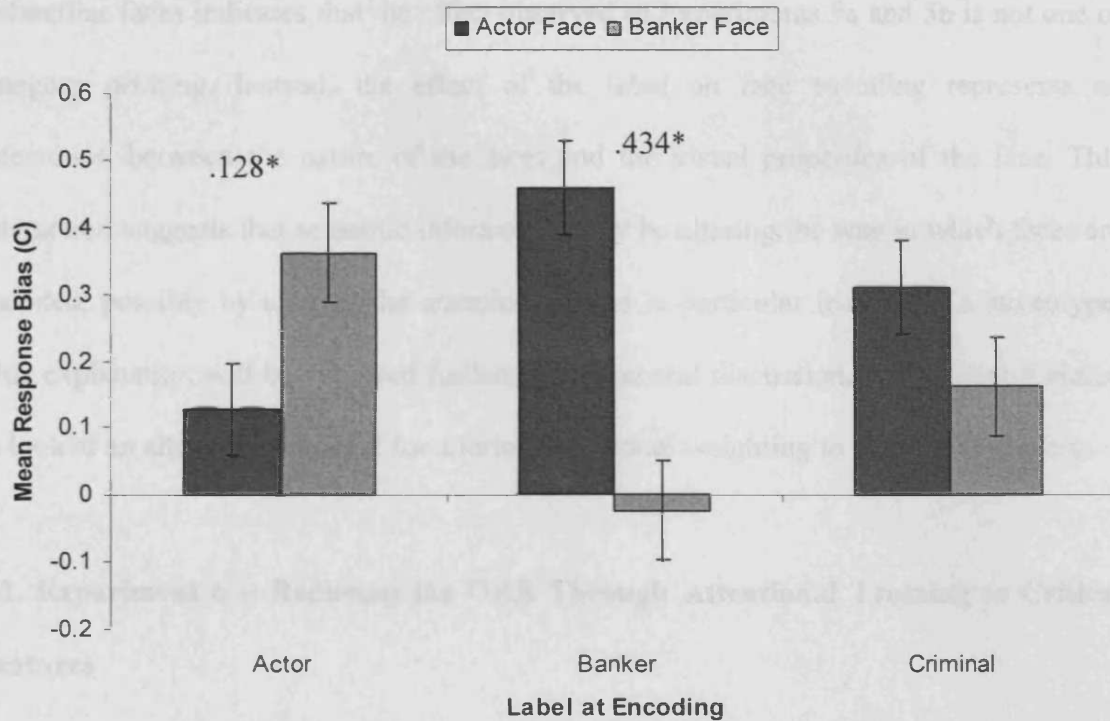
Figure 3.5. Mean response bias (*C*) for actor and banker faces when the labels were actor, banker, or criminal (irrelevant) at encoding. Error bars represent standard error. Asterix represents significant differences between type of face for the particular label at encoding.

3.1.5. Discussion

The results from this study clearly show a reliable stereotype priming effect, whereby faces that are learnt with a congruent stereotype label are better recognised subsequently. This effect, apparent in an analysis of *d'* scores, was solely a consequence of the presentation of the stereotype congruent label during encoding. This priming effect, however, is also associated with a more liberal response bias for faces congruent with the label, as evident in more false hits for stereotype congruent distractors (c.f., Klatzky et al., 1982a). The fact that the effect of presentation of a label at encoding is on previously

unfamiliar faces indicates that the effect observed in Experiments 5a and 5b is not one of category priming. Instead, the effect of the label on face encoding represents an interaction between the nature of the label and the visual properties of the face. This interaction suggests that semantic information may be altering the way in which faces are encoded, possibly by altering the attentional focus to particular features of a stereotype. This explanation will be explored further in the general discussion. Experiment 6 aimed to look at an alternative method for altering attentional weighting to particular features.

3.2. Experiment 6 – Reducing the ORB Through Attentional Training to Critical Features

Given that experience is not the sole mechanism that reduces the ORB (Lavrakas, et al., 1976), there must be some alternative reason that it occurs. Often people will process those of their own culture in an individuating manner and pay attention to the unique features (c.f., Levin, 2000). As such, they are attending to relevant features that aid individuation and thus discrimination. Experiment 6, therefore, aimed to train White participants to attend to the features that are most diagnostic for recognising Black faces. The dimensions hypothesised to be relevant come from evidence summarised in section 1.2.2. (see Ellis, et al., 1975; Shepherd & Deregowski, 1981) and include the mouth, cheeks, and nose. Using a perceptual learning paradigm, White participants were trained to recognise and distinguish faces based on these features. To control for mere exposure, a condition was included that involved seeing the same learning faces, but these were extraneous to the task. As such, guided attention can be compared to mere exposure and no exposure in the ORB.

3.2.1. Method

3.2.1.1. Participants

One-hundred-and-twenty-four Cardiff University undergraduates were randomly divided with the condition that equal numbers of participants were in each of four training groups: feature-critical, no training, non-critical feature, and colour blob discrimination. All participants had normal or corrected vision and were ethnically White. The data from two participants were removed since they had been raised in Africa, and were replaced by additional participants.

3.2.1.2. Apparatus & Stimuli

All stimuli were presented on a high-resolution colour monitor, displayed using DirectRT™ Research Software (Empirisoft™). Participant responses were recorded on a standard keyboard. Participants viewed the monitor at a distance of 60 cm.

Two face sets were used to assess the ORB before and after training. These were the MacBrain Face Stimulus Set[10], and a face set collected by Paul Hutchings at Cardiff University. Each face had two pictures taken of it, which were both neutral in expression and frontal views. One of these pictures was used in the presentation phase, while the other face was used in the recognition phase of the recognition paradigm. This was to avoid pictorial recognition. All stimuli had a resolution of 72 dpi and were 150 mm by 200 mm.

Three different training sets of faces were constructed using a face composite computer package FACES™ (Interquest™). This package uses 16 features to make a

---

[10] Development of the MacBrain Face Stimulus Set was overseen by Nim Tottenham and supported by the John D. and Catherine T. MacArthur Foundation Research Network on Early Experience and Brain Development. Please contact Nim Tottenham at tott0006@tc.umn.edu for more information concerning the stimulus set.

face. Within each feature a choice of between 10 and 100 features can be made. For each set, 40 different target faces were created. From these targets, 10 look-a-likes were created which were identical in certain features depending on condition (described below). These were all presented with 72 dpi resolution and dimensions 130 mm by 130 mm.

For the feature-critical faces, all the faces had similar hair style and colour. The look-a-likes only differed according to 'lower' facial features: the chin, nose, mouth, and cheeks. They were identical to the targets according to 'White' facial features. The coloured blob discrimination condition used these faces; however, each face had a coloured blob placed on it in a random place. The non-critical feature faces were created in the same way as the Feature-critical faces, except all the look-a-likes were identical in the 'lower' regions: chin, nose, mouth, and cheeks. Only the eyes and hair styles differed. An example of each of the faces is shown in Figure 3.6.

As an additional control, all participants filled out the Modern Racism Scale (McConahay, 1989). This scale is a reliable test measuring White participants' subtle racism (i.e., not explicit racist views).
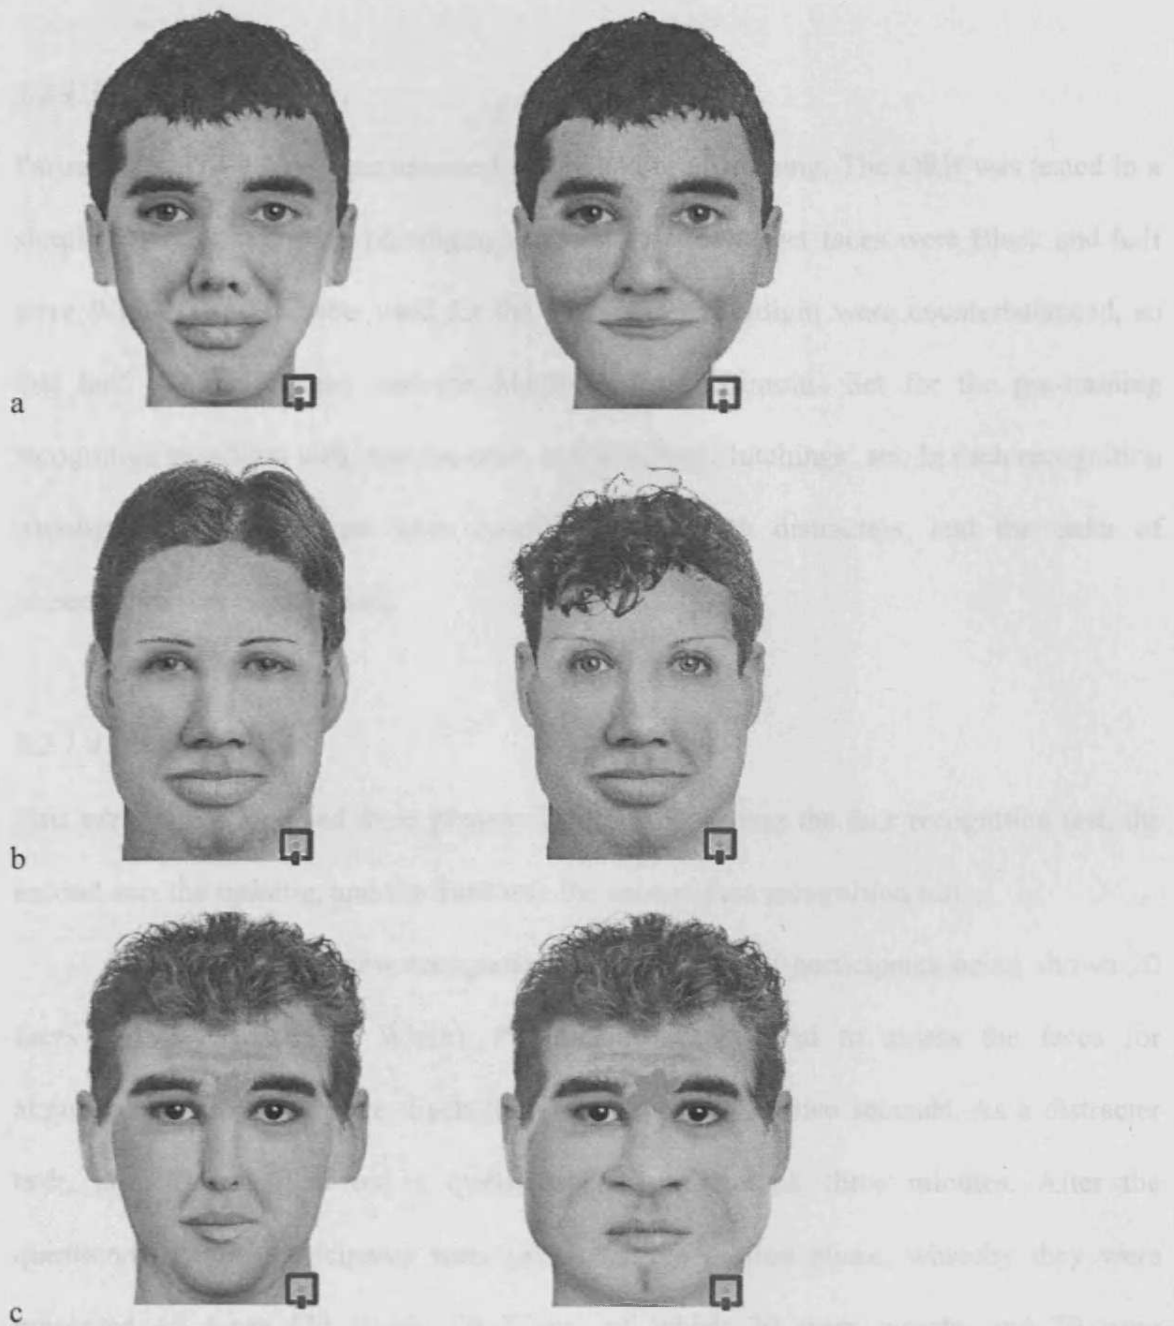
Figure 3.6. Examples of the face-stimuli used in the training session. The left face is the target face. The right face is a look-a-like. Feature-critical faces are shown in a. They differ only on nose, chin, mouth, and cheeks. Non-feature critical faces are shown in b. They differ only in hair style, eyes and eyebrows. Coloured blob faces are shown in c. They are the same as in the Feature-critical condition, except the coloured blobs differ in hue.

3.2.1.3. Design

Participants' ORB level was assessed before and after training. The ORB was tested in a simple yes/no recognition paradigm, whereby half the target faces were Black and half were White. The face sets used for the recognition paradigm were counterbalanced, so that half the participants had the MacBrain Face Stimulus Set for the pre-training recognition paradigm task, and the other half had Paul Hutchings' set. In each recognition paradigm task, the targets were counterbalanced with distracters, and the order of presentation was randomised.

3.2.1.4. Procedure

This experiment involved three phases. The first phase was the face recognition test, the second was the training, and the third was the second face recognition test.

The standard old/new recognition test consisted of participants being shown 20 faces (10 Black and 10 White). Participants were asked to assess the faces for attractiveness on a 0-9 scale. Each face was presented for two seconds. As a distracter task, participants filled out a questionnaire, which took three minutes. After the questionnaire, the participants were given the recognition phase, whereby they were presented 40 faces (20 Black, 20 White; of which 20 were targets, and 20 were distracters). They were asked for a yes or no response as to whether they had seen the face previously.

The training phase consisted of a perceptual learning paradigm. The participants were pre-exposed to the target face, and the target face was assigned a name. The participants were then presented with 20 further faces sequentially. For each face, participants had to state whether the face matched the target face (or colour in the colour

blob discrimination condition). The ratio of targets to distracters was 1 to 4. Ten target faces made up one training block. The training blocks got progressively harder, using more similar faces. There were a total of four training blocks, lasting a total of 60 minutes. Those participants in the no-training condition did not undergo this phase of the experiment. Instead, they were asked to leave the experimental booth and return after 60 minutes.

The final recognition test was identical to the first phase of the experiment, but using a different questionnaire and a different face set.

3.2.2. Results

All scores on the Modern Racism Scale were between 0 and 4 implying that no participants held racist views, and therefore this was not considered any further. Recognition accuracy (d') was converted into an ORB score for each participant using Formula 2.1. Figure 3.7 shows the mean ORB scores before and after the training regimes. The ORB is present for all participants, except those having completed critical-feature training.

These data were entered into a mixed factorial ANOVA, with type of training as a four-level between-subjects variable and before and after being a two-level within-subjects variable. The interaction between level of training and condition was significant, F (3, 124) = 3.303, MSE = 0.201, p < .05. This was indicated by the ORB being significantly reduced after training for those receiving feature-critical training (mean difference = 0.36, p < .05). After feature-critical training, participants' ORB was not significantly different from zero, suggesting equivalent performance on Black and White faces, as tested using an one sample t-test, t (31) = 0.759, p > .45. No other comparisons

were significant. The main effect of before and after training was not significant, F(1,

124) = 1.011, MSE = 0.201, $p > .31$. The effect size for the removal of the ORB was

large, r = 0.9 (Cohen's d = 4.2). Incidentally, using a simpler method for calculating the

ORB (d'_w - d'_B) the results are identical.
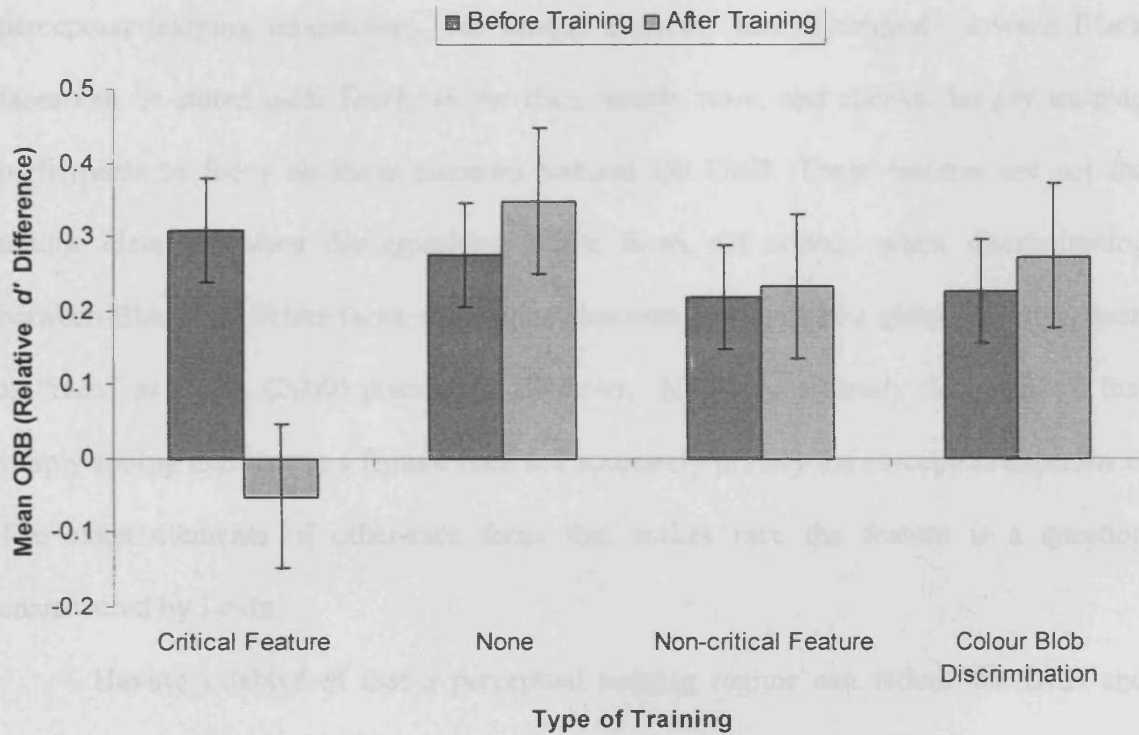
Figure 3.7. ORB scores before and after training. Error bars represent standard error.

3.2.3. Discussion

By altering the focus of attention of White participants, the ORB has been reduced in this

experiment. Participants were trained, in a perceptual learning paradigm, to focus on

lower facial features (i.e., used by Black participants to describe Black faces, Shepherd

and Deregowski, 1981). Since the ORB was not reduced when the training faces were

presented without deliberate attention paid to them, the results indicate that this

procedure directed the participants' attention toward the lower facial features. It was not

possible, however, to increase the ORB using training on White face diagnostic

123

information. This may be due to the ORB already being an established mechanism within face processing.

The present data are compatible with perceptual learning mechanisms suggested by such researchers as McLaren and Mackintosh (2000, 2002) and Hall (1980). In perceptual learning terminology, the unique elements that distinguish between Black faces can be stated quite firmly as the chin, mouth, nose, and cheeks. Simply training participants to focus on these elements reduces the ORB. These features are not the unique elements when distinguishing White faces. Of course, when discriminating between Black and White faces, the unique elements may well be a global construct such as "race" as Levin (2000) postulates. However, this study certainly demonstrates that simply stating that race is a feature does not accurately portray the perceptual experience. The exact elements of other-race faces that makes race the feature is a question unanswered by Levin.

Having established that a perceptual training regime can reduce the ORB and theorising that this is due to attentional mechanism, it becomes important to directly test attentional mechanisms in face recognition. Experiment 7 aims to do this, using a fixation cross procedure borrowed from the attention literature.

3.3. Experiment 7 – Reducing the ORB with Fixation Crosses Preceding Diagnostic Features

Shepherd and Deregowski (1981) have identified the features in the lower portion of the face as more critical for Black faces than features in the upper part of the face, which are more diagnostic for White faces. The training in Experiment 6 involved a perceptual learning paradigm whereby participants had to distinguish faces based on features of the

lower portion of the face and crucially mere exposure was not sufficient to reduce the ORB. Indeed, attention had to be paid to the critical features of the faces to cause the reduction. This attentional explanation of the ORB was tested more directly in Experiments 7, 8, and 9. Using a paradigm borrowed from research on attention (c.f., Ogawa, Takeda, and Kumada, 2007), participants were shown a series of Black and White faces. These real faces were preceded by a fixation cross that appeared in the upper portion of the face or the lower portion of the face. These fixation crosses were designed to attract attention to either portion of the face. If attention is the underlying cause of the ORB, then attention paid to the lower portion of Black faces should cause a significant reduction in the ORB.

3.3.1. Method

3.3.1.1. Participants

An opportunity sample of 24 participants from the Psychology undergraduate population at Cardiff University took part in this study for partial fulfilment of a course requirement. All had normal vision and were White women.

3.3.1.2. Apparatus & Stimuli

All stimuli were presented on a high-resolution colour monitor, displayed using SuperlabPro2™ Research software. Participant responses were recorded on a standard keyboard. Participants viewed the monitor at a distance of 60 cm. A total of 128 faces were to be tested from the database used in Experiment 6. All the images were presented in 72 dpi resolution, 125 mm wide by 188 mm high. Participants also filled out the Modern Racism Scale (McConahay, 1989).

3.3.1.3. Design

The fixation cross could precede the upper portion of the face or the lower portion of face at learning and could precede the upper portion of the face or lower portion of the face at test in a 2 by 2 within-subjects experimental procedure. Faces were counterbalanced so that each face appeared as a target and a distracter an equal number of times. The order of presentation of the faces was randomised. Position of the fixation cross was randomised so that each face had an equal chance of the fixation cross preceding the upper portion or the lower portion at learning and at test. Half of the face stimuli shown were of Black faces and the other half were of White faces. Recognition accuracy for Black and White faces was measured using d'.

3.3.1.4. Procedure

A standard old/new recognition paradigm similar to that employed in Experiment 1 was used here. Sixty-four faces were presented on screen for 2000 ms, each preceded by a fixation cross lasting 200 ms (either in the upper or lower half of the face). Participants were instructed to rate each face for distinctiveness, by giving it a score out of 9 referring to how easy the face would be to spot in a crowd (c.f., Light, et al., 1979). Faces were presented sequentially. After the learning phase, participants were given the Modern Racism Scale questionnaire as a distractor task. Once they had completed the questionnaire, the participants were given the test phase. In this, the participants were shown all 128 faces sequentially. Prior to each face, a fixation cross appeared for 200 ms (either in the upper or lower half of the face). The faces were on screen until the response was made. Finally, the participants were thanked and debriefed. A schematic representation of this procedure is shown in Figure 2.8.
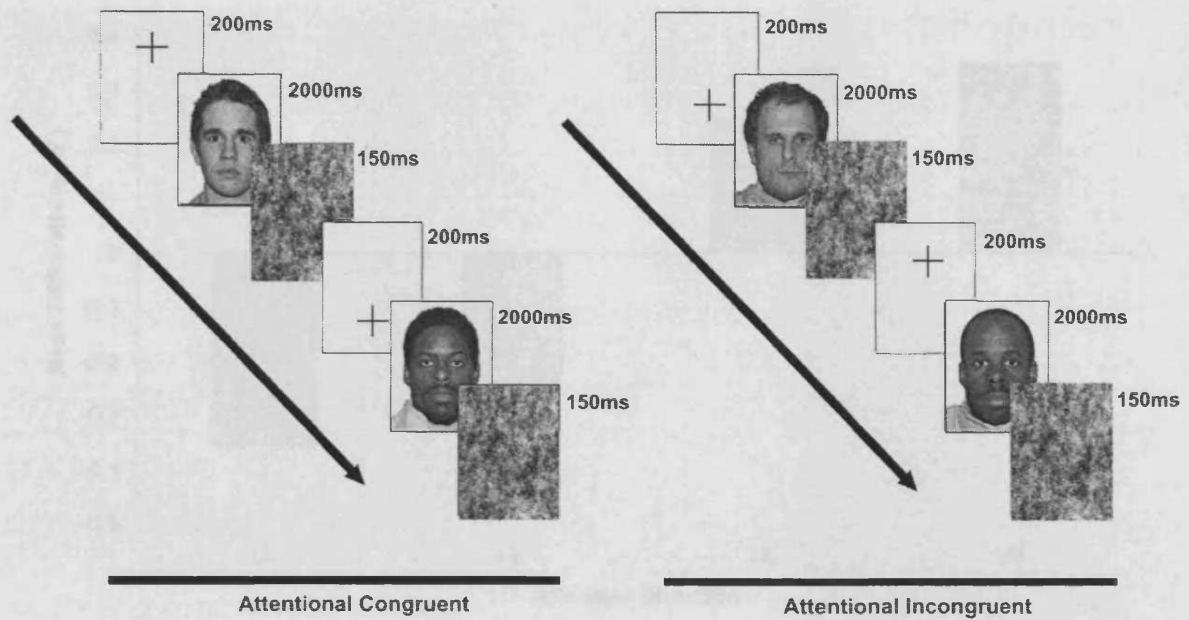
<u>Figure 3.8.</u> Procedure for Experiment 7. The left panel shows an example of Black and White faces being preceded by fixation crosses that are congruent with the most diagnostic features. The right panel shows an example of Black and White faces being preceded by fixation crosses that are incongruent with the most diagnostic features.

3.3.2. Results

All scores on the Modern Racism Scale were between 0 and 6 implying that no participants held racist views, and therefore this was not considered any further. ORB scores were calculated as in Experiment 6 and are presented in Figure 3.9.

The ORB score data were subjected to a 2 x 2 within-subjects ANOVA. The ANOVA revealed the interaction between attention directed to upper and lower portions of the face at learning and test was not significant, $F(1, 23) = 0.027$, MSE = 0.080, $p > .87$. There was main effect of attention at learning, $F(1, 23) = 93.909$, MSE = 0.090, $p < .05$, but not at test, $F(1, 23) = 3.460$, MSE = 0.087, $p > .05$.

Figure 3.9. Mean ORB when attention is directed to either the upper (U) or lower (L) portion of the face at learning or test (first and second digit respectively). Positive numbers indicate an ORB; negative numbers indicate an other-race bias. Error bars show standard error.

Figure 3.9 shows the transformed data from raw hit and false alarm rates into d' and subsequently into an ORB score. Due to the multiple transformations required, some information is lost: specifically whether the effect is one of improving recognition of Black faces or a detriment in the recognition of White faces. As such, a more detailed presentation of the results is presented in table 3.1. This presentation explores the effect in more detail but due to its complexity is not as simple to understand as the findings summarised in Figure 3.9.

Table 3.1.

Hit, false alarm, and d' data for both Black and White faces and the ORB scores used in the main analysis. Standard errors of the mean are shown in parentheses.

	White Faces				Black Faces				ORB
	Hit rate	False Alarm rate	d'		Hit rate	False Alarm rate	d'		d' difference
UU	0.847 (.104)	0.094 (.092)	2.462 (.821)		0.675 (.218)	0.329 (.180)	1.215 (.772)		0.348 (.411)
UL	0.844 (.084)	0.135 (.079)	2.307 (.619)		0.709 (.205)	0.281 (.117)	1.422 (.586)		0.246 (.239)
LU	0.732 (.176)	0.246 (.092)	1.457 (.692)		0.859 (.083)	0.177 (.086)	2.154 (.458)		-0.237 (.282)
LL	0.636 (.148)	0.313 (.128)	1.157 (.0691)		0.823 (.086)	0.135 (.084)	2.208 (.634)		-0.358 (.287)

3.3.3. Discussion

By altering White participants' focus of attention, their ORB was reduced in this experiment. Indeed, an interesting pattern of results has been discovered. The ORB is less if the face is learnt with attention directed toward the lower portion of the face. In fact, if faces are learnt with attention directed to the lower portion of the face, a significant other-race bias appears. This indicates that attentional shift is sufficient to explain, remove, and reverse the ORB. Furthermore, this result is indicative that the ORB is produced as a result of the way the faces are observed when being encoded rather than when being retrieved.

Examination of the more detailed presentation of the data in Table 3.1 reveals that the effect of the fixation cross preceding the bottom of the face is in improving the recognition of Black faces and causing a detriment in recognition of White faces. Indeed, it is a reversal of the classic ORB. The magnitude of the mean accuracy for White faces when the fixation cross preceded the top of the face is similar to that of the Black faces when the fixation cross preceded the bottom of the face (compare d' of 2.462 to 2.208 respectively).

During post-test briefing, it became clear that the fixation crosses were sufficient to grab attention. When the fixation crosses were presented in the lower portion of the face, participants reported that their eyes drifted upward whilst the face was on screen. Of course, a face is likely to grab attention and eye movements will result and so only the initial position of the gaze will be different in the conditions. This suggests that participants' initial glance at a face is most important for the way they encode it. Experiment 8 aimed to explore this in more detail.

3.4. Experiment 8 – Reducing the ORB with fixation crosses when instructed to be fast or accurate

Experiment 7 demonstrated that participants' ORB can be reduced by presenting participants with fixation crosses in the lower portion of gaze prior to Black faces. Due to the fact that participants reported that their eyes drifted toward the top of the face whilst it was on screen irrespective of the location of the fixation cross, a second experiment was conducted in which this was either encouraged or discouraged by experimental instructions. When participants are told to be accurate, their recognition accuracy is higher than without this instruction. However, their time to response is significantly longer (c.f., Light, et al., 1979; Shepherd, Gibling, and Ellis, 1991). This speed-accuracy trade-off is reversed when the instructions are to be as fast as possible. It was surmised that the instruction to be quick would force participants to make distinctiveness judgements and recognition responses before their eyes had chance to drift up to the upper portion of the faces. With accuracy instructions, however, participants would be more willing to wait until their eyes settled in the upper portion of the face. Thus, it was

hypothesised that the fixation cross in the lower portion of the face will only remove the ORB for those participants instructed to be as fast as possible.

3.4.1. Method

3.4.1.1. Participants

An opportunity sample of 24 participants from the Psychology undergraduate population at Cardiff University took part in this study for partial fulfilment of a course requirement. All had normal vision and were brought up in the UK. Four were male.

3.4.1.2. Apparatus & Stimuli

To assess the ORB, 64 faces (32 Black and 32 White) were taken from the Minear and Park (2004) face set. All faces were typical and contained no extraneous features such as clothing or make-up. The faces chosen ranged in age between 18 and 28. All other stimuli were identical to Experiment 7.

3.4.1.3. Design & Procedure

The same experimental procedure used in Experiment 7 was applied here, except that half the participants were instructed to be as fast as possible whereas the other half of participants were instructed to be as accurate as possible just prior to the presentation and the test phases. As such, there was an additional between-subjects variable in this Experiment. This created a 2 (fixation cross at learning) x 2 (fixation cross at test) x 2 (instructions) mixed design.

3.4.2. Results

As in Experiment 7, d' was recorded and the ORB was calculated using Equation 2.1. No participant scored above 4 on the Modern Racism Scale so this was not considered further. The mean ORB scores are presented in Figure 3.10. A Mixed ANOVA revealed that there was a main effect of fixation cross location at learning, $F (1, 22) = 25.593$, $MSE = .285$, $p < .05$, whereby the ORB was lower when the fixation cross was in the bottom half of the face. None of the interactions reached significance, the main effects of the fixation cross at test nor the main effect of instruction type, largest $F = 3.371$, $p > .08$.

As in Experiment 7, a more detailed presentation of the data can be found in Table 3.2. This revealed a similar pattern of findings to those found in Experiment 7, whereby the effect of the fixation cross was to lower the recognition of White faces and improve the recognition of Black faces. The data also indicate that the mean accuracy scores in Experiment 8 were lower than those found in Experiment 7. This will be discussed subsequently.
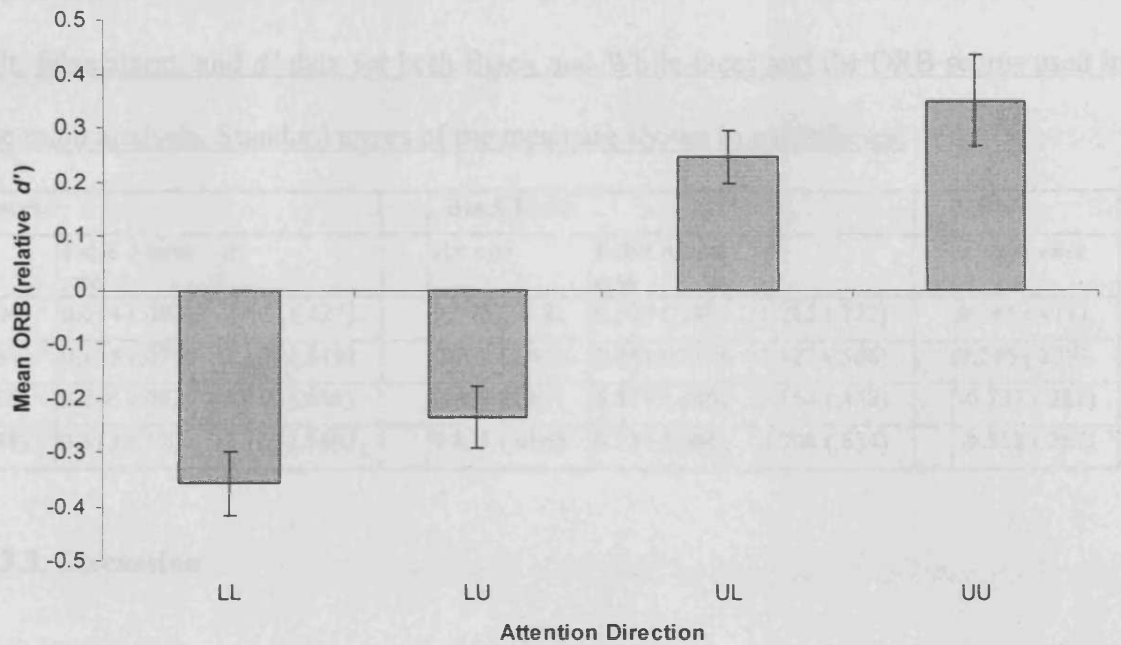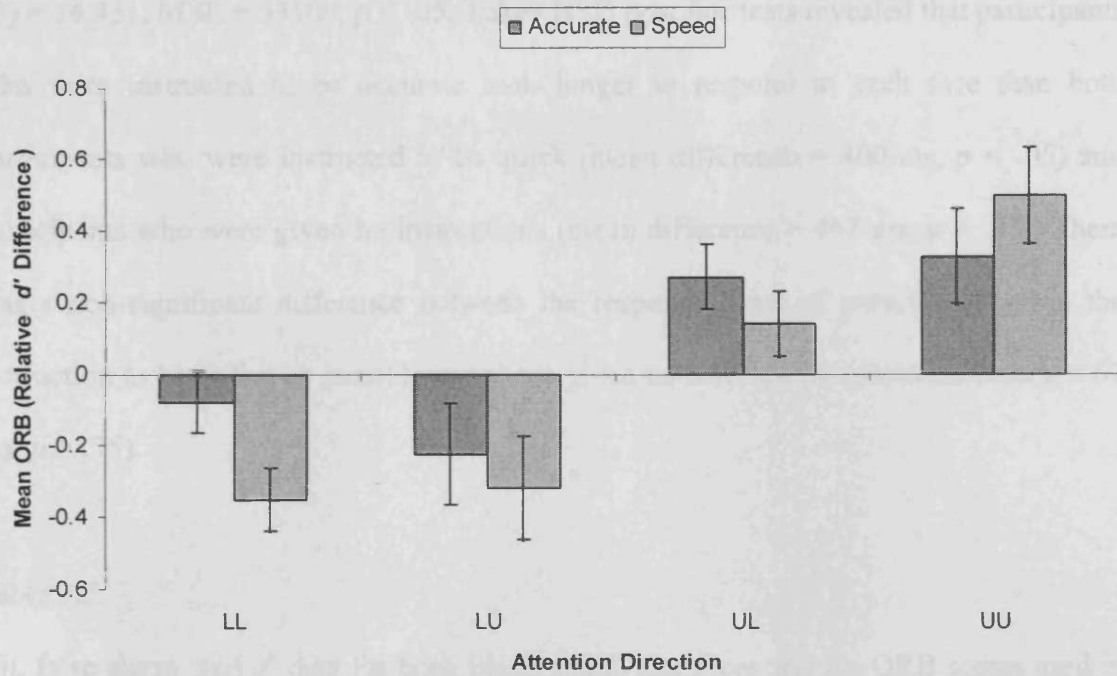
Figure 3.10. Mean ORB when attention is directed to either the upper (U) or lower (L) portion of the face at learn or test (first and second digit respectively), split by experimental instructions. Positive numbers indicate an own-race bias; negative numbers indicate an other-race bias. Error bars show standard error.

Due to the failure to find an effect of instructions, an analysis on the reaction-time data was conducted as a manipulation check. The reaction times to make a response to the target faces from Experiments 7 and 8 were combined[11]. The data are summarised in Figure 3.11, demonstrating that when participants were instructed to be accurate, they did indeed take longer to give a distinctiveness rating than when they were instructed to be quick. However, participants instructed to be quick were no quicker than participants given no instructions. The data summarised in Figure 3.11 were subjected to a 3-level one-way ANOVA. This revealed a significant effect of instruction on reaction time, $F(2,$

---

[11] It must be noted that combining the data from Experiments 7 and 8 involves comparing responses across different face stimuli and may not be appropriate. Nevertheless, it is shown since it highlights an important finding.

33) = 14.431, MSE = 53100, $p < .05$. Tukey HSD *post hoc* tests revealed that participants who were instructed to be accurate took longer to respond to each face than both participants who were instructed to be quick (mean difference = 400 ms, $p < .05$) and participants who were given no instructions (mean difference = 467 ms, $p < .05$). There was a non-significant difference between the response times of participants given the instruction to be as fast as possible and those given no instructions (mean difference = 67 ms, $p > .75$).

Table 3.2.

Hit, false alarm, and d' data for both Black and White faces and the ORB scores used in the main analysis. Standard errors of the mean are shown in parentheses.

		White Faces			Black Faces			ORB
		Hit rate	False Alarm rate	d'	Hit rate	False Alarm rate	d'	d' *difference*
Speed	UU	0.772 (.079)	0.258 (.094)	1.436 (.364)	0.567 (.223)	0.319 (.162)	0.673 (.736)	0.496 (.593)
	UL	0.729 (.231)	0.208 (.144)	1.567 (.682)	0.597 (.170)	0.217 (.094)	1.08 (.415)	0.140 (.349)
	LU	0.550 (.243)	0.190 (.089)	1.051 (.816)	0.771 (.139)	0.250 (.131)	1.521 (.472)	-0.319 (.327)
	LL	0.589 (.142)	0.215 (.121)	1.077 (.650)	0.813 (.065)	0.167 (.062)	1.904 (.342)	-0.350 (.566)
Accurate	UU	0.770 (.079)	0.178 (.016)	1.687 (.260)	0.667 (.123)	0.325 (.136)	0.960 (.571)	0.326 (.267)
	UL	0.750 (.131)	0.271 (.175)	1.408 (.547)	0.700 (.191)	0.417 (.123)	0.834 (.408)	0.270 (.263)
	LU	0.604 (.219)	0.233 (.130)	1.088 (.792)	0.729 (.175)	0.250 (.131)	1.408 (.547)	-0.226 (.274)
	LL	0.677 (.135)	0.292 (.075)	1.054 (.492)	0.688 (.146)	0.271 (.175)	1.217 (.342)	-0.079 (.405)
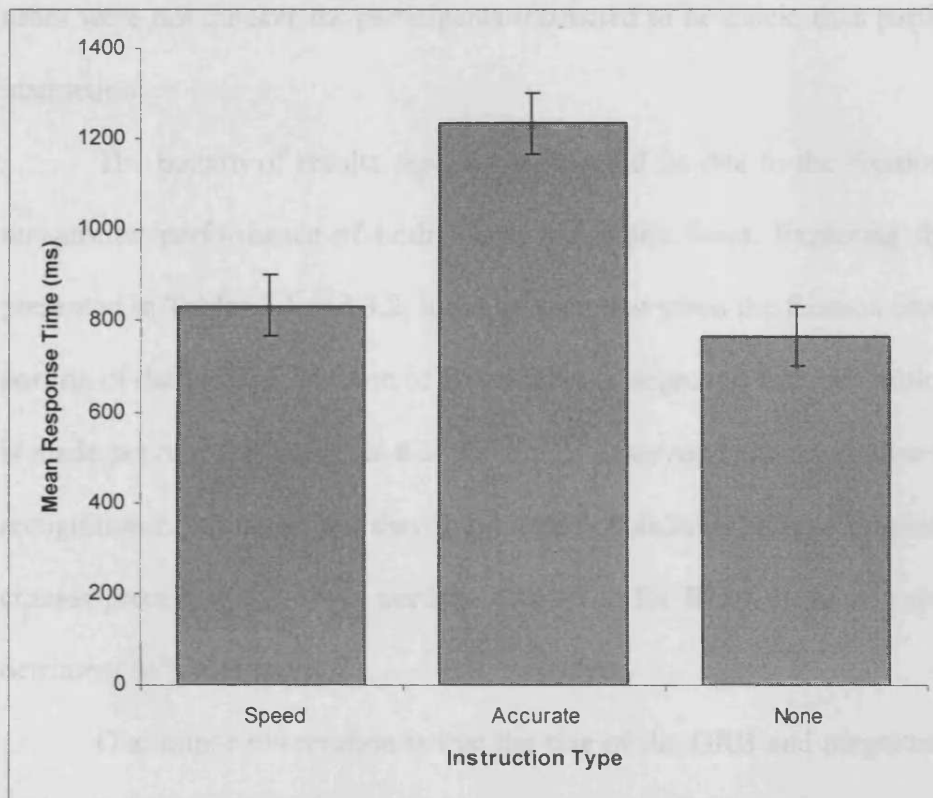
Figure 3.11. Mean response time (ms) for participants given the instructions: to be as fast as possible; to be as accurate as possible; and no instructions. Error bars show standard error.

3.4.3. Discussion

Experiment 8 partially replicated the findings of Experiment 7. The ORB was significantly reversed when the fixation cross that precedes a face was in the lower half during initial learning of the face. This effect occurs irrespective of the type of instructions given at the presentation and the recognition phase. This finding is inconsistent with the hypothesis that decisions made faster would be more susceptible to the effects of a fixation cross, however, the reaction time analysis revealed that reaction

135

times were not quicker for participants instructed to be quick, than participants given no instructions.

The pattern of results reported here could be due to the fixation cross lowering recognition performance of both Black and White faces. Exploring the extended data presented in Tables 3.1 and 3.2, it can be seen that when the fixation cross is in the lower portion of the face, recognition of Black faces is improved and recognition of White faces is made poorer. This suggests that the effects observed here are not due to a detriment in recognition of all faces, but they are due to beneficial effects of presentation of fixation crosses preceding the lower portion of the face for Black faces and at the same time a detriment to White faces.

One minor observation is that the size of the ORB and magnitude of d' obtained in Experiment 8 is considerably lower than that in Experiment 7. An explanation offered for this is that Experiment 8 used the Minear and Park (2004) face set whereas Experiment 7 used the NimStim face database. It is conceivable that the NimStim face database contains White faces that are more distinctive than the Minear and Park set. This would account for the apparent discrepancy. Nevertheless, both sets of stimuli produce the ORB and its reduction with fixation cross preceding the face in the lower portion.

A larger issue is why participants instructed to be quick did not show an exaggerated effect of the fixation cross. The instructions were clear to the participants but they did not respond faster to these faces than participants given no instructions in Experiment 7. This suggests that participants in Experiment 7 were already responding at the optimal speed (although stimuli differences cannot be discounted). To address this issue, an alternative method can be applied to test whether the drifting of attention removes the effect. Rather than instructing participants to be more accurate (and thus spend more time examining the face), the face can remain on screen for a set amount of

time. This would thus prevent the participants from allowing their eyes to drift, if the face was presented for a brief enough duration.

3.5. Experiment 9 – Reducing the ORB with Fixation Crosses with Limited or Extended Face Processing Time

The purpose of Experiment 9 was to investigate the effect that varying exposure time had on the effect of fixation crosses on the ORB. In order to do this, the faces were either on screen for 4000 ms (to allow for fixation drift), 2000 ms (to allow for some fixation drift), or 1000 ms (to prevent fixation drift). Only after this delay were participants permitted to make a response. Thus, it is expected that the fixation cross will have a larger effect in reversing the ORB in the 1000 ms condition than either of the other two conditions.

3.5.1. Method

3.5.1.1. Participants

Thirty-eight psychology undergraduates from Cardiff University participated in this experiment as partial fulfilment of a course requirement. All were White British nationals of which eight were male. All self-reported they had normal or corrected vision. Participants were randomly divided into one of six experimental groups with the condition that approximately equal numbers of participants in each condition (N = 6/7).

3.5.1.2. Apparatus & Stimuli

The same apparatus and materials used in Experiment 8 were used in Experiment 9, except that participants saw 128 faces (64 Black and 64 White) from the Minear and Park (2004) face set.

3.5.1.3. Design & Procedure

Experiment 9 employed a 3 by 2 by 2 mixed factorial design, whereby participants were randomly divided into whether they would have faces on screen for 1000 ms, 2000 ms, or 4000 ms. As in Experiment 7, the fixation crosses could appear in the upper or lower half of the faces at learning or at test. A control condition was implemented, whereby three additional groups of participants were not given the fixation cross, but were tested with faces on screen for either 1000 ms, 2000 ms, or 4000 ms. The faces were counterbalanced such that each appeared as a target the same number of times, and the order of presentation was randomised.

The procedure for Experiment 8 was similar to Experiments 10 and 11. In Experiment 9, the faces were on screen for a set period of time (1000 ms, 2000 ms, or 4000 ms), after which the participants were prompted to make their response. This occurred at the learning and at the test phases. All other aspects of the procedure were identical to Experiment 7.

3.5.2. Results

Figure 3.12 shows the mean ORB for those participants who had the fixation crosses and those participants who did not have the fixation crosses. The data for no fixation cross are presented in Figure 3.12 and were compared to the data from the upper fixation cross at

learning and test in a 2 by 3 mixed ANOVA with the factors: fixation cross (no versus preceding the upper portion of the face) and time (1000 ms, 2000 ms, and 4000 ms). This showed no significant effect of fixation cross, F (1, 35) = 0.015, MSE = 0.116, $p > .90$, nor a significant interaction, F (2, 35) = 0.360, MSE = 0.116, $p > .70$. For the purposes of the analysis, the no fixation cross condition was not entered into the main 3 by 2 by 2 mixed ANOVA. This ANOVA revealed that there was a main effect of presentation fixation cross, whereby the ORB was greater when the fixation cross preceded the upper portion of the face than the lower portion of the face (mean difference = 0.355), $F(1, 35)$ = 33.419, MSE = .143, $p < .05$. This main effect was qualified by a significant interaction with the length of presentation, $F(2, 35)$ = 23.987, MSE = .143, $p < .05$, whereby the ORB was not reduced by lower fixation crosses for participants who saw the faces for 2000 ms (mean difference = 0.320, $p > .20$) and 4000 ms (mean difference = 0.142, $p >$ .36), but it was reduced for 1000 ms (mean difference = 0.887, $p < .05$). There was no effect of fixation cross at test, largest $F(2, 35)$ = 2.124, $p > .13$. There was a main effect of length of time the faces were on screen for, $F(2, 35)$ = 5.013, MSE = 0.267, p < .05, whereby faces on screen for 4000 ms had a significantly larger ORB than those on screen for 1000 ms (mean difference = 0.322, $p < .05$) but not larger than those on screen for 2000 ms (mean difference = 0.144, $p > .16$). Faces on screen for 2000 ms had a non-significantly larger ORB than those on screen for 1000 ms (mean difference = 0.178, $p >$ .10).
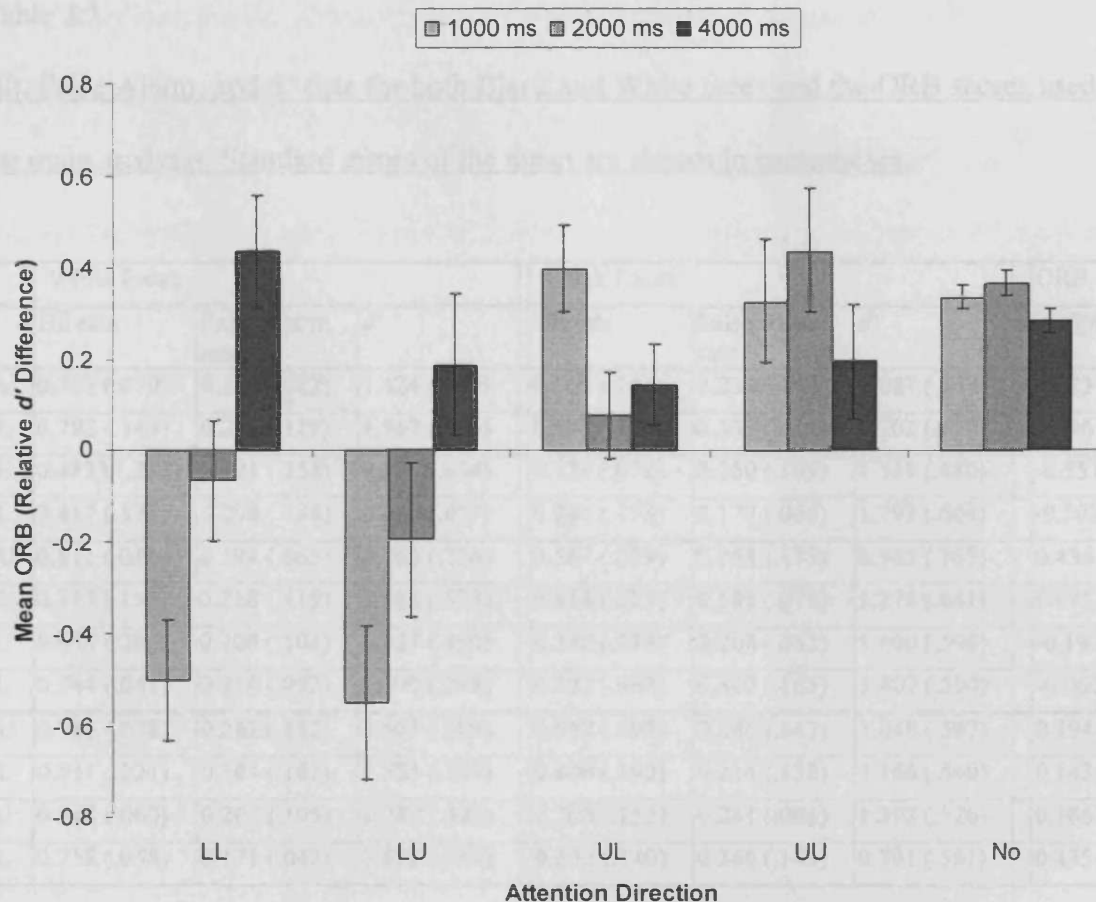
Figure 3.12. Mean ORB when attention is directed to either the upper (U) or lower (L) portion of the face at learn or test (first and second digit respectively), split by length of time the face was on screen for. Positive numbers indicate an own-race bias; negative numbers indicate an other-race bias. Error bars show standard error.

As in Experiment 7, a more detailed presentation of the results is presented in Table 3.3. This presentation explores the complexity of these findings insofar as to determine whether the effect is an improvement in the recognition of Black faces or a detriment in the recognition of White faces. The results replicate those observed in Experiments 10 and 11.

Table 3.3.

Hit, False Alarm, and d' data for both Black and White faces and the ORB scores used in the main analysis. Standard errors of the mean are shown in parentheses.

		White Faces			Black Faces			ORB
		Hit rate	False Alarm rate	d'	Hit rate	False Alarm rate	d'	d' difference
1000 ms	UU	0.733 (.070)	0.222 (.082)	1.424 (.189)	0.608 (.164)	0.233 (.140)	1.087 (.849)	0.323 (.592)
	UL	0.792 (.144)	0.271 (.119)	1.567 (.466)	0.592 (.149)	0.333 (.111)	0.702 (.470)	0.396 (.311)
	LU	0.4833 (.235)	0.321 (.158)	0.503 (.644)	0.729 (.072)	0.250 (.169)	1.339 (.480)	-0.551 (.691)
	LL	0.417 (.174)	0.208 (.138)	0.652 (.659)	0.781 (.178)	0.177 (.064)	1.792 (.603)	-0.502 (.687)
2000 ms	UU	0.811 (.0164)	0.194 (.065)	1.762 (.236)	0.567 (.229)	0.258 (.173)	0.905 (.767)	0.436 (.426)
	UL	0.713 (.115)	0.218 (.119)	1.424 (.513)	0.614 (.223)	0.186 (.071)	1.278 (.641)	0.075 (.365)
	LU	0.658 (.203)	0.200 (.108)	1.327 (.850)	0.750 (.238)	0.208 (.062)	1.600 (.598)	-0.196 (.687)
	LL	0.744 (.041)	0.214 (.092)	1.490 (.248)	0.792 (.062)	0.302 (.155)	1.402 (.590)	-0.067 (.238)
4000 ms	UU	0.786 (.078)	0.262 (.182)	1.505 (.569)	0.667 (.092)	0.286 (.147)	1.046 (.507)	0.194 (.367)
	UL	0.811 (.121)	0.304 (.182)	1.525 (.589)	0.600 (.190)	0.214 (.138)	1.166 (.540)	0.143 (.312)
	LU	0.814 (.060)	0.207 (.105)	1.787 (.320)	0.705 (.152)	0.241 (.091)	1.292 (.526)	0.186 (.300)
	LL	0.758 (.058)	0.171 (.042)	1.672 (.263)	0.652 (.140)	0.366 (.143)	0.791 (.561)	0.435 (.338)

3.5.3. Discussion

Experiment 9 replicates the findings of Experiments 7 and 8, and suggests that the difference in size of the ORB in Experiments 7 and 8 is due to the face database being used, since the magnitude of the ORB in Experiment 9 is similar to that in Experiment 8. Experiment 9 also found evidence that the attentional element the ORB is transient, and when participants were forced to delay their responses, the fixation cross did not have any effect. Thus, Experiment 9 deals with two major problems with Experiment 8, in that it indicates that the fixation cross guides immediate attention, but White participants' eye gaze drifts up when looking at faces.

As in Experiment 7, the data presented in Table 3.3 reveal that recognition of Black faces is poor when the fixation cross precedes the upper portion of the face, or

141

when the faces are on screen for a prolonged duration. Recognition of Black faces is improved if the face is on screen for a short amount of time and the fixation cross precedes the lower portion of the face. Indeed, the recognition of White faces is made worse when the presentation time is shorter and the fixation cross precedes the lower portion of the face. As such, this effect is one of reversing the ORB, not solely causing a detriment to the recognition of White faces worse than Black faces.

There are several important issues for the theoretical understanding of the ORB from these experiments. The results can be consistent with Levin's (2000) feature-selection hypothesis of the ORB if one suggests that the fixation cross at the chin prevents the "race feature" being processed. Levin suggested that this "race as a feature" was attended to first in other-race faces. If this explanation generalises to the reversal found when the fixation cross precedes the chin, then own-race faces can be processed according to race if attention is paid to the chin. Another facet of these results is that the effect of fixation crosses dies away quickly. This means the initial encoding that maybe good in the processing of other-race faces is over-ridden by that which is usually used. Experiment 10 explored another effect the fixation cross on the encoding of another class of faces, namely the inversion effect.

3.6. Experiment 10 – Reducing the Inversion Effect with Fixation Crosses Preceding the Forehead Part 1 – The Fixation Cross in Different Positions

Experiment 9 established that the first glance at a face is important for its subsequent recognition. First glances are usually to the top of the face (e.g., Althoff and N. Cohen, 1999). It is conceivable that the inversion effect in face recognition is due to the focus of primary attention being in the less diagnostic facial features. Experiment 10 tested this

using fixation crosses preceding upright and inverted faces either in the upper portion or the lower portion of the stimulus at learning and at test. Thus, a 2 by 2 by 2 design was employed, whereby the inversion effect was measured by recognition accuracy performance on inverted and upright faces.

3.6.1. Method

3.6.1.1. Participants

An opportunity sample of 20 Undergraduates from the School of Psychology at Cardiff University took part in this study as partial fulfilment of a course requirement. All participants reported that they had normal vision.

3.6.1.2. Materials

Eighty faces from the EvoFit database[13] (Frowd, Hancock, and Carson, 2004) were used to assess the inversion effect. These were of typical looking males, aged between 18 and 35. All images were presented in 72 dpi resolution and with dimensions 75 mm by 100 mm for either learning or test and with dimensions 125 mm by 167 mm for either test or learning (this was counterbalanced). Inverted versions of the faces were also generated. Fixation crosses were 1 mm thick and 20 mm high and wide. All stimuli were presented on a high-resolution colour monitor and were displayed using SuperlabPro[TM] 2 Research Software. Participants' responses were recorded on a standard computer keyboard, and the participants sat at a distance of 60 cm from the monitor.

---

[13] I would like to thank Charlie Frowd and David A Ross for providing these stimuli.

3.6.1.3. Design

A 2 by 2 by 2 Within-subjects design was employed, whereby the three independent variables were: the orientation of the face at learning and test (upright or inverted); the spatial location of the preceding fixation cross at learning (upper part of the subsequent image or lower part of the subsequent image); and the spatial location of the preceding fixation cross at test (upper or lower). The face stimuli were counterbalanced, such that each appeared as a target and a distractor on equal number of times and each appeared upright and inverted an equal number of times. This created four counterbalanced groups that had equal numbers of participants within them. The location of the fixation cross was randomised according to the rule that it appeared upper the same number of times as it appeared lower for both upright and inverted faces. Recognition accuracy was measured using the signal detection theory measure of stimulus discriminability, d'.

3.6.1.4. Procedure

A standard recognition paradigm was employed similar to that used in Experiment 7, except that the faces were on screen for 1500 ms (selected due to the findings of Experiment 8) each in the learning phase, and the fixation cross appeared for 150 ms. Moreover, as a distractor task, participants filled out the Visual Imagery Questionnaire. All other aspects of the procedure were identical to that in Experiment 7.

3.6.2. Results

The old/new responses were converted into the measure of recognition accuracy d' using the Macmillan and Creelman (2005) method. The means recognition accuracy scores for upright and inverted faces preceded by a fixation cross that was upper or lower are

presented in Figure 3.13. This indicates that when the fixation cross preceded the upper portion of the face, recognition accuracy was better for upright faces than inverted faces. However, when the fixation cross preceded the lower portion of the face, recognition accuracy was lower for upright faces than inverted faces.
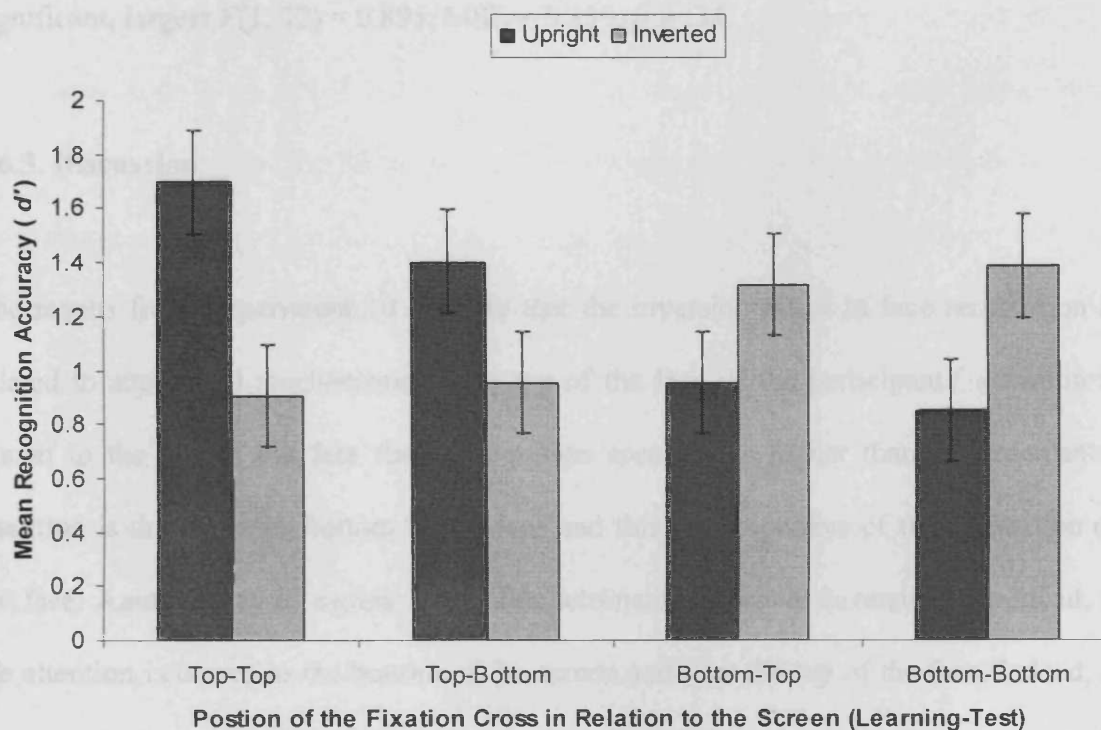
Figure 3.13. Mean recognition accuracy (d') for upright and inverted faces preceding fixation crosses in the upper or lower portion of the face at learning-at test. Error bars represent standard error.

The data summarised in Figure 3.13 were subjected to a 2 by 2 by 2 within-items ANOVA with the factors orientation of the face (upright and inverted); the spatial location of the preceding fixation cross at learning (upper or lower); and the spatial location of the preceding fixation cross at test (upper or lower). This analysis revealed an interaction between orientation of the face and the location of the fixation cross at learning, $F(1, 72) = 15.962$, MSE = 0.359, $p < .05$. Analysis of simple effects revealed

that when the fixation cross was upper at learning, recognition accuracy was greater for upright faces than inverted faces (mean difference = 0.650, p < .05). However, when the fixation cross was lower at learning, inverted faces were better recognised than upright faces (mean difference = 0.422, p < .05). No other interactions or main effects were significant, largest $F(1, 72) = 0.895$, MSE = 0.359, p > .34.

3.6.3. Discussion

The results from Experiment 10 indicate that the inversion effect in face recognition is related to attentional mechanisms to the top of the face. If the participants' attention is drawn to the top of the face then recognition accuracy is higher than if participants' attention is drawn to the bottom of the face and this is irrespective of the orientation of the face. Another way of saying this is the detrimental effect of inversion is reduced, if the attention is drawn to the bottom of the screen and thus the top of the face. Indeed, if attention is directed to the lower portion of an upright face, the recognition accuracy is similar to that of an inverted face that has attention drawn to the eyes and hairline. It is possible, therefore, that previous results relating to the inversion may be mediated by a shift in attention to the lower part of the face caused by inversion.

It is worth noting that the effect of the fixation cross only had an influence at learning and not at test. In previous studies, the orientation of the faces at test was deemed more critical than the orientation of faces at learning (e.g., Valentine and Bruce, 1986). However, these authors and others (e.g., Scapinello and Yarmey, 1970; Yarmey, 1971) always gave participants upright faces to learn. As such, a comparison with such previous research is unwarranted. To the author's knowledge, only one study has examined the effects of learning faces inverted. This study found that the orientation of

the face at learning was comparatively more critical to subsequent face recognition than testing the faces inverted (Hills and Lewis, resubmitted). In this way, the detrimental effect of testing inverted faces was much smaller than the detrimental effect of learning faces inverted.

There is at least one possible concern with the findings from Experiment 10. The faces were always presented in the middle of the screen and the fixation crosses were presented in the upper or lower portion of the screen. This may have provided the participants with a cue as to where the top of the face should be. Experiment 11 was conducted, whereby the fixation cross always remained in the same place and the faces were in a high or low position on the screen. Thus, the fixation cross will not provide a cue as to where the top of the face will be.

3.7. Experiment 11 – Reducing the Inversion Effect with Fixation Crosses Preceding the Forehead Part 2 – The Faces in Different Positions

Experiment 10 indicated that there may be an attentional mechanism to the inversion effect, and that this mechanism has a larger effect at learning than at test. Experiment 11 was conducted whereby the fixation cross was in the centre of the screen preceding all faces. The faces however, were positioned high or low with respect to the screen and either upright or inverted. Thus, when the face is positioned high and upright or low and inverted, the fixation cross preceded the chin and when the face is positioned low and upright or high and inverted, the fixation cross preceded the forehead. Thus, if the effects in Experiment 10 are reliable, there would be an interaction between orientation of the face and the position of the face. Experiment 11 assessed the findings of Experiment 10

that the attentional mechanism influenced the inversion effect at learning more so than at test. As such, faces could be high or low at learning and high or low at test.

3.7.1. Method

3.7.1.1. Participants

An opportunity sample of 20 participants from the Participation Panel at Cardiff University took part in this study for money. All self-reported normal or corrected vision.

3.7.1.2. Materials

The same experimental set-up and face database was used for this Experiment, except that a different set of 64 faces were used. These were either presented raised sufficiently for the fixation cross to appear before the chin, or lowered such that the fixation cross would appear before the forehead for upright faces (and vice versa for inverted faces). An example is presented in Figure 3.14.

3.7.1.3. Design and Procedure

Experiment 11 employed a 2 by 2 by 2 within-items design whereby the independent variables were orientation of the face (upright or inverted), position of the face at learning (high or low), and position of the face at test (high or low). Recognition accuracy was measured using the SDT measure d'. The faces were counterbalanced such that each appeared as a target the same number of times as a distractor. The presentation order of the stimuli was randomised. All other aspects of the procedure were identical to Experiment 10.
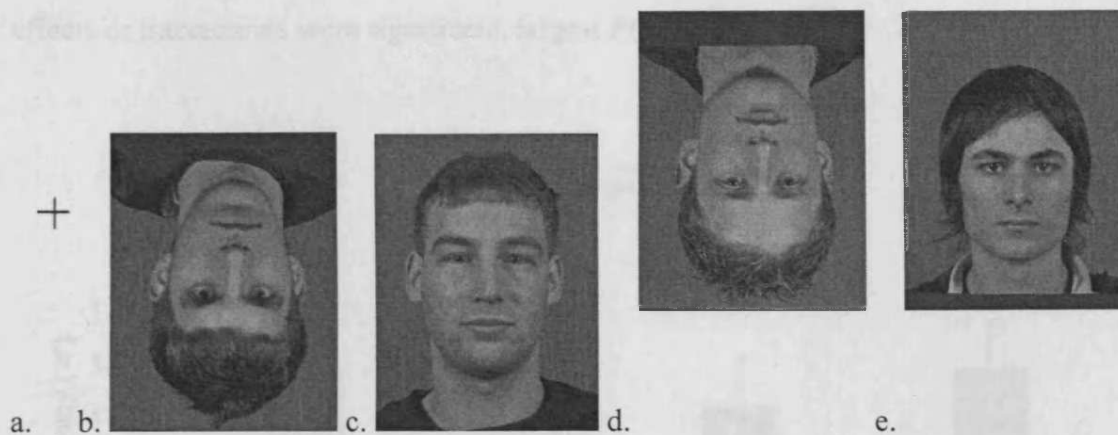
Figure 3.14. Examples of the positioning of the stimuli for Experiment 11: a. fixation cross presented at the same height for all stimuli; b. inverted low; c. upright low; d. inverted high; and e. upright high. The fixation cross is always in the same place thus preceded the: b. mouth level; c. eye level; d. eye level; and e. mouth level.

3.7.2. Results

The mean recognition accuracy (d') data are summarised in Figure 3.15. The pattern of results indicates that upright faces are better recognised if they have been learnt in the low position (recall that this is with the fixation cross preceding the forehead) than inverted faces.

The data summarised in Figure 3.15 were subjected to a 2 by 2 by 2 within-items ANOVA with the factors: orientation of the face (upright or inverted), position of the face at learning (high or low), and position of the face at test (high or low). There was an interaction between orientation of the face and position of the face at learning, $F(1, 56) = 17.149$, MSE = 0.292, $p < .05$. Simple effects were used to explore this interaction and revealed that the upright faces were better recognised than inverted faces when the faces were learnt in the low position (mean difference = 0.513, $p < .05$), but inverted faces

149

were learnt better in the high position (mean difference = 0.605, $p < .05$). No other main effects or interactions were significant, largest $F(1, 56) = 1.375$, $p > .24$.
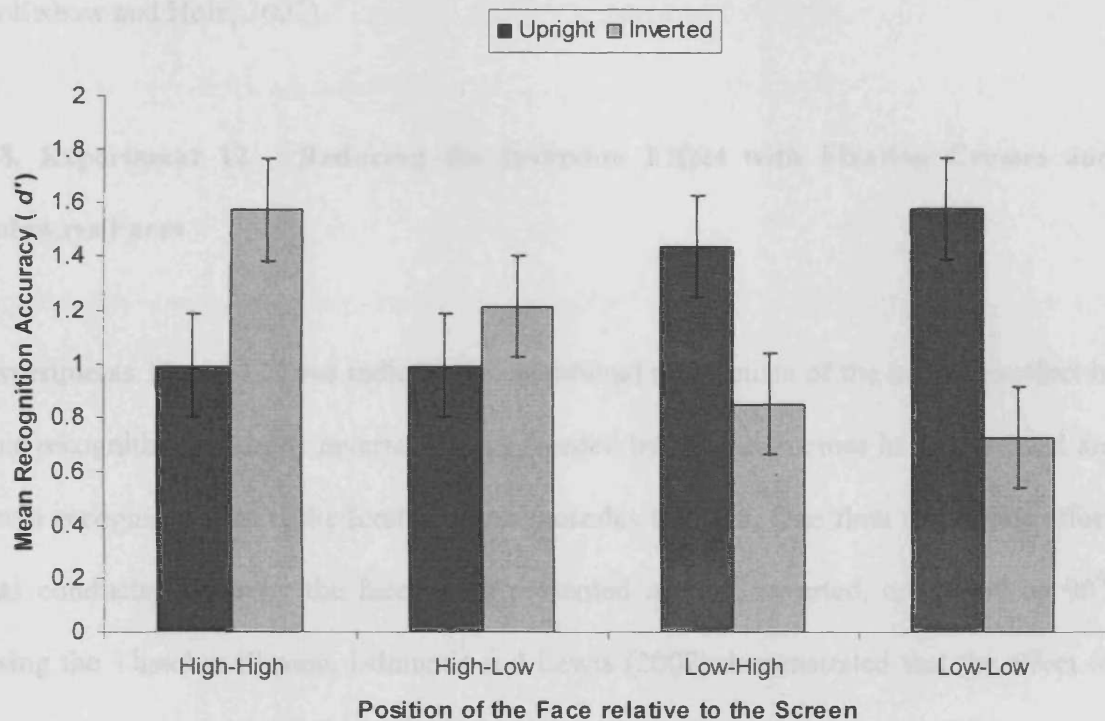
Figure 3.15. Mean recognition accuracy (d') for upright and inverted faces presented either high or low at learning-at test. Error bars represent standard error.

3.7.3. Discussion

The findings from Experiment 11 were consistent with those of Experiment 10. Fixation crosses preceding the forehead at learning of an inverted face prevented the detrimental effect of inversion, whereas fixation crosses preceding the chin at learning of an upright face caused it to be recognised at levels similar to those of inverted faces. This attentional effect is one of learning rather than a test phenomenon. This replication of the effect observed in Experiment 10 suggests that the position of the fixation cross within the

screen was not responsible for the original findings. A third experiment was carried out to assess the effects of this attentional cueing on faces rotated by 90^0 since research has indicated they will be recognised at a level between upright and inverted faces (e.g., Collishaw and Hole, 2002).

3.8. Experiment 12 – Reducing the Inversion Effect with Fixation Crosses and Sideways Faces

Experiments 10 and 11 have indicated an attentional mechanism of the inversion effect in face recognition, whereby inverted faces preceded by a fixation cross in the forehead are better recognised than if the fixation cross precedes the chin. One final test of this effect was conducted whereby the faces were presented upright, inverted, or rotated by 90^0. Using the Thatcher illusion, Edmunds and Lewis (2007) demonstrated that the effect of inversion is a gradual one: The effect of rotation is less for faces rotated by 90^0 than those rotated by 180^0. In a recognition paradigm, it was found that there is a linear disruption to accuracy with rotation (Collishaw and Hole, 2002). Experiment 12 tested the fixation cross procedure in upright, inverted, and sideways faces.

Due to the sideways orientation, it was impossible to use the same position of the fixation cross variable as in Experiments 10 and 11. Instead the position of the fixation cross was coded as preceding the chin or preceding the forehead in all orientation of the faces. As such, an interaction between orientation of the face and position of the fixation cross at learning is not predicted. Rather, a main effect of the position of the preceding fixation cross at learning is predicted, if the results are consistent with the previous experiments.

3.8.1. Method.

3.8.1.1. Participants

An opportunity sample of 20 participants from the Participation Panel at Cardiff University took part in this study for cash. All self-reported normal or corrected vision.

3.8.1.2. Materials

One-hundred-and-twenty randomly selected faces from the EvotFit database (as used in Experiments 10 and 11) were used in this Experiment. These were either presented upright, inverted, or rotated 90^0 to the right or to the left. All other materials were the same as those used in Experiment 10.

3.8.1.3. Design and Procedure

This study employed a slightly different design from Experiments 10 and 11. A 3 by 2 by 2 within-items design was employed, whereby the independent variables were orientation of the face (upright, inverted, or sideways), position of the preceding fixation cross at learning (forehead or chin), and position of the preceding fixation cross at test (forehead or chin). In all other respects, the design and the procedure of this Experiment were identical to those in Experiments 10 and 11.

3.8.2. Results

The mean recognition accuracy (d') data are summarised in Figure 3.16. The pattern of results indicates that a fixation cross preceding the forehead is beneficial to the recognition of faces, irrespective of the orientation of the face.

The results summarised in Figure 3.16 were subjected to a 3 by 2 by 2 within-items ANOVA with the factors orientation of the face (upright, inverted, or sideways), position of the preceding fixation cross at learning (forehead or chin), position of the preceding fixation cross at test (forehead or chin). There was a significant main effect of the position of the preceding fixation cross at learning, $F(1, 108) = 15.286$, MSE = 0.789, $p < .05$, whereby recognition accuracy was greater when the fixation cross preceded the forehead, irrespective of the orientation of the face (mean difference = 0.634). No other main effects, or interactions were significant, largest $F(1, 108) = 0.395$, MSE = 0.789, $p > .65$.
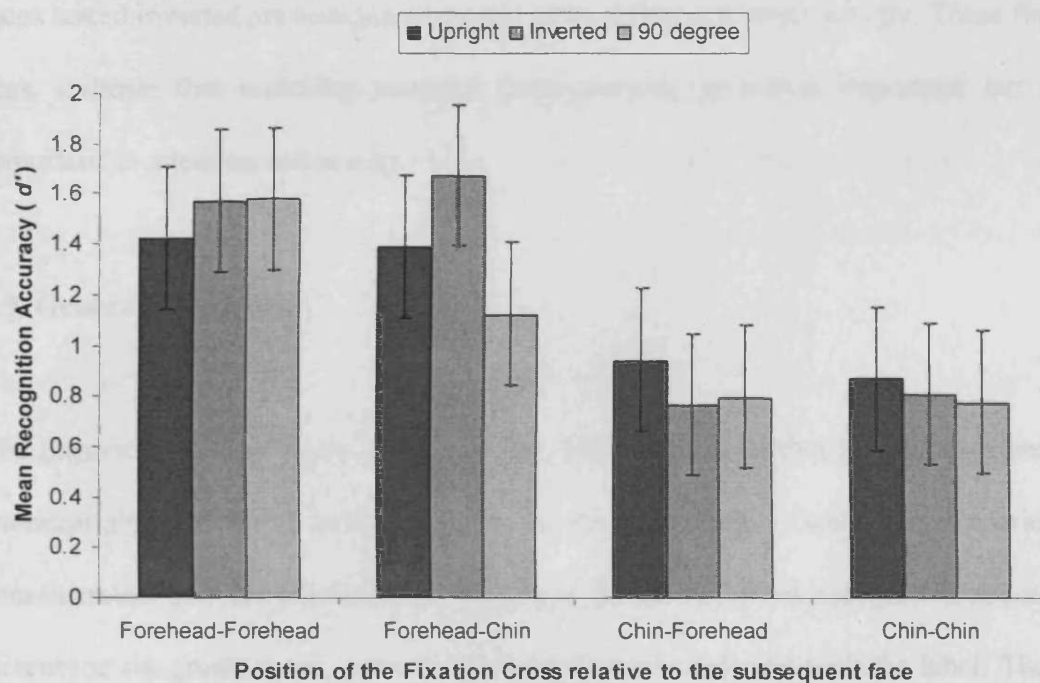
Figure 3.16. Mean recognition accuracy (d') for upright, inverted, and sideways faces split by spatial location of the preceding fixation cross at learning-at test. Error bars show standard error of the mean.

153

3.8.3. Discussion

The results of Experiment 12 are consistent with those of Experiments 10 and 11. Learning a face is more successful if the face is preceded by a fixation cross to the forehead irrespective of the face's orientation. There are some minor inconsistencies in the data between these three experiments when the spatial location of the fixation cross is not matched between learning and test. The mismatch conditions are broadly consistent with the same attentional effect at learning. However, the effect of mismatch in attention does appear to reduce the accuracy with which the face will be recognised in some cases but not all. This is not surprising given the findings of Valentine and Bruce (1986b) that faces tested inverted are recognised poorly even if they are learnt upright. These findings, thus, indicate that matching attention from learning to test is important, but not as important as attention at learning.

3.9. General Discussion

The Experiments reported in this chapter are indicative of the fact that the distribution of attention alters the way in which faces are encoded within face-space. Experiment 5 demonstrated that the provision of stereotype labels improved recognition accuracy of stereotype congruent faces, provided that the faces were learnt with the label. The ORB was reduced in Experiments 6 to 9 through training White participants to focusing on features diagnostic for the recognition of Black faces (Experiment 6) or by presenting a fixation cross preceding the diagnostic features of Black faces (Experiments 7 to 9). The reversal of the ORB due to the presentation of a fixation cross preceding diagnostic features of Black faces was not found if the participants were forced to delay their

154

responses during the learning of the faces. Experiments 10 to 12 found that focusing attention to the forehead of an inverted face removed the inversion effect.

Taken together the findings from the Experiments presented in this chapter indicate that attentional weighting to particular features of faces alters the way in which they are encoded and subsequently recognised. If the original encoding is paid toward the most diagnostic visual features then recognition accuracy will be greater than if the original encoding is paid toward less diagnostic visual features. Since the visual features are considered to be dimensions within the face-space, this research indicates that attention paid toward relevant dimensions of face-space improves recognition memory for faces, whereas attention paid to irrelevant dimensions of face-space causes a detriment to face recognition.

As in Chapter 2, it is important to analyse these results within the three face-space frameworks presented in the introduction. Shrinking-face-space can account for these results since it suggests that the adult face-space contains dimensions that can be activated for particular tasks. It suggests that attention to particular dimensions can be changed depending on cognitive, affective, and perceptual activation. Thus, this model allows for two types of flexibility: flexibility in the number of dimensions and flexibility in the use of dimensions. Constant-face-space based on the idea that all dimensions are used equally (from face-space-r, Lewis, 2001), does not explicitly allow for the weighting of attention to particular dimensions to be changed. However, the original constant-face-space (Valentine, 1991) can allow for attentional weighting to alter which dimensions are attended for a given task. It may be, for example, that attention paid to a particular dimension is proportional to ease of encoding on that dimension. Ease of encoding is akin to exemplar strength within face-space-r (Lewis, 2001) and is inversely related to error associated with such encoding. Expanding-face-space also does not preclude this second

type of flexibility, in that different dimensions can be used for different tasks, however, it has not been explicitly stated in previous models of face-space.

The findings of this Chapter also demonstrate the importance of attention paid during the initial learning of faces given that other-race faces learnt with the correct attentional guidance are better recognised irrespective of the attentional guidance during the test phase of an experimental procedure. What is most interesting is that attentional during learning is more important than a mismatch in attention from learning to test, suggesting that once a face is learnt it has relatively robust representation.

Following on from the importance of learning, the results from Experiment 9 demonstrate that although the initial learning is important, when participants are forced to delay their responding to a face during learning, they override the experimentally induced attentional guidance and implement their preset scan. Thus, participants will actually replace a largely effective scanpath over faces with a less effective scanpath and any features learnt during the effective scan are replaced. This indicates that the attentional system will inhibit the features not necessarily used for encoding.

Interpreting these findings within the face-space metaphor, it suggests that the way faces are encoded is dependent on the cognitive, perceptual, and attentional state. Attention directs participants to encode faces according to a particular subset of dimensions. These may or may not be the most diagnostic in the recognition of the faces being encountered. When attention is being paid to the most relevant dimensions of face-space, recognition accuracy will be superior due to the increased accuracy of encoding (c.f., Valentine, 1991), and thus exemplar strength (c.f., Lewis, 2004). Presently, the face-space model does not explicitly account for any semantic information to affect processing within it. To allow for this, one must accept that face-space is the best representation for the FRU pool in the IAC (Valentine, et al., 1995).

156

The results of Experiment 5 are largely difficult to explain within the current IAC architecture. One possible explanation is based upon the idea that the stereotype label, as an SIU, activates all known faces of that particular stereotype. In order for semantic information to affect the processing of a given face then a PIN for that face needs to mediate the spread of activation from an SIU to an FRU. When a face is novel it will neither have an FRU nor a PIN. However, the updated version of the IAC (Burton, et al., 1999) suggests that FRUs are connected to Eigenfaces in a PCA analysis. Thus, an SIU for "actors" will activate all known actor PINs and FRUs. If stereotypes were accurate then all the FRUs for actors would be similar and thus the Eigenfaces associated with known and unknown "actors" will be have increased activation. However, Bull and Green (1980) note that stereotype labels are not accurate reflections of individuals who actually are of that particular label. If one assumes that the stereotype label only brings to mind known FRUs that actually fit with the prototype of that particular stereotype (for example, the stereotype label "criminal" would bring to mind a thug-like actor who is type-cast into portraying criminals such as Steve McFaddon), then the explanation provided above works. However, this analysis presupposes that the strength of the links between SIUs and PINs have different values to each other and those FRUs that are more visually similar to the prototype SIU will have stronger connections between them. It also presupposes that there are verbal SIUs and visual SIUs that represent the prototype (or verbal SIUs that describe the visual). Since Eigenfaces can be represented by dimensions within face-space, this explanation can be illustrated within face-space (Figure 3.17).

In the above explanation, the Eigenface associated with stereotypical criminals will have narrow and deeply set eyes. Phrasing this within face-space, the dimensions of face-space that are associated with stereotypical criminals are the narrowness of the eyes and how deeply set they are. If these dimensions, or regions along these dimensions, have

raised activation, then encoding onto these regions becomes easier. Thus, the face-space

has flexibility in how the dimensions are used as well as when they are used.
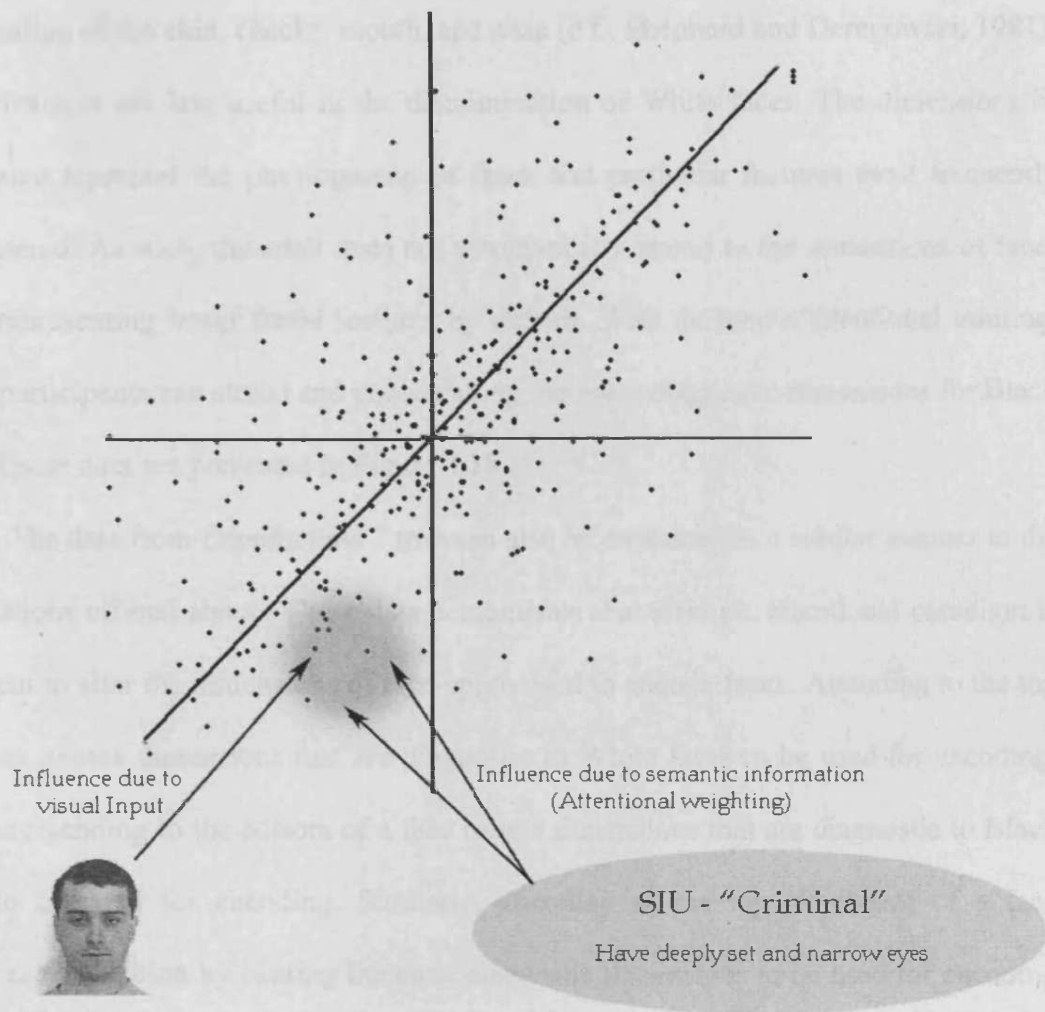
Figure 3.17. A representation of the proposed interaction between semantic and visual

information in face-space. The SIU "criminal" activates all known and stereotypical PIN

"criminals" which in turn activates Eigenfaces or a region of face-space representing the

visual features that are stereotypically criminal (e.g., narrow and deeply set eyes). This

allows for ease of encoding into this region by an incoming "criminal" face. In this

example, the Y axis represents distances between the eyes and the Z axis represents how

deeply set the eyes are and the X represents, for example, hairline.

The results of Experiment 6 are broadly consistent with this theoretical treatise. Black faces are best discriminated if the lower facial features are used involving a combination of the chin, cheeks, mouth, and nose (c.f., Shepherd and Deregowski, 1981). These features are less useful in the discrimination of White faces. The dimensions of face-space represent the physiognomy of faces and particular features most frequently encountered. As such, the adult does not automatically attend to the dimensions of face-space representing lower facial features by default. With an hour's attentional training, White participants can attend and encode along the more diagnostic dimensions for Black faces. These data are presented in Figure 3.18.

The data from Experiments 7 to 9 can also be explained in a similar manner to the explanations offered above. These data demonstrate that a simple attentional paradigm is sufficient to alter the dimensions of face-space used to encode faces. Attending to the top of a face causes dimensions that are diagnostic to White faces to be used for encoding, whereas attending to the bottom of a face causes dimensions that are diagnostic to Black faces to be used for encoding. Similarly, attending to the top (forehead) of a face improves recognition by causing the most diagnostic dimensions to be used for encoding irrespective of the orientation of the face as demonstrated in Experiments 10 to 12.
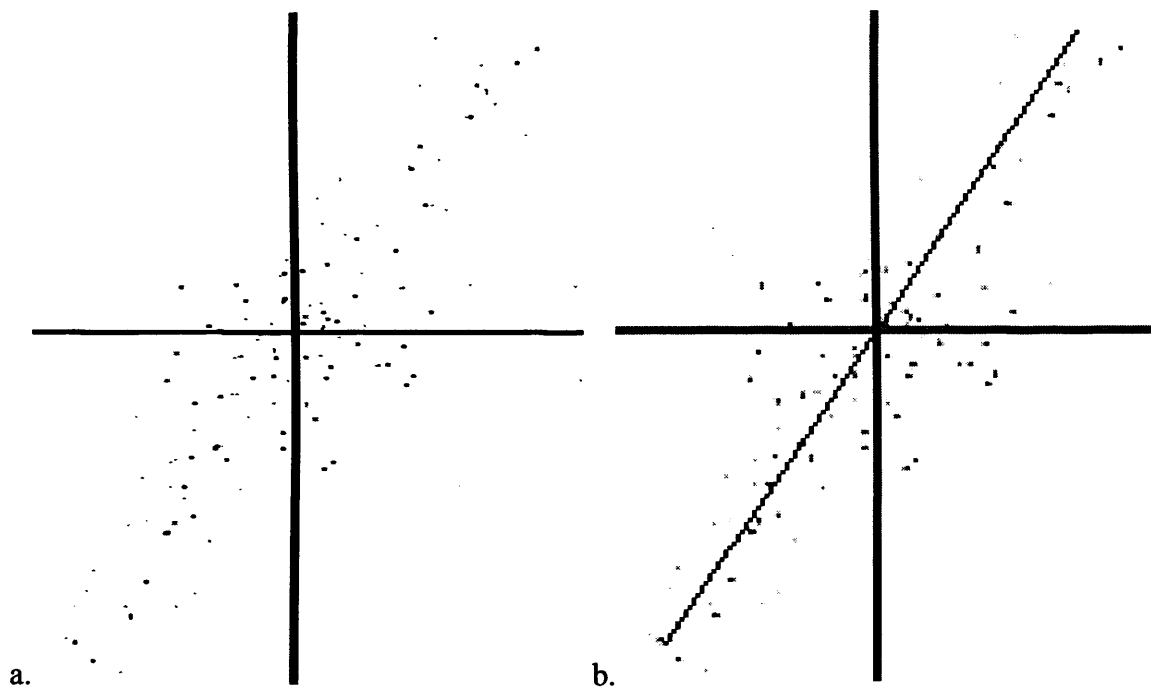
Figure 3.18. A 2-dimensional representation of face-space to explain the results of Experiment 6 showing the distribution of other-race faces (Black faces for White participants) and the default dimensions used to encode faces along the x and y axes. The X axis may represent hairline and the Y axis may represent eye colour. In a., the other-race faces are not easily discriminated by these two dimensions. Activating a more appropriate dimension, as in b., makes the other-race faces easier to discriminate. In the example provided, the Z-axis may represent width of nose.

These results are indicative of the face-space being a relatively flexible system, whereby whole dimensions or parts of them can be attended to or ignored depending on the task at hand, in addition to dimensions being added or removed. Such an idea contradicts the two assumptions of Lewis' (2004) face-space-r, which suggests that face-space a sufficient number of dimensions to recognise all faces and that these are all necessary (the sufficiency and necessity assumptions) leading to an implied assumption that all dimensions are used equally. The results presented here suggest that a set of dimensions

or region within a dimension that is more diagnostic for discriminating the face presented will lead to increased recognition accuracy.

Conclusions

The experiments reported in this chapter established that attentional mechanisms can alter the dimensions used to encode faces. Attentional direction can be caused through priming, training, or fixation crosses. The importance of the first glance in face recognition has been established here, although the first glance can be replaced by default attention. In the following chapter, adaptation shall be used to alter how faces are encoding along particular dimensions of face-space.

Chapter 4 – Recalibration of Face-Space

Six experiments reported in this chapter explored how the weighting along particular dimensions can be recalibrated using visual and non-visual adaptation. Experiment 13 employed a two-alternative forced-choice paradigm to demonstrate visual after-effects due to adaptation to identity-specific non-visual information. A replication using a free response procedure is presented in Experiment 14. Experiment 15 explored the involvement of imagination in the visual face after effects using a psychophysical procedure. Experiments 16 and 17 also employed a psychophysical procedure and tested visual after-effects following adaptation to various types of visual and non-visual stimuli in familiar faces (Experiment 16) and unfamiliar faces (Experiment 17). Experiment 18 assessed the after effects following adaptation involving configural or featural processing. Experiment 19 explored the duration of the FIAE.

4.1. Experiment 13 – Demonstration of Identity Adaptation Part 1 – Forced Choice

Face after-effects have been shown by several researchers whereby after examining an image for some time, a perceptual midpoint appears shifted in the opposite direction (see section 1.2.3.). A demonstration is shown in Figure 4.1, in which examining Pierce Brosnan for one minute causes the midpoint in a morphed continuum from Pierce Brosnan to Harrison Ford to appear to be made up of mostly Harrison Ford. An explanation of these effects is that prolonged presentation of an identity causes a recalibration of particular dimensions in face-space (e.g., Hurlbert, 2001). This

recalibration may be caused by adaptation to neuronal populations in the face-space (thereby located in the visual cortex, see Robbins et al., 2007) or specific to the identity of the individual (thereby located in the memory centres of the brain). It is possible to find behavioural evidence as to the location of the face after effects, and this is addressed in Experiment 13.
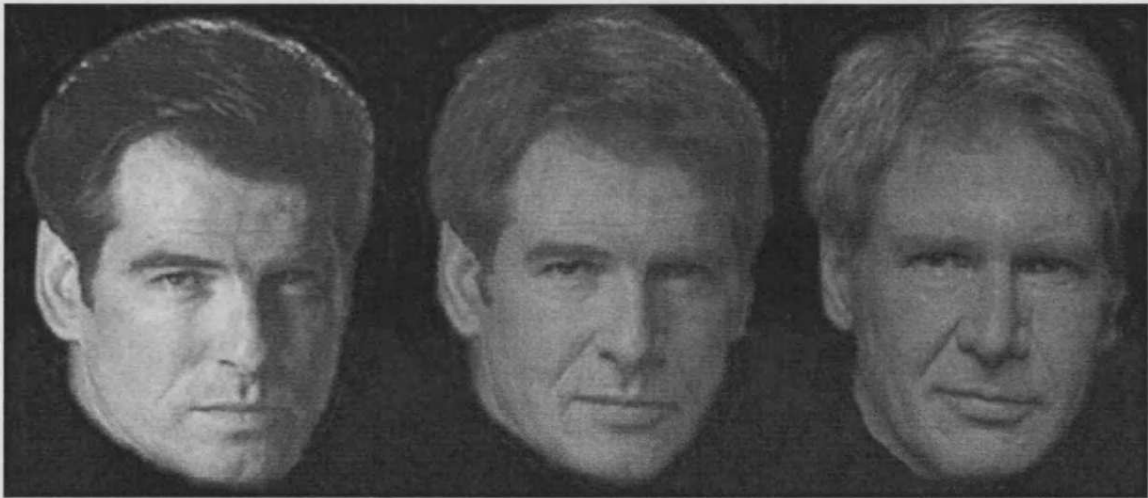
Figure 4.1. Demonstration of identity after effects. The middle figure is ambiguous, but after adapting to the figure on the left, he becomes obvious.

Priming research indicates that there are times when non-visual (names and voices) information speeds up the recognition of faces. This cross-modal priming effect has been reported several times (see e.g., Bruce and Valentine, 1986; Johnston and Barry, 2006; Young, Flude, Hellawell, and A. Ellis, 1994; Young, Hellawell, and de Haan, 1988) and is neatly accounted for by the IAC (see Section 1.1.4.). Experiment 13 was conducted to assess whether the face after-effect finding is due to the visual properties of the face solely or whether the effect is due to neuronal populations representing identity. Experiment 13 used a two-alternative forced-choice procedure to identify cross-modal adaptations within the IAC, by pre-exposeing participants to the identities used. To do this, two familiar pictures of faces were morphed together, creating a 50% identity.

Familiar faces were used, since this requires less training and fewer instructions. Single faces were used to prevent demand characteristics, as after one condition, participants would be aware of the experimental hypothesis.

4.1.1. General Method

4.1.1.1. Participants

Participants consisted of 100 staff and students at Cardiff University with normal or corrected vision who participated voluntarily. Twenty participants took part in each condition.

4.1.1.2. Materials

Two different images of Pierce Brosnan and Harrison Ford were collected. They were paired based on dimension (pair A was 83 mm by 140 mm; pair B was 100 mm by 100 mm), pose (pair A were frontal with slight tilt to the right; pair B were three-quarters facing left), and lighting (pair A was lit from the left; pair B was lit from the front). All images had a resolution of 72 dpi. Each pair was morphed together using Smartmorph™ Software, using 200 anchor points. A series of morphs of each were created that ranged from 100% Harrison Ford to 100% Pierce Brosnan in increments of 5%. These were pretested using a staircase procedure, such that the point that was rated as Harrison Ford on 50% trials was found. It was found to be 55% Harrison Ford for image pair A, and 45% Harrison Ford for image pair B. These perceptual midpoints were used as the test stimuli. The 100% images for each identity and each pair were used as the adaptation stimuli. All stimuli were presented on an RM PC running Microsoft™ Windows XP™ using Powerpoint™ Presentation software.

4.1.1.3. Design

The dependent variable for the experiments was identity response given to the midpoint image. The adaptation stimuli were varied between-subjects. Half the participants were adapted to Harrison Ford, the other half were adapted to Pierce Brosnan. Half the participants were tested on image pair A, whereas the remaining participants were tested on image pair B. A Chi-square test revealed that there was no effect of identity (48% Harrison Ford responses) or image pair. As such, the data was pooled across these variables. In each experiment, chance levels would be indicated by 50% (10 out of 20) participants giving a response as the non-adaptor.

4.1.1.4. Procedure

Participants were tested individually or in groups of up to three. The procedure consisted of three short phases: the instructions; adaptation; and the test phase. In the instructions, participants were shown a picture of Harrison Ford and informed who he was, followed by a picture of Pierce Brosnan and his name. The faces and the names were on screen for 1 second each. All participants knew who they were. Participants were then instructed to examine adaptor for the duration it was on screen. Immediately following adaptation, the test morph appeared on screen. Participants were instructed to write down the initials of the person the morph looked most like, and could take as long as they required. This never took longer than 20 seconds. Figure 4.2 shows a demonstration of this task.
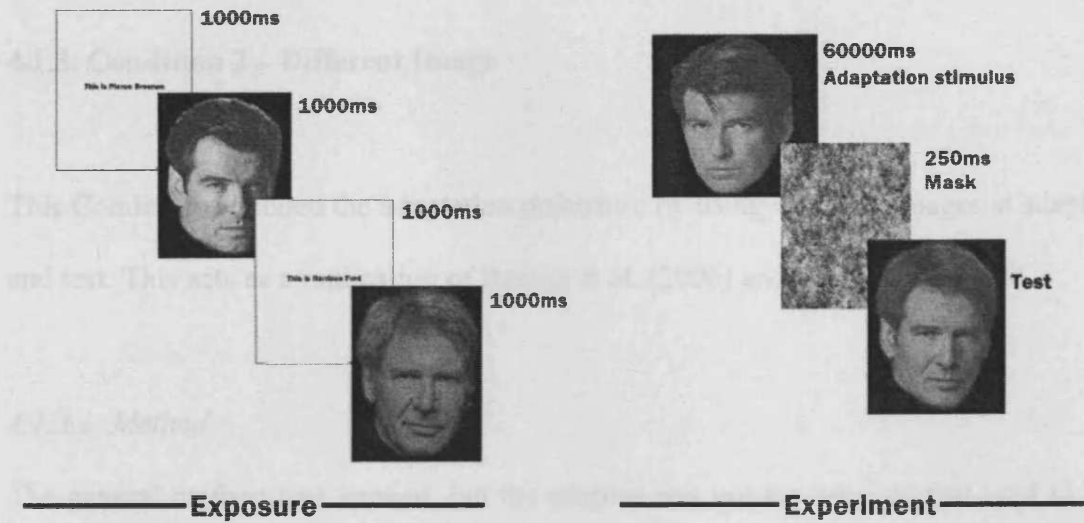
Figure 4.2. Procedure for Experiment 13, with Pierce Brosnan as the adaptor.

4.1.2. Condition 1 – Same Image

This condition acted as a replication of the Leopold et al. (2006) study, testing whether FIAEs are reliable for familiar faces.

4.1.2.1. Method

The method was as described above. The adaptor was the same image as that used to make the morph.

4.1.2.2. Results

If participants were unaffected by the adaptor, then the midpoint image would be identified as either identity at chance levels (50%, 10 out of 20). This was not the case. Ninety percent (18 out of 20) of participants gave a response that was not the adaptation identity. This was significantly different to chance and reflects adaptation taking place, χ^2 (1) = 7.619, $p < .01$.

166

4.1.3. Condition 2 – Different Image

This Condition extended the adaptation procedure by using different images at adaptation and test. This acts as a replication of Benton et al. (2006) and Jiang et al. (2006).

4.1.3.1. Method

The general method was applied, but the adaptor was not the same as that used to make the morph. That is, there was a pose change from the adaptation to the test phase.

4.1.3.2. Results and Discussion

Adaptation to a different image than the one tested with caused a significant after-effect, χ^2 (1) = 7.619, p < .01, again with 90% of participants giving adapted responses (18 out of 20). This effect is indicative of the viewpoint-independent representation of familiar faces (c.f., Benton et al., 2006). Simple image-based adaptation effects have not been shown to occur if the test image is rotated.

4.1.4. Condition 3 – Written Name

Identity is made up from many different aspects, including voices, names, biographical information, and the visually-derived aspects. Condition 3 set out to test whether the FIAEs can be caused by other aspects of identity.

4.1.4.1. Method

Instead of using an image as the adaptor, the written name was used (presented in the middle of the screen for 60s). All other aspects of the method were as described in the general method.

4.1.4.2. Results and Discussion

Eighty percent of responses given were adapted when the adaptor was the name (16 out of 20). This was not significantly different from the same or different images χ^2 (1) = 1.318, p > .51, but was significantly different from chance, χ^2 (1) = 3.956, p < .05.

When the adaptor is a written name, it causes FIAEs of similar magnitude to those caused by images. This finding is evidence for non-visually caused face identity after effects, whereby responses to a visually presented face are being altered due to the prolonged exposure of the adaptation identity, irrespective of the modality of presentation. This important effect may well be controversial and one must be wary of demand characteristics playing a role in the present study.

4.1.5. Condition 4 – Written Nationality

Since the mode of presentation in Condition 3 was different from Conditions 1 and 2 (image to name), an additional condition was run, whereby the written nationality of the faces' identity was used as the adaptor. Nationality is linked to identity, but is linked to many identities. As such, it is not a finely tuned piece of semantic information. It is unlikely to affect adaptation since low-level adaptation effects are usually quite finely tuned (c.f., Webster & Miyahara, 1997). This condition controlled for mode of

presentation. This condition also controlled for demand characteristics, since participants are aware of the 'opposite' identity.

4.1.5.1. Method

An identical method was used as in Condition 3, except the adaptation stimuli were written nationalities (American or British). Since few participants were aware that Pierce Brosnan is Irish, it was felt that the nationality British would suffice as the opposite of American in this experiment.

4.1.5.2. Results and Discussion

Responses for each identity were exactly at chance levels: 50% (10 out of 20) responding with the identity associated to the adaptation nationality. This condition is significantly different from the three conditions described above.

4.1.6. Condition 5 – Brief Presentation

Condition 4 indicated that the mode of presentation does not influence demand characteristics. However, there is still a possibility that negative priming (c.f., Wiggs & Martin, 1998) might be occurring, whereby the prolonged stimulation of one identity actually primes the alternative identity. Negative priming requires a presentation time of less than 1000 ms to occur. To control for this, the adaptor was presented for 250 ms rather than 60 seconds.

4.1.6.1. Methods

A similar method to the general method was employed, except the adaptor was on screen for 250 ms, which is long enough for it to be registered by the participants, but is not prolonged enough to cause adaptation (c.f., Wiggs & Martin, 1998). Immediately following the adaptor, the test stimulus appeared, as such there was no gap between the two stimuli. The adaptor was the same image as that used to make up the morph.

4.1.6.2. Results and Discussion

Thirty percent of responses were the non-adaptor identity (6 out of 20). This trend was in the opposite direction to the adaptation effects, but was not significantly different from chance, χ^2 (1) = 1.667, p > .19. This finding indicates the possibility of priming occurring when the stimuli are presented briefly – a finding that is similar to the effect of repetition priming observed by A. Ellis, Young, Flude, and Hay (1987). The effect was not significant, which is unusual, and may indicate that some demand characteristics are occurring whereby participants believe they are supposed to respond with the opposite identity. Alternatively, this methodology may not be appropriate for bringing about strong priming effects.

4.1.7. Overview and Discussion

Five experimental conditions demonstrated that participants can be adapted to a specific identity from the face and the name. The results in terms of number of responses of the non-adapted identity are shown in Figure 4.3. There were no significant differences across the same face, different face, and name adaptation stimuli. These three conditions demonstrated significant adaptation.
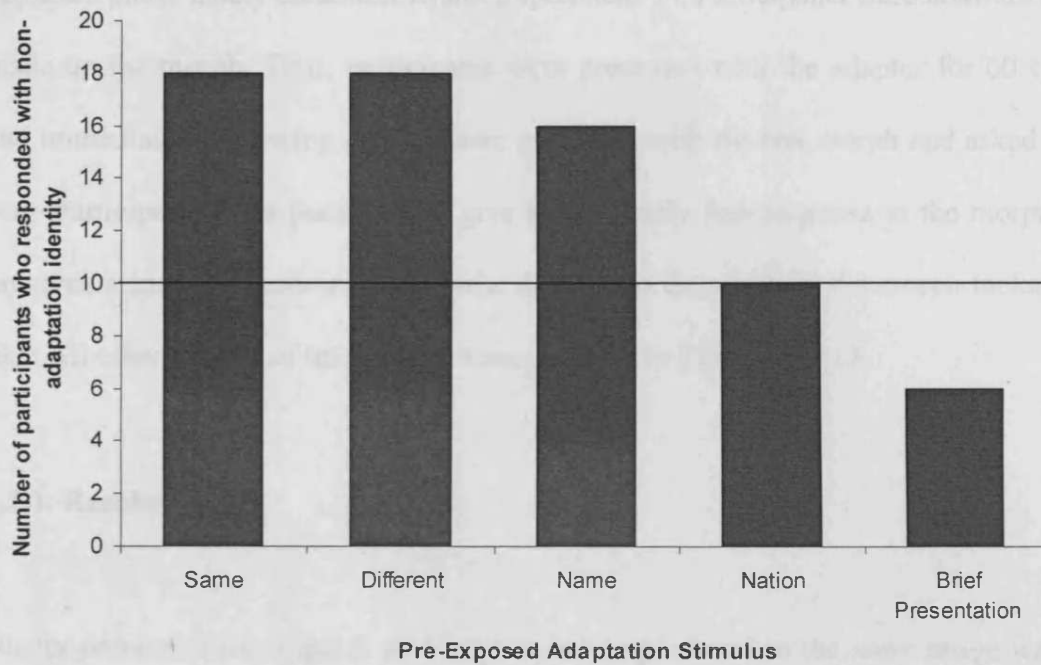
Figure 4.3. Number of participants who responded with the non-adapted identity split by type of adaptor for Experiment 13 (pre-exposed). Chance levels are 10 participants.

There is a possibility, however, that these effects are due to a demand characteristics. Since the participants in the conditions were pre-exposed to both of the face identities that made up the morph, they may have felt that they were supposed to give the "opposite" answer. Indeed, the methodology was virtually a two-alternative forced-choice paradigm, since the participants knew that there were only two possible responses. Due to this, Experiment 14 was carried out.

4.2. Experiment 14 – Demonstration of Identity Adaptation Part 2 – Free Response

Experiment 14 replicates Experiment 13 using a new set of participants controlling for possible demand characteristics (N = 100, 20 in each condition, see Experiment 13). The only difference between Experiment 14 and Experiment 13 is that there was no pre-

exposure phase in any condition within Experiment 14. Participants were unaware of who made up the morph. Thus, participants were presented with the adaptor for 60 seconds and immediately following it, they were presented with the test morph and asked who it was. Participants were permitted to give a completely free response to the morph. This involved asking the participants to write down who they thought the morph looked most like. All other aspects of the method were identical to Experiment 13.

4.2.1. Results

Ninety percent of participants gave an adapted response when the same image was used at test and adaptation (18 out of 20), akin to Experiment 13: Condition 1. This was significant as before, χ^2 (1) = 7.619, p < .01. Eighty-five percent of participants gave adapted responses when the images were not the same at adaptation and test (17 out of 20), akin to Experiment 13: Condition 2, χ^2 (1) = 3.956, p < .05. Seventy-five percent of responses were adapted when the adaptor was the name (15 out of 20), similar to Experiment 13: Condition 3. This proportion is significantly different from chance, χ^2 (1) = 3.467, p < .05. When the adaptor was nationality, 55% percent of responses were adapted (11 out of 20), akin to Experiment 13: Condition 4. This proportion is not significantly different from chance. For brief presentation (Experiment 13: Condition 5), 100% of responses were the adaptor (20 out of 20), indicating total priming occurred. This was significantly different from chance, but in the opposite direction to the adaptation effects observed in all other conditions, χ^2 (1) = 13.333, p < .01. These results are presented in Figure 4.4. The results from Experiment 13 and 17 were combined and subjected to a logistic regression analysis. This analysis determined that there was not an

effect of Experiment on the data (χ^2 (1) = 1.428, p > .23), suggesting that pre-exposure

and no-pre-exposure did not affect the pattern of results.
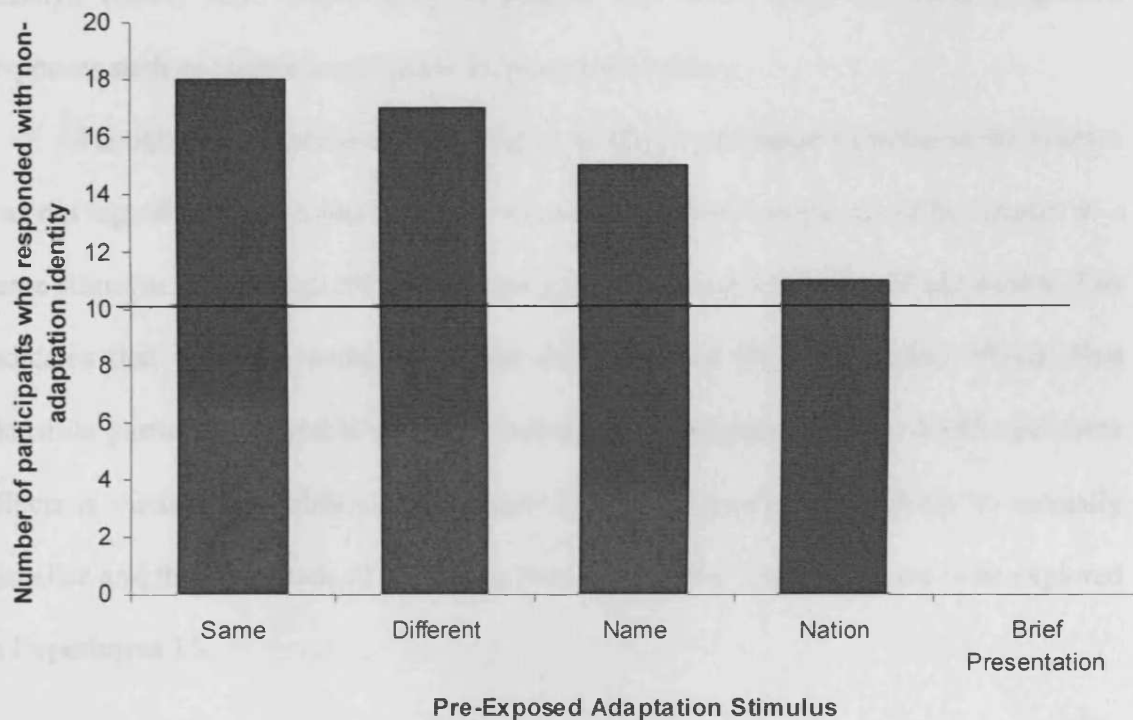
Figure 4.4. Number of participants who responded with the non-adapted identity split by

type of adaptor for Experiment 14 (free response). Chance levels are 10 participants.

4.2.2. Discussion

The adaptation effects observed here have changed the detection thresholds of perceived

images (Webster & MacLin, 1999) when the adaptor was an image or a name. This

finding is important because it suggests the existence of neuronal populations that

173

selectively respond to particular identities. Nevertheless, it is unclear what participants are doing when they are presented with a name. When presented with a name, participants are likely to visualise the face belonging to it. This may mean that the route of this identity adaptation is due to imagination. Indeed, O'Craven & Kanwisher (2000) have shown that, using fMRI techniques, there is selective activation in the fusiform gyrus for perceiving *and* thinking about faces. Furthermore, Ganis, Thompson, and Kosslyn (2004) have noted how imagination can affect even low-level adaptation responses such as contrast and spatial frequency adaptation.

Though the difference between the same image and name stimulus as the adaptor was not significant, there was a trend to suggest that fewer people could be adapted to a name stimulus. Indeed, not all participants gave a response indicative of adaptation. This indicates that there are some individual differences in these adaptation effects. One plausible participant variable that may influence the magnitude of the FIAE and these effects is visualisation ability. Experiment 15 looks at participants' ability to mentally visualise and the magnitude of adaptation they experience. These areas are to be explored in Experiment 15.

4.3. Experiment 15 – FIAEs are Mediated and Moderated by Visualisation Abilities

Questions remain unanswered regarding the FIAEs reported in Section 1.2.3 and those found in Experiments 13 and 14. This is the cause of the individual differences in the magnitude of the reported FIAEs (c.f., Jiang et al., 2006; 2007). One possible explanation for individual differences in the magnitude of these after effects between participants is their ability to mentally visualise the image they are presented with (Stillman & Kemp, 1993). It is possible for people to identify faces in frontal views even if they have only

ever seen them in profile views previously, but recognition performance is significantly lower following slight changes in viewpoints from learning to test (Bruce, Valentine, & Baddeley, 1987). It has been suggested that when presented with a face in either frontal or profile views, people mentally rotate the image into a ¾ view, thus visualising and creating a ¾ view representation (e.g., Kersten, Troje, & Bulthoff, 1996; Valentine & Bruce, 1988; see also, Troje & Bulthoff, 1996). There is a great deal of variation in people's ability to generate mental images both in terms of quality of mental image (Richardson, 2000) and in terms of types of mental imagery (Hasnain & Husain, 1980; Richardson & McAndrew, 1990).

Mental imagery and visualisation is known to activate the visual centres of the brain (e.g., Kosslyn, et al., 1993; 1994). O'Craven & Kanwisher (1999) demonstrated that the fusiform gyrus is active when perceiving faces *and* when thinking about faces. Since there are individual differences in the activation of the visual cortex (c.f., Ishai, Schmidt, & Boesiger, 2005), there are likely to be individual differences in the magnitude of fusiform gyrus activation due to imagination.

A multiple response psychophysical procedure was employed for Experiment 15, as the magnitude of adaptation can be assessed for different participants and different adaptation stimuli. Moreover, this method requires fewer participants than a single response procedure. Indeed, many psychophysical studies involve only the authors (such as, Benton, et al., 2006) or only three participants. Furthermore, it made it possible to replicate the identity adaptation effects in a second methodology using different face identities, thus extending the generalisability of the results. If visualisation abilities affect the magnitude of adaptation, the participants who have higher visualisation abilities as measured by the Visual Imagery Questionnaire (VIQ, Marks, 1973) should demonstrate a greater magnitude of adaptation than participants who scored lower on the VIQ. The VIQ

measures the ability of individuals to visualise specific scenes such as a sunrise, a person, and a shop. To assess the mental processes, an additional adaptation condition was incorporated, whereby participants were explicitly told to think about the person. Two faces morphed together are characterised by categorical perception between the two (Beale, and Keil, 1995; Rotshtein, Henson, Treves, Driver, and Dolan, 2005). The perceptual midpoint in the continuum from one face to another is called the Point of Stimulus Equality (PSE) and reflects the strength of identity required to perceive one identity in the morph. The change in PSE pre- and post-adaptation is used as the measure of the magnitude of adaptation. Psychophysical functions were fitted to calculate the PSE.

4.3.1. Method

4.3.1.1. Participants.

Sixteen Cardiff University students undertook this experiment as partial fulfilment of a course requirement. All participants had normal or corrected vision.

4.3.1.1. Materials

Two different images of George Bush and Tony Blair were collected. They were matched for dimensions (100 mm by 160 mm) and resolution (72 dpi). Image one of George Bush was matched for pose and lighting with image one of Tony Blair. A series of morphs were created for image pair one and image pair two, using Smartmorph™ Software with 200 anchor points. Fifty morphs were created that ranged from 100% George Bush to 100% Tony Blair in increments of 2%. Image two of George Bush was in a different pose and under different lighting conditions from image one of George Bush and matched to

image two of Tony Blair. The image two pairs were morphed together in the same way as the image one pair. The 100% images for each identity and each pair were also used as the adaptation stimuli.

Names were displayed on screen in Palatino Font, size 20, black on white. All stimuli were presented using SuperlabPro 2™ Research Software on an RM PC.

4.3.1.2. Design

The four adaptation stimuli (same image as at test, different image to the one at test, name, and imagined) were varied between subjects. Scores on the VIQ were split by the median score of 12 creating a second independent variable, ability to visualise. This created a 4 by 2 between-subjects experiment. Participants were randomly selected to be in one of the adaptation conditions in such a way to maintain an equal number of participants in all conditions. Participants were either adapted to George Bush or Tony Blair and either image pair one was used or image pair two was used. There was no effect of identity or image identity, as such all effects are collapsed across these variables for the analysis.

4.3.1.3. Procedure

Participants were introduced to pictures of George Bush and Tony Blair that they would see in the experiment. The Experiment had three consecutive phases: baseline, adaptation, and test. The baseline phase involved the participants seeing all the morphs 10 times in a random order. They had to make a decision based on whether they thought the image looked more like George Bush (by pressing the G key) or Tony Blair (by pressing T) based on the methodologies in Levitt (1971). Each morph was on screen until

the participant responded. Between each morph a 100 ms white noise mask was on screen. The procedure for this is presented in Figure 4.5.
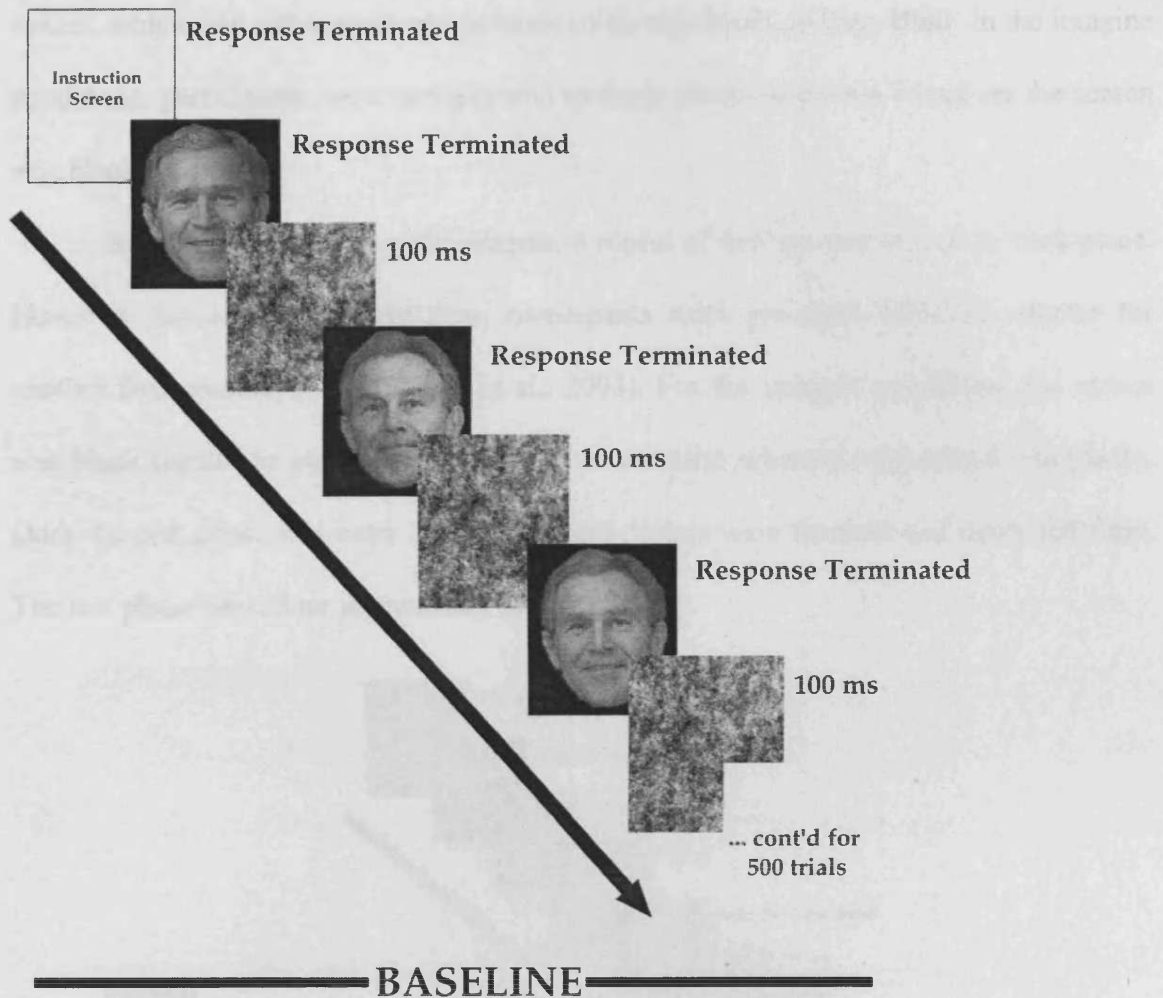
Figure 4.5. The procedure for the baseline phase of Experiment 15.

Once the baseline had finished, the participants were instructed to rest for two minutes. Following this, participants were given the VIQ. As instructed by Marks (1973), the VIQ was read out by the experimenter who marked down how elaborate the visual memory was that the participants commented on. Prior to the experimental research, two experimenters assessed visual imagery separately using the VIQ and were found to have an inter-rater reliability of r = .82.

Once the questionnaire had been completed, participants were presented with the adaptor for 60 seconds. They were told to examine the image that was presented on screen, which was either an image or name of George Bush or Tony Blair. In the imagine conditions, participants were verbally told to think about the person whenever the screen was blank.

Immediately following the adaptor, a repeat of the baseline procedure took place. However, preceding each test face, participants were presented with the adaptor for another five seconds (c.f., Rhodes, et al., 2003). For the imagine conditions, the screen was blank (recall the participants were told to visualise whenever the screen was blank). Once the test phase had been completed, participants were thanked and debriefed fully. The test phase procedure is presented in Figure 4.6.
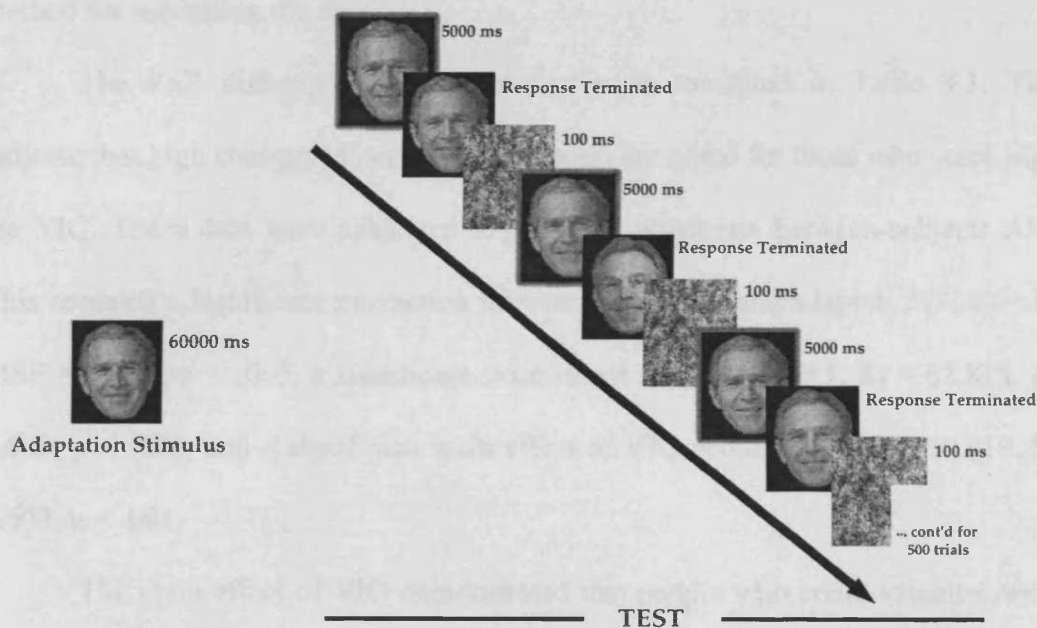
Figure 4.6. The procedure for the test phase. Those surrounded by the box are the adaptor repeated during the test phase.

4.3.2. Results

Psychometric functions[14] were fitted before and after adaptation for each participant using Matlab™ (Version 7.2) mathematical programming software. These are presented in Figure 4.7. Of interest is the location of perceptual midpoint (the PSE), the homogeneity of variance of the function, and the slope shape. For all the participants who scored highly (the upper median after a median-split) on the VIQ and when the same image is used at adaptation and test for low VIQ scorers, the psychometric function fits the observed data reliably. However, the psychometric functions were not always reliable for participants who scored low on the VIQ, and thus may not be the most appropriate method for modelling the data.

The PSE differences post-adaptation were tabulated in Table 4.1. The data indicate that high changes in perceptual midpoint are noted for those who score higher on the VIQ. These data were subjected to a 2 by 4 univariate between-subjects ANOVA. This revealed a significant interaction between VIQ score and adaptor, $F(3, 8) = 10.013$, MSE = 3.603, p < .005, a significant main effect of adaptor, $F(3, 8) = 62.818$, MSE = 3.603, p < .001, and a significant main effect of VIQ score, $F(1, 8) = 130.019$, MSE = 3.603, p < .001.

The main effect of VIQ demonstrated that people who could visualise well were significantly more likely to show adaptation than those who were less good at visualisation (mean difference = 10.823, $p < .05$). The main effect of adaptor revealed that the differences between all the stimuli were significant after Tukey HSD *post hoc* tests, with significant greater adaptation to the same stimulus at adaptation and test over

---

[14] I would like to thank Rhodri Woodhouse (Cardiff University) for providing the Matlab code to fit the heterogenous Probit and homogenous Probit functions.

the different stimulus (mean difference = 8.825, $p < .05$), name stimulus (mean difference = 17.158, $p < .05$), and imagined adaptation (mean difference = 14.143, $p < .05$). The interaction revealed itself in that for low visualisation ability participants, the differences between adaptation to a different image, the name, or an imagined stimulus were all not significantly different from each other nor chance ($p > .5$), whereas for those with high visualisation abilities, the mean differences were all significantly different from each other and different from chance ($p < .05$).

Table 4.1.

Perceptual midpoint changes post-adaptation. Positive numbers indicate more of the adaptation identity was needed in the morph to see that identity. Significant shift in the PSE ($p < .05$) is denoted by an asterix (calculated using t-tests).

VIQ Score (Visualisation Ability)	Adaptor (compared to test images)	Adapted to:		Mean
		George Bush	Tony Blair	
High	Same image	23.68%	24.54%	24.11%*
	Different image	21.17%	20.37%	20.77%*
	Name	4.42%	6.61%	5.52%*
	Imagined	13.38%	7.85%	10.62%*
Low	Same image	15.21%	16.06%	15.64%*
	Different image	2.43%	0.22%	1.33%
	Name	1.71%	-1.88%	-0.09%
	Imagined	1.62%	0.07%	0.85%
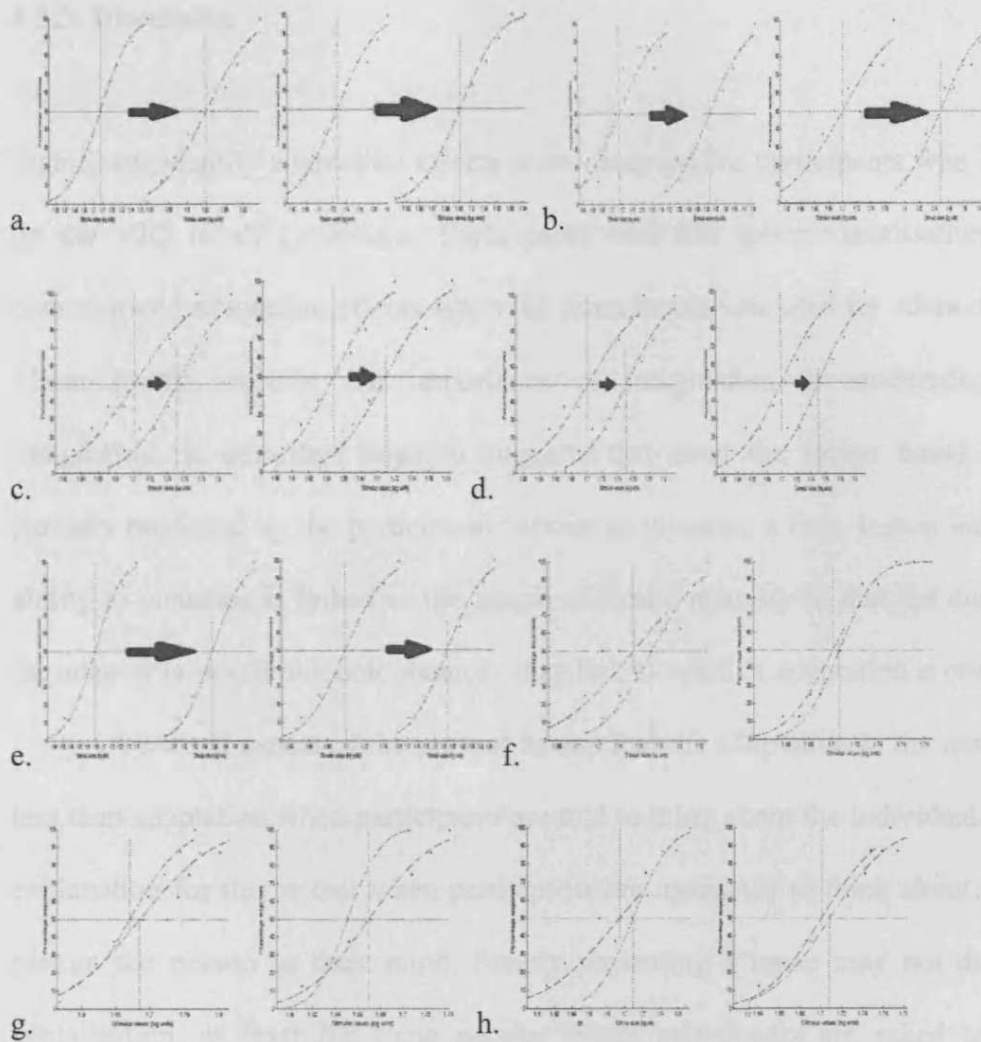
Figure 4.7. Results and psychophysical functions before and after adaptation. Sample solution based on median PSE for that condition. The arrow indicates significant changes in PSE. Left panels are those adapted to George Bush. Right panels are those adapted to Tony Blair. Top four panels are participants who are scored above 12 on the VIQ. a. Same adaptation image as that used at test. b. Different adaptation image to those used at test. c. Name used as adaptation stimuli. d. Imagined adaptation stimuli. Bottom four panels are participants who scored less than 12 on the VIQ. e. Same adaptation image as that used at test. f. Different adaptation image to those used at test. g. Name used as adaptation stimuli. h. Imagined adaptation stimuli.

182

4.3.3. Discussion

Significant identity adaptation effects were observed for participants who scored higher on the VIQ in all conditions. Participants who had lower visualisation scores only demonstrated adaptation effects when the same image was used for adaptation as at test. These results indicate the importance of imagination in moderating the FIAE. Imagination is important since it suggests that even the image based adaptation is partially mediated by the participants' ability to visualise a face. It may indicate that the ability to visualise is linked to the length of iconic memory in that the duration of time the adaptor is stored in iconic memory may link to whether adaptation is observed or not.

A second aspect of the present data is that the adaptation to the name stimulus is less than adaptation when participants are told to think about the individual. One possible explanation for this is that when participants are instructed to think about someone they picture the person in their mind. Simply presenting a name may not directly link to visualisation, at least for some people. When participants are asked to think about someone, they will usually picture them rather than think about semantic information (c.f., Stillman & Kemp, 1993).

The results from this Experiment contrast with those of Moradi, et al. (2005), who found that FIAEs required conscious awareness of the adaptor. However, Moradi et al.'s work used faces that were unfamiliar to the participants and thus harder to visualise. Moreover, the accuracy with which the participants visualised in Moradi et al.'s study was not tested with any scientific accuracy. This limitation is especially important given the wide ranging individual differences in mental imagery (Marks, 1973). To explore differences between the present data and those of Moradi et al., Experiments 16 and 17 were conducted.

183

4.4. Experiment 16 – Cross-Modal FIAE for Familiar Faces

The identity of an individual is expressed through numerous methods. The most obvious and frequently used aspect to distinguish identity is the face. However, if facial information is obscured in some way, identity can still be extracted from other aspects, including body language (Ambady, Hallahan, & Conner, 1999), surface information (Bruce et al., 1993), head templates (Sinha and Poggio, 1996), and voice (B. Clifford, 1980). Voice recognition is not as efficient as face recognition, and typically takes longer to accomplish and is more prone to misidentification (D. Yarmey, 1995). Moreover, biographical information (such as occupation) is harder to obtain from voices (Hanley, Smith, and Hadfield, 1998).

Within face recognition, it has long been known that the processing of one aspect of identity aids in the recognition of other aspects of identity (e.g., Bruce & Valentine, 1986). A. Ellis, et al. (1987) reported three studies where participants were instructed to state whether or not a face was familiar. When participants had seen a photograph of a familiar face in a preceding stage of the experiment they were faster at identifying that the face was familiar by 71 ms than if they had not seen the face in the experimental session (Experiment 1). However, familiarity judgements were not made quicker if the participant had seen the face's written name prior to the face (Experiment 2). If the prime photograph was the same as that used as the target, the familiarity judgement was quicker by 196 ms in Experiment 3, compared to 163 ms faster if the photograph was only similar. A dissimilar photograph reduced reaction times by 104 ms. These results indicate the importance of perceptual similarity in the repetition priming effect (see also Ellis, D. Ellis, & Hosie, 1993).

These results are slightly different from those subsequently found by A. Young, et al. (1994), whereby written names sped up familiarity judgements of faces only if the prime preceded the face by a short period of time. Moreover, Ellis, Jones, and Mosdell (1997) found that voices can prime faces. However, this cross-modal repetition priming effect only occurs if the prime precedes the target face by a maximum of five seconds. Johnston and Barry (2006) have shown that cross-modal repetition priming effects are significantly smaller than within-modal priming effects, but are still significant themselves.

Bruce and Valentine (1986) found that there are also associative priming effects, whereby the face or name of a person will speed up the identification of a name or face of a related person. For example, the face of Ronnie Barker sped up the familiarity judgements made to Ronnie Corbett's face because they are a comedy duo and often seen together. Furthermore, Young, et al. (1988) have shown an associative form of cross-modal priming, whereby the name of a famous person speeds up recognition of a highly associated famous face. For example, reading the name of "Stan Laurel" will speed up the recognition of the face of "Stan Laurel" and the recognition of the face of "Oliver Hardy". Category priming (Carson and Burton, 2001) is where priming of the category label speeds up recognition of all known exemplars of that category. For example, the category label "comedians" will speed up recognition of "Stan Laurel" and all other known comedians. Category priming is considerably weaker than associative or semantic priming.

These cross-modal priming effects form the basis for the hypothesised face after effects in Experiment 16. Since priming effects have been observed that are cross-modal and the neurological architecture of the visual system is set up to allow for non-visual information to impact on the processing of visual information, after effects are likely to

be caused by non-visual modalities. The present work, therefore, aimed to test whether cross-modal adaptation effects can be found and whether they are similar in their form to the priming effects in terms of the types of stimuli that will cause them. Thus, names, semantic information, associated people, and voices can be used to adapt the perception of a target face. Identity threshold of a particular facial identity (c.f., Leopold et al., 2001) in a morphed continuum between two faces can be assessed pre- and post-adaptation, whereby after effects would be revealed by a higher identity threshold post-adaptation.

Familiar (famous) faces are used in Experiment 16 since they have established semantic, visual and often auditory information associated with them. This information was used to establish the degree to which it led to adaptation of the identity. Twelve identities were used in the present experiments. For names, voices, category information, and imagination, four highly familiar faces were used. These were George Bush, Tony Blair, Harrison Ford, and Pierce Brosnan. For the caricatures and the associated people, comedy double acts were used. Two types of caricature present themselves to be used for such a study. Photographic caricatures are realistic computer-generated images that have been exaggerated according to a mathematical formula. These would be useful to test visually-derived identity adaptation. Alternatively, there are natural caricatures, drawn by caricature artists. These are not as similar to veridical images as photographic caricatures, and may contain the artist's own 'filling in.' However, they can get across features that are joked about which may have semantic information associated with them. For example, there is a piece of semantic information that Prince Charles has big ears. A natural caricature will make these even more prominent than they would be in a photographic caricature. Moreover, natural caricatures are used for practical reasons (easiness to obtain).

4.4.1. Method

4.4.1.1. Participants.

Participants were 118 Cardiff University students who took part as partial fulfilment of a course requirement or for £8 cash. All participants had normal or corrected vision and were White British nationals who were familiar with the target identities. Participants were randomly divided into one of the conditions. There were unequal group sizes, due to the number of adaptor identities used in some conditions. The number of participants in each condition is presented in Table 4.2. Unequal group sizes will affect the variance of the experimental groups, yet this is inevitable due to the identity pairings.

4.4.1.2. Materials

Morph continua between the identity pairs described in Table 4.2. were created as described in Experiment 15. For the face, name, voice, category information, and imagination adaptors, faces of four identities were collected from the Internet. These were George Bush, Tony Blair, Harrison Ford, and Pierce Brosnan. The category information used was: largely relevant (President, Prime Minister, Indiana Jones, and James Bond); quite relevant (Politicians and Actors); and largely irrelevant (American and British). The morphed continua used were George Bush to Pierce Brosnan and Tony Blair to Harrison Ford.

For the voice stimuli, two sound clips were collected from Internet. They contained no interruption from the audience. The clips were cut down to 1 minute in length, chosen to be the second to third minute of the speeches in question. This meant that the clip came in mid sentence, and cut out mid sentence. Two further 5 second clips were made from the next minute of the clips in the same manner.

Table 4.2.

The types of adaptor and test morph continua associated with each and the number of

participants in each condition for Experiment 16.

Adaptor	Identities Used	Morph Continua	N
Face: Same image	George Bush Tony Blair Harrison Ford Pierce Brosnan Ant Macfarland and Dec Donnerly Matt Lucas and David Walliams Eric Morecombe and EarnieWise The Two Ronnies: Ronnie Barker and Ronnie Corbett	GB – TB GB – PB TB – HF AM – ML AM – DW DD – ML DD – DW EM – RB EM – RC EW – RB EW - RC	22
Face: Different image	George Bush Tony Blair Harrison Ford Pierce Brosnan	GB – TB GB – PB TB – HF PB - HF	8
Name	George Bush ·Tony Blair Harrison Ford Pierce Brosnan	GB – TB GB – PB TB – HF PB - HF	8
Voice	George Bush Tony Blair Harrison Ford Pierce Brosnan	GB – TB GB – PB TB – HF PB - HF	8
Specific Semantic Information	George Bush Tony Blair Harrison Ford Pierce Brosnan	GB – PB TB – HF	4
Occupation Semantic Information	George Bush Tony Blair Harrison Ford Pierce Brosnan	GB – PB TB – HF	4
Nationality Semantic Information	George Bush Tony Blair Harrison Ford Pierce Brosnan	GB – PB TB – HF	4
Associated Person	Ant Macfarland and Dec Donnerly Matt Lucas and David Walliams Eric Morecombe and EarnieWise The Two Ronnies: Ronnie Barker and Ronnie Corbett	AM – ML AM – DW DD – ML DD – DW EM – RB EM – RC EW – RB EW - RC	32
Caricature	Ant Macfarland and Dec Donnerly Matt Lucas and David Walliams Eric Morecombe and EarnieWise Ronnie Barker and Ronnie Corbett	AM – ML AM – DW DD – ML DD – DW EM – RB EM – RC EW – RB EW - RC	16
Imagination	George Bush Tony Blair Harrison Ford Pierce Brosnan	GB – TB GB – PB TB – HF PB - HF	12

Names and category information were displayed on screen in Palantino Font, size 20, black on white. All stimuli were presented using a SuperlabPro 2™ Research Software on a Toshiba™ Tecra M4 Tablet PC™. Voices were played into the laboratory using headphone speakers.

For the associated people, pictures of famous double-acts were collected from the Internet. They were: Anthony McPartlin (AM) and Declan Donnelly (DD) from British television show "Saturday Night Takeaway"; David Walliams (DW) and Matt Lucas (ML) from the British television show "Little Britain"; Eric Morecombe (EM) and Earnie Wise (EW) from the British television show "The Morecombe and Wise show"; and Ronnie Barker (RB) and Ronnie Corbett (RC) from the British television show "The Two Ronnies". The adaptor face was either one of the morphed continua extremities, or the double act partner. As such, AM acted as the actual person in the AM-DW morph and the associated person in the DD-DW morph. This created 32 adaptation stimuli-morph continua pairings for the associated people condition.

Caricatures of the 8 double act pairings were collected from the Internet and served as the adaptors, with the morphs from the associated people continua used above.

4.4.1.3. Design

The experiment was designed as a series of individual experiments, each investigating the relative size of the observed adaptation effect for a particular class of stimuli. Given the amount of similarity between these experiments, they are now expressed as a single experiment, albeit, with different numbers of participants in each level of the main manipulation. The adaptor was the only experimentally manipulated variable in this experiment. The identity of the adaptation stimuli (that is either identity 1 or 2 from the adapting continuum) was counterbalanced. The types of adaptation stimuli were

manipulated between-subjects and were: Face (same image as that used at test), Face (different image from that used at test), Name, Voice, Specific semantic information, Occupation semantic information, Nationality semantic information, Associated person, Caricature, or Imagination adaptation. The order of presentation of the faces during the pre-test and the test phase was randomised. Each participant saw a morphed continuum made up of only two identities and saw (or heard) only one adaptor.

4.4.1.4. Procedure

The procedure for Experiment 16 was similar to the procedure in Experiment 15, except for the response keys[15]. Moreover, when the adaptor was a voice, the screen was blank and the sound clip was played through headphone speakers.

4.4.2. Results

The measurement of the magnitude of adaptation was based upon the logic used by Moradi et al. (2005), whereby the required identity strength to perceive an identity post-adaptation can be taken from the identity strength to perceive the same identity pre-adaptation. The identity strength is equated to the PSE in a psychometric function. Psychometric functions were fitted before and after adaptation, using MatLab™, for each of the 118 participants (due to space limitations these are not presented here). The differences between the perceptual midpoints are summarised in Figure 4.8.

---

[15] The response keys were: G for George Bush; T for Tony Blair; H for Harrison Ford; P for Pierce Brosnan; A for Ant; D for Dec; W for David Walliams; L for Matt Lucas; M for Eric Morecambe; E for Earnie Wise; B for Ronnie Barker; and C for Ronnie Corbett. Each Participant only had two identities to choose from.
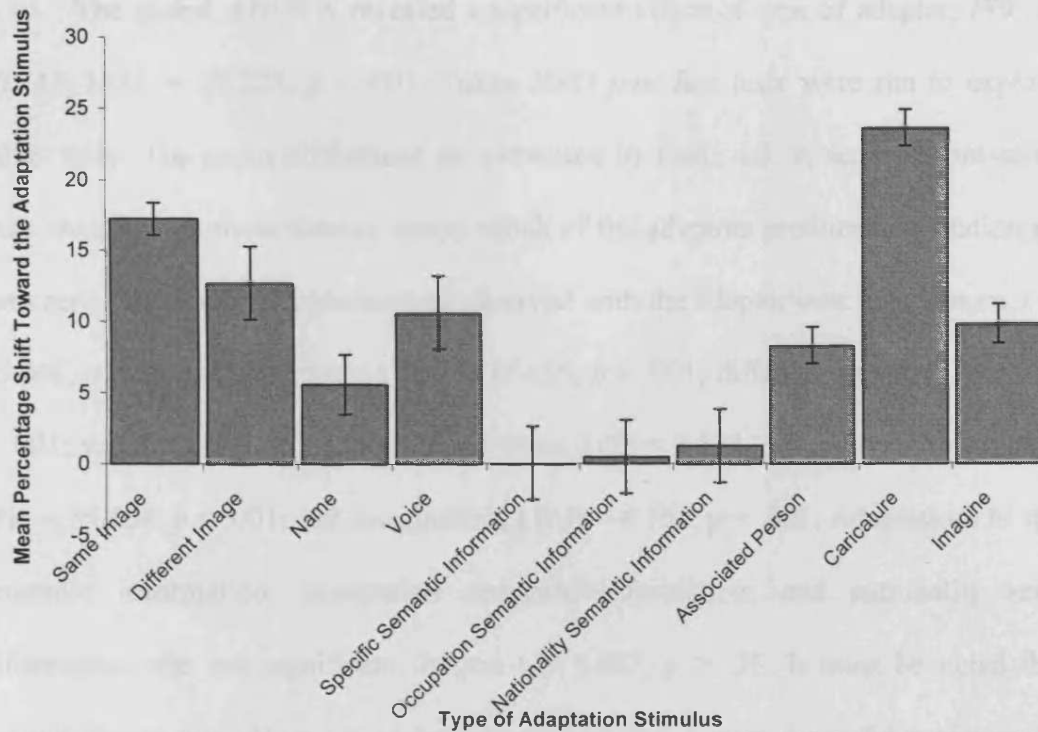
Figure 4.8 Mean percentage shift of the PSE toward the adaptor post adaptation for all the types of adaptation stimuli. Error bars represent standard error. Zero is indicative of no adaptation.

Post-adaptation responses tended to include more opposite identity responses than pre-adaptation responses. The perceived perceptual boundary between the two identities was thus shifted in the direction toward that identity which the participants had been adapted to. The data summarised in Figure 4.8. were subjected to a 10-level one-way univariate ANOVA with the single factor being the type of adaptor. As mentioned in the method, this Experiment had unequal group sizes, which led to unequal standard error. Due to this, Levine's test of homogeneity of variance was conducted and revealed that significant heterogenous variance, $F(9, 108) = 8.958$, $p < .05$. The F_{max} test was consistent with this, $F_{max} = 24.7$. Due to this, the alpha level for the analysis was lowered to $\alpha = .001$.

191

The global ANOVA revealed a significant effect of type of adaptor, $F(9, 108) = 30.743$, MSE $= 19.227$, $p < .001$. Tukey HSD *post hoc* tests were run to explore this effect fully. The mean differences are presented in Table 4.3. A series of one-sample t-tests were run on these data to assess which of the adaptors produced adaptation greater than zero. Significant adaptation was observed with the adaptor was: a caricature, t (15) = 16.068, $p < .001$; same image, t (22) = 17.155, $p < .001$; different image, t (7) = 6.464, $p < .001$; voice, t (7) = 15.112, $p < .001$; name, t (7) = 3.524, $p < .01$; associated person, t (31) = 35.654, $p < .001$; and imagination, t (11) = 4.754, $p < .001$. Adaptations to specific semantic information, occupation semantic information, and nationality semantic information was not significant, largest t = 1.482, $p > .23$. It must be noted that the present design was somewhat underpowered and thus a more powerful design may have revealed significant adaptation to semantic information.

Table 4.3.

Mean differences in magnitude of adaptation between the 10 different types of adaptor. Significant differences are marked with * ($p < .001$).

	Different	Voice	Name	Specific Semantic Information	Occupation Semantic Information	Nationality Semantic Information	Associated Person	Imagination	Caricature
Same	2.846	7.154*	11..079*	14.720*	18.440*	17.736*	9.474*	9.523*	-5.787*
Different		4.308	8.233*	11.874*	15.593*	14.889*	6.628*	6.677*	-8.633*
Voice			3.925	7.566*	11.286*	10.582*	2.321	-2.369	-12.941*
Name				3.641	7.381*	6.657*	-1.604	1.556	-16.866*
Specific Semantic Information					3.719	3.016	-5.246*	-5.197*	-20.507*
Occupation Semantic Information						-0.704	-8.965*	-8.917*	-24.227*
Nationality Semantic Information							-8.261*	-8.213*	-23.523*
Associated Person								-0.485	-15.261*
Imagination									-15.310*

4.4.3. Discussion

Significant face identity adaptation effects were observed for all the types of adaptation stimuli tested here, except semantic information. These results indicate that FIAE can be caused by non-facial information (i.e., voices and names) in addition to facial information. The adaptation effects observed here are therefore cross-modal. The within-modal adaptation was found to have a greater magnitude than that of cross-modal adaptation.

The magnitude of adaptation was greatest to caricatured images. This finding is intriguing since caricatures are not veridical images and thus should not cause such strong low-level adaptation. It can be hypothesised that the better the image is processed, the more it activates the personal information. In other words, though the same image is visually most veridical, some of the semantic information concerns how caricatured the face could be. Using Bruce and Young's (1986) terminology, faces can be recognised based on a verbal code as well as a visual code. It maybe that a caricature contains information about the visual properties of a face, but also contain some aspect of the verbal code, since they have specific features enlarged as would be expected in a verbal description (Ant has a big forehead – his caricature has this feature enlarged). Adaptation to non-identity specific information was not possible in the present study. Names may well be a special form of semantic information (c.f., Burton & Bruce, 1993). In this case, names can cause adaptation to faces, but other forms of semantic information apparently cannot.

These results also indicate that imagination is sufficient to cause adaptation. These results are consistent with the evidence that visual imagination activates the visual centres of the brain but are inconsistent with the results of Moradi et al. (2005). The

reason for this discrepancy is probably based on the different methodologies. In the present study, familiar faces that participants could visualise in a multitude of poses, views, and actions were used. The visualisation is likely to be more elaborate than a visualisation of a facial image that has only been seen in one pose. Thus, the results of the present study are likely to reflect a more realistic effect of mental imagery on adaptation than that presented in the Moradi et al. paper. This possibility was explored in Experiment 17 by exploring adaptation for faces that are unfamiliar prior to the experiment.

4.5. Experiment 17 – Cross-Modal FIAE for Unfamiliar Faces

Experiment 16 demonstrated that it was possible to cause adaptation in face representations by prolonged presentation of visual and specific non-visual stimuli. The purpose of Experiment 17 was to generalise the findings from familiar faces to unfamiliar faces, dealing with the potential problem of prior semantic knowledge of the identities.

A series of unfamiliar faces were associated with semantic information of varying degrees of specificity. If the findings of Experiment 16 generalise then prolonged exposure to a name should cause adaptation. Furthermore, the two additional levels of specificity of semantic information can show whether learnt non-specific semantic information can also produce adaptation effects. This type of learning paradigm has been employed in much previous research and is considered reliable (Moradi et al., 2005).

4.5.1. Method

4.5.1.1. Participants

Twelve Cardiff University Psychology undergraduates took part in this study as partial fulfilment of a course requirement. All were White British Nationals. Each had normal or corrected vision.

4.5.1.2. Materials

A set of 24 faces was collected from the PICs at Stirling face database (available at www.pics.stir.ac.uk). These were male faces and each had the dimensions 110 mm by 90 mm with a resolution of 72 dpi. Two average faces in terms of distinctiveness and attractiveness were chosen to be used as the target faces. These were morphed together as described in Experiment 15. All 24 faces were given 3 pieces of semantic information: one was unique to each face (a name or specific semantic information) and was a letter of the alphabet; one was shared with a quarter of the other faces (semi-relevant semantic information) and were Arabic numerals; and one was shared with half of the other faces (largely irrelevant semantic information) and were Greek letters. Finally, a completely irrelevant adaptor was used to act as a control. The associated pieces of semantic information were presented on screen in size 36 Palatino font, black in colour.

4.5.1.3. Design

A six-level between-subjects one-way design was implemented, whereby the adaptor type was the independent variable: face; specific information; semi-relevant semantic information; largely irrelevant semantic information; totally irrelevant semantic information; and imagination. Face identity was not a significant predictor of adaptation

strength and thus not considered. The dependent variable was the magnitude of adaptation that occurred, in terms of percentage shift in the perceptual boundary from one identity to the other away from the identity to which prolonged exposure had been given.

4.5.1.4. Procedure

This experiment had two main phases. The first was a learning phase. In this phase, the participants were shown all 24 faces in sequence for ten seconds each, coupled with a piece of semantic information and the appropriate key to respond. Once all 24 faces had been presented once with the first piece of semantic information, the participants were given an associative learning task. In this task, participants were presented with a face and had to respond with the appropriate key for the piece of semantic information associated with that particular face. The face remained on screen until the participant gave a response. If participants gave an incorrect response, they were informed of the correct response on screen. The error screen was displayed with the face for 3 seconds. This task continued until the participants reached an accuracy level of 90% on all faces and 95% on the two crucial target faces. Typically this took 200 trials (similar to Moradi et al., 2005). This procedure was repeated with the second piece of semantic information until the participants obtained an accuracy level of 90% on this piece of semantic information (95% on the crucial target faces). The participants were then presented with the third piece of semantic information and were given the learning task until the 90% accuracy level was obtained for this piece of information (95% for the crucial target faces). To complete the learning phase, the participants were then presented with each face and asked about one of the pieces of semantic information. All participants achieved an accuracy of at least 95% at this point and so it was assumed that the identity had

indeed been established. Participants were allowed breaks every five minutes for 30 seconds. This whole procedure typically took one hour and twenty minutes to administer.

Following the learning phase, the experimental phase began. This was identical to the three phases used in Experiment 16, except that the response keys were C for face identity A and L for face identity B.

4.5.2. Results

Psychometric functions were fitted for each of the participants before and after adaptation (these are not shown for space limitations). The percentage shift of the PSE toward the adaptation identity for each category of adaptor is shown in Table 4.4. The data were subjected to a 6-level between-subjects ANOVA. This revealed a significant effect of type of adaptor, $F(4, 5) = 26.165$, MSE $= 4.425$, $p < .05$. Adaptation to the facial image was significantly greater than specific semantic information (mean difference = 8.599%, $p < .05$), semi-relevant semantic information (mean difference = 19.173%, $p < .05$), vaguely-relevant semantic information (mean difference = 18.742%, $p < .05$), imagination (mean difference = 17.536%, $p < .05$), and irrelevant information (mean difference = 16.203%, $p < .05$). The magnitude of adaptation was greater for specific-semantic information than semi-specific semantic information (mean difference = 10.574%, $p < .05$), vaguely-relevant semantic information (mean difference = 10.143%, $p < .05$), imagination (mean difference = 8.937%, $p < .05$), and irrelevant information (mean difference = 7.604%, $p < .05$). No other differences were significant after Tukey HSD *post hoc* tests. A series of one-sample t-tests were run on the data, revealing that adaptation due to image was greater than zero, t (11) = 47.366, $p < .05$, and adaptation

due to specific-semantic information was greater than zero, t (11) = 18.173, $p < .05$). No other one-sample t-tests were significant, largest t = 0.931, $p > .52$.

Table 4.4.

Percentage shift in PSE toward the adaptor identity for Experiment 17. Positive numbers indicate adaptation whereas negative numbers indicate priming. Significant effects are marked with an * for $p < .05$ and ** for $p < 0.01$ (assessed using t-tests).

Adaptor	PSE shift
Image	17.73%**
Name (Specific Semantic Information)	10.63%*
Semi-Specific Semantic Information	0.05%
Largely-Irrelevant Semantic Information	0.48%
Irrelevant (Control)	1.25%
Imagination	1.69%

4.5.3. Discussion

Identity thresholds were significantly adapted when the adaptor was the image or specific-semantic information but not non-specific semantic information. These results are broadly consistent with Experiment 16, in that adaptation to non-specific semantic information did not occur, whereas adaptation to specific information was found. Moreover, within-modal adaptation produced stronger adaptation than cross-modal adaptation. The one finding that was not replicated from Experiment 16 was the imagination finding. In Experiment 17, imagination did not cause significant adaptation,

which is consistent with Moradi et al.'s (2005) findings, and suggests that the reason for this is that faces only learnt in one viewpoint are not able to be visualised well enough for imagination.

Experiment 16 demonstrated that adaptation for famous faces could be caused by faces, names, voices, associated people, caricatures, and imagination, but not non-specific semantic information. Experiment 17 found that adaptation for recently learnt faces could be caused by faces and names, but not non-specific semantic information nor imagination. Thus, adaptation to familiar and unfamiliar faces is similar. Faces that are familiar; however, can be imagined better. This conclusion is consistent with both behavioural (Jiang, et al., 2007; Burton, et al., 1999) and neurological (Eger, Schweinberger, Dolan, and Henson, 2005; Pourois, et al., 2005) findings that familiar faces have more stable, viewpoint-independent representations.

The finding that natural caricatures produce larger adaptation effects than veridical images is somewhat counter-intuitive. One explanation for the advantage of caricatures in recognition (e.g., Rhodes, Brennan, and Carey, 1987) is based on the idea that memory for faces is not veridical. Memory for faces may exaggerate the most distinguishing features (e.g., Harvey, 1986). Caricatures may thus be a more veridical image of the representation of a face stored in memory, especially if computer generated. The present study used natural caricatures that may not be quite so close to the stored representations of faces. Thus, it may be that caricatures also contain identity-specific semantic information that can add to the visual adaptation effects. In this way, there is a visually coded and a semantic coded adaptation occurring. These two types of adaptation add together to form a greater adaptation effect.

4.6. Experiment 18 – FIAEs are Mediated by Type of Processing

Face recognition is characterised by an expert processing mechanism that relies on the configuration of two eyes above a nose above a mouth. This configuration is disrupted in an inverted face (Yin, 1969), making it harder to recognise. Photographic negation (reversed contrast polarity) of a facial image causes it to be recognised less accurately, but does not alter the type of processing engaged (e.g., Galper, 1970; Phillips, 1972). Negative images are generally associated with more error in encoding rather than a change in processing (e.g., George, Dolan, Fink, Baylis, Russell, and Driver, 1999; Valentine, 1991).

Few studies have looked at the effect type of processing has on the magnitude of the FIAE. In a comprehensive study of types of changes from the adaptor to the test stimuli in FDAEs, Yamashita, et al. (2005) demonstrated that the after-effects transferred across stimuli differing in size, colour, contrast, but not contrast polarity or spatial frequency. Rhodes, Jeffery, Watson, Jacquet, Winkler, and Clifford (2004) demonstrated that it was possible to show a contraction FDAE for upright faces and expansion FDAE for inverted faces at the same time, arguing that the processing of upright and inverted faces relies on different neuronal populations. Indeed, Mazard, Schiltz, and Rossion (2006) have shown that recovery from adaptation in the fusiform gyrus is larger for upright faces than inverted faces.

Experiment 18 aimed to examine how the FIAE is affected by configural processing. Eight different adaptation stimuli were used comparing the effects of same and different adaptor image from that used at test, whilst also comparing the effects of inversion and negation of the magnitude of adaptation. Two hypotheses can be made regarding this study. Image based adaptation may occur, whereby after effects are greater

when the adaptor and test stimuli are matched, regardless of what the adaptor is. However, if the FIAE is based on some form of face-specific[16] coding mechanism, then it is likely to be observed for upright rather than inverted faces.

4.5.1. Method

4.5.1.1. Participants

Thirty-two Cardiff University students undertook this experiment as partial fulfilment of a course requirement. All participants had normal or corrected vision. All were White British nationals who were familiar with George Bush and Tony Blair.

4.5.1.2. Materials

The same images of George Bush and Tony Blair used in Experiment 16 were employed here. These two sets of morphs were inverted into two additional sets. Two negated sets were also created using Adobe Photoshop™ image manipulation software. These negated sets were subsequently inverted to create the final two sets. The 50% image of each type of stimulus is presented in Figure 4.9. All stimuli were presented using SuperlabPro 2™ Research Software on an RM PC.

---

[16] In this sense, face-specific reflects the expert nature of face recognition and configural encoding. I am not suggesting that face recognition is special.

a.   b.

c.   d.
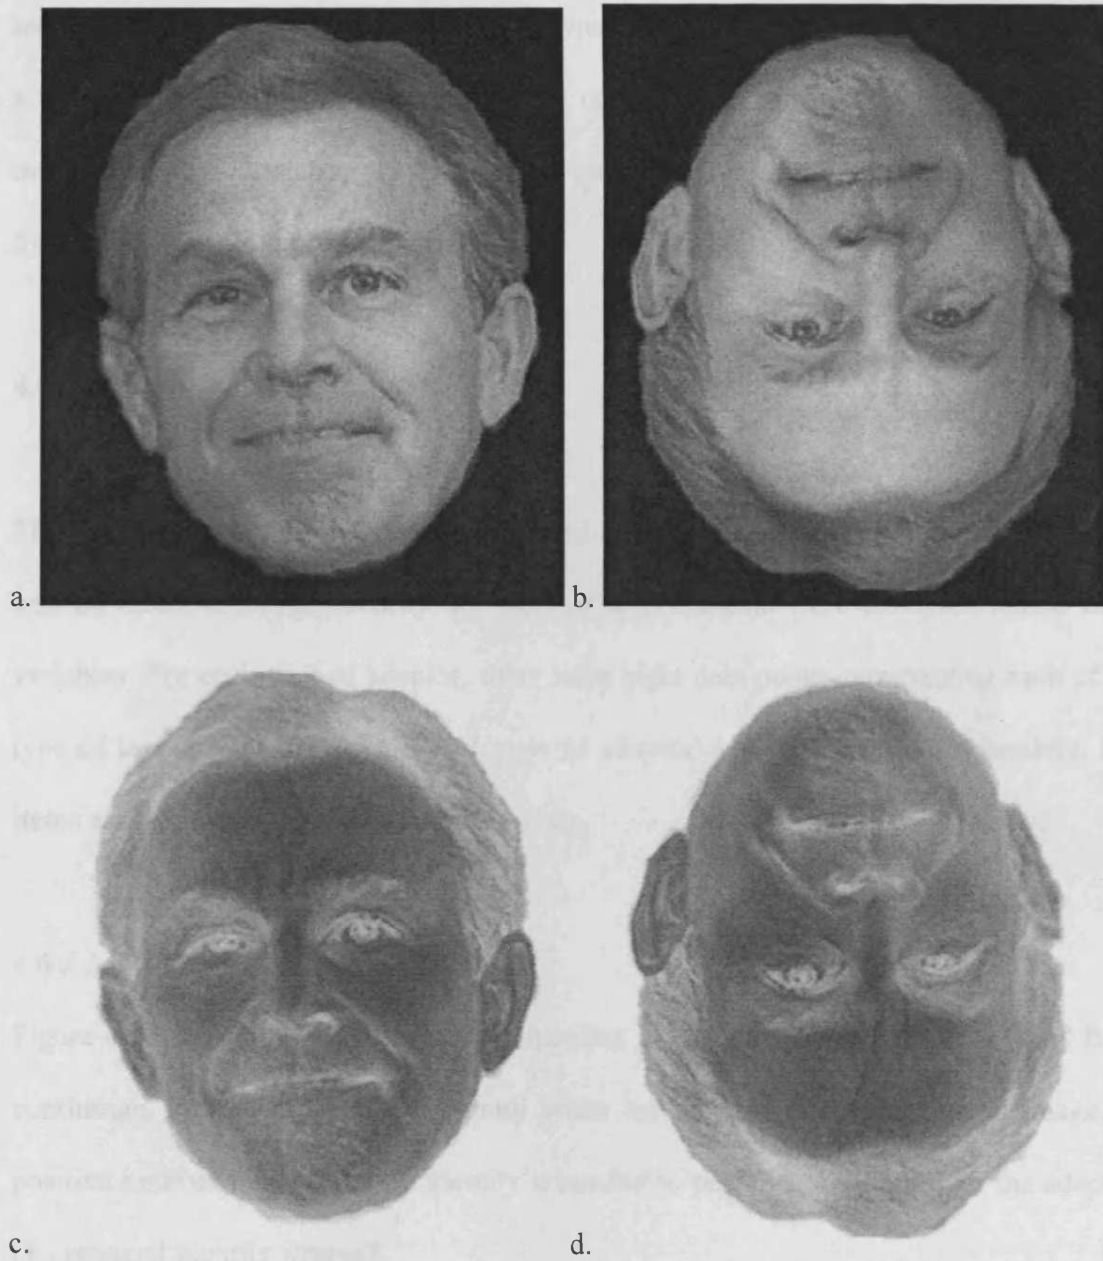
Figure 4.9. Examples of the stimuli used in this Experiment. a. The unaltered 50% midpoint. b. The inverted 50% midpoint. c. The negated 50% midpoint. d. The inverted and negated 50% midpoint.

4.6.1.3. Design and Procedure

The adaptor was manipulated between-subjects with four levels (same image, different image, negated image, or inverted image). A within-subjects manipulation was also

implemented, whereby participants saw 8 types of test faces: 2 (same or different image) x 2 (inverted or upright) x 2 (negated or control). The magnitude of adaptation was measured as the change in the PSE pre- to post-adaptation. The procedure was identical to the image conditions in Experiment 16.

4.6.2. Results

The magnitude of adaptation was calculated in the same way as Experiment 16. There was no effect of image identity nor pair, as such the data were collapsed across these variables. For each type of adaptor, there were eight data points representing each of the type of test stimuli. Data from each type of adaptor was first analysed separately. By-items analyses were conducted on these data.

4.6.2.1. Unaltered Adaptor

Figure 4.10 shows the mean percentage increase in PSE in the George Bush – Tony Blair continuum for each of the test stimuli when the adaptor was an unaltered image. A positive number indicates more identity is needed to perceive the identity of the adaptor, i.e., reduced identity strength.

Figure 4.10. Mean PSE shift of identity needed to respond with that identity post-adaptation. Darker bars represent upright test stimuli, lighter bars represent inverted test stimuli. Error bars represent standard error.

Figure 4.10 indicates that greater adaptation occurred when the test stimuli were upright. The data were subjected to a 2 by 2 by 2 Within-subjects ANOVA. This revealed that greater adaptation occurred when the same image was used for both adaptation and test, $F(1, 3) = 28.475$, MSE = 0.905, $p < .05$. Greater adaptation was observed for upright test stimuli than inverted test stimuli, $F(1, 3) = 713.538$, MSE = 1.761, $p < .05$. There

were no significant differences in the magnitude of adaptation for negated test stimuli, $F(1, 3) = 1.473$, MSE = 7.534, $p > .1$. There was a significant interaction between image and negation, $F(1, 3) = 440.437$, MSE = .31, $p < .05$, revealing that greater adaptation was found for same view unaltered test stimuli than same view negated stimuli (mean difference = 5.312, $p < .05$) and different negated test stimuli than different unaltered stimuli (mean difference = 2.956, $p > .05$). There was also an interaction between negation and inversion, $F(1, 3) = 29.356$, MSE = 1.043, $p < .05$. Simple effects showed that the effect of inversion was larger for unaltered stimuli (mean difference = 14.490, $p < .05$) than for negated stimuli (mean difference = 10.577, $p < .05$). Finally, there was a three-way interaction, $F(1, 3) = 71.673$, MSE = 1.297, $p < .05$.

4.6.2.2. Negated Adaptor

A parallel analysis was run for when the adaptor was a negated image. The results are presented in Figure 4.11. This revealed a significant effect of image, whereby greater adaptation was observed when the same image was used at adaptation and test than when a different image was used, $F(1, 3) = 123.650$, MSE = .859, $p < .05$. There was also a significant effect of inversion, whereby greater adaptation was observed when the test stimuli were upright, $F(1, 3) = 150.024$, MSE = 7.978, $p < .05$. There was a significant interaction between image and inversion, $F(1, 3) = 53.067$, MSE = 2.197, $p < .05$. Simple main effects showed that the magnitude of adaptation was stronger for negated images than unadjusted images when the same image was used as the adaptor as those that made up the test morph continua (mean difference = 2.628, $p < .05$), whereas the magnitude of adaptation was stronger for unadjusted images than negated images when a different image was used as the adaptor to that at test (mean difference = 2.787, $p < .05$). There was also an interaction between negation and inversion, $F(1, 3) = 12.051$, MSE = 1.570,

$p < .05$, which revealed itself in a greater magnitude of adaptation for negated upright stimuli than inverted stimuli (mean difference = 13.465, $p < .05$) which was greater than when the stimuli were unadjusted (mean difference = 10.694, $p < .05$).
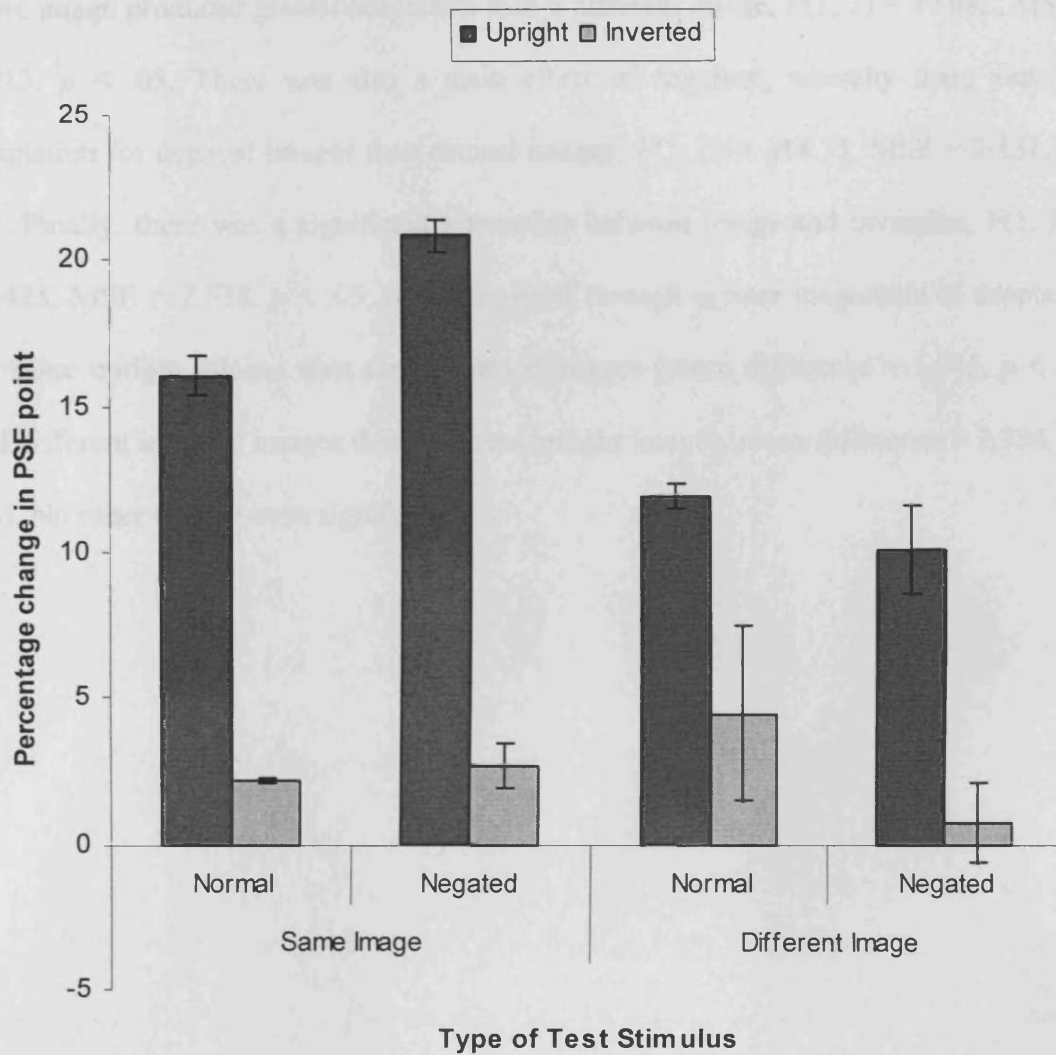
Figure 4.11. Mean perceptual midpoint shift of identity needed to respond with that identity post-adaptation. Darker bars represent upright test stimuli, lighter bars represent inverted test stimuli. Error bars represent standard error.

4.6.2.3. Inverted Adaptor

A further parallel analysis was run on the data when the adaptor was inverted. These data are presented in Figure 4.12. This revealed a significant effect of image, whereby the same image produced greater adaptation than a different image, $F(1, 3) = 49.682$, MSE = 8.913, $p < .05$. There was also a main effect of negation, whereby there was less adaptation for negated images than control images, $F(1, 3) = 314.35$, MSE = 2.131, $p < .05$. Finally, there was a significant interaction between image and inversion, $F(1, 3) = 14.425$, MSE = 2.538, $p < .05$, revealing itself through greater magnitude of adaptation for same upright images than same inverted images (mean difference = 1.945, $p < .05$) and different inverted images than different upright images (mean difference = 2.334, $p < .05$). No other effects were significant.
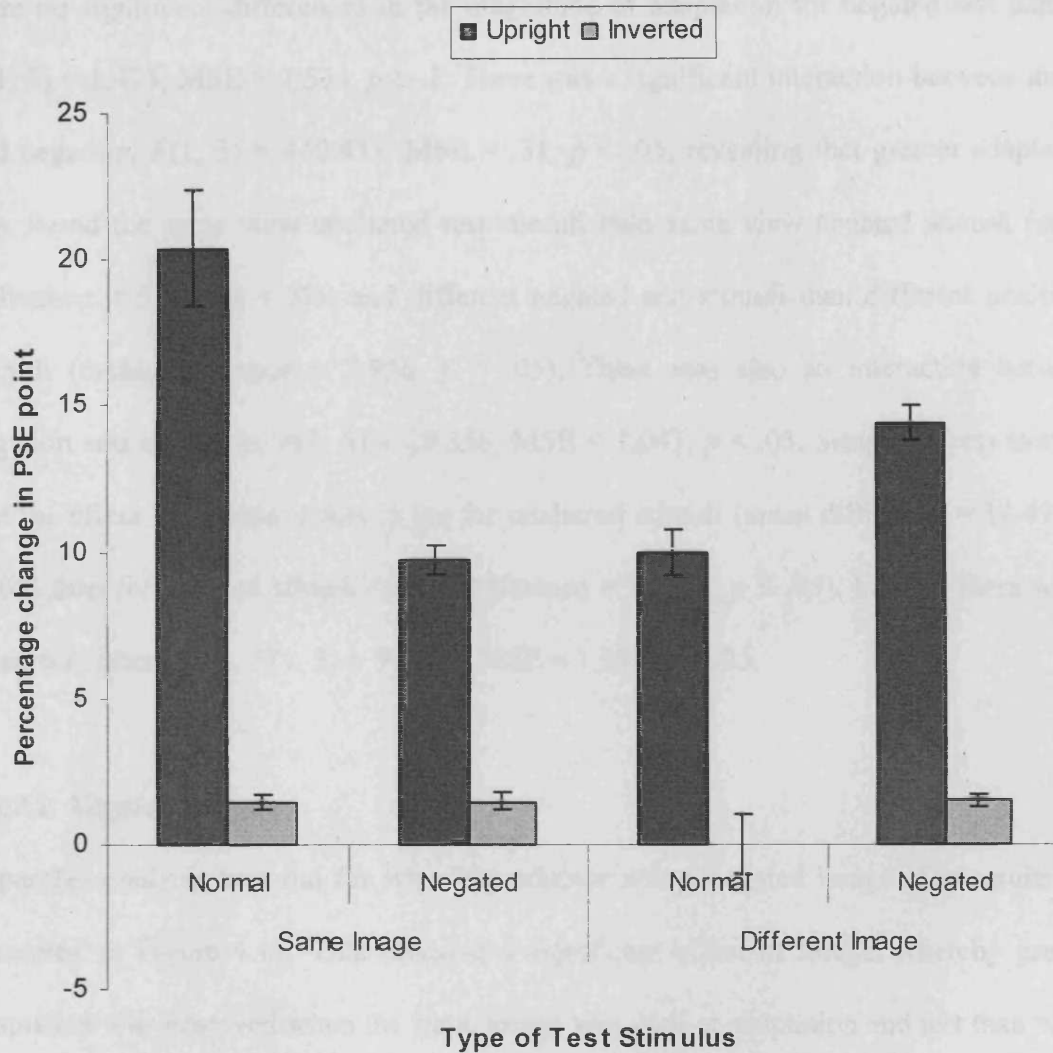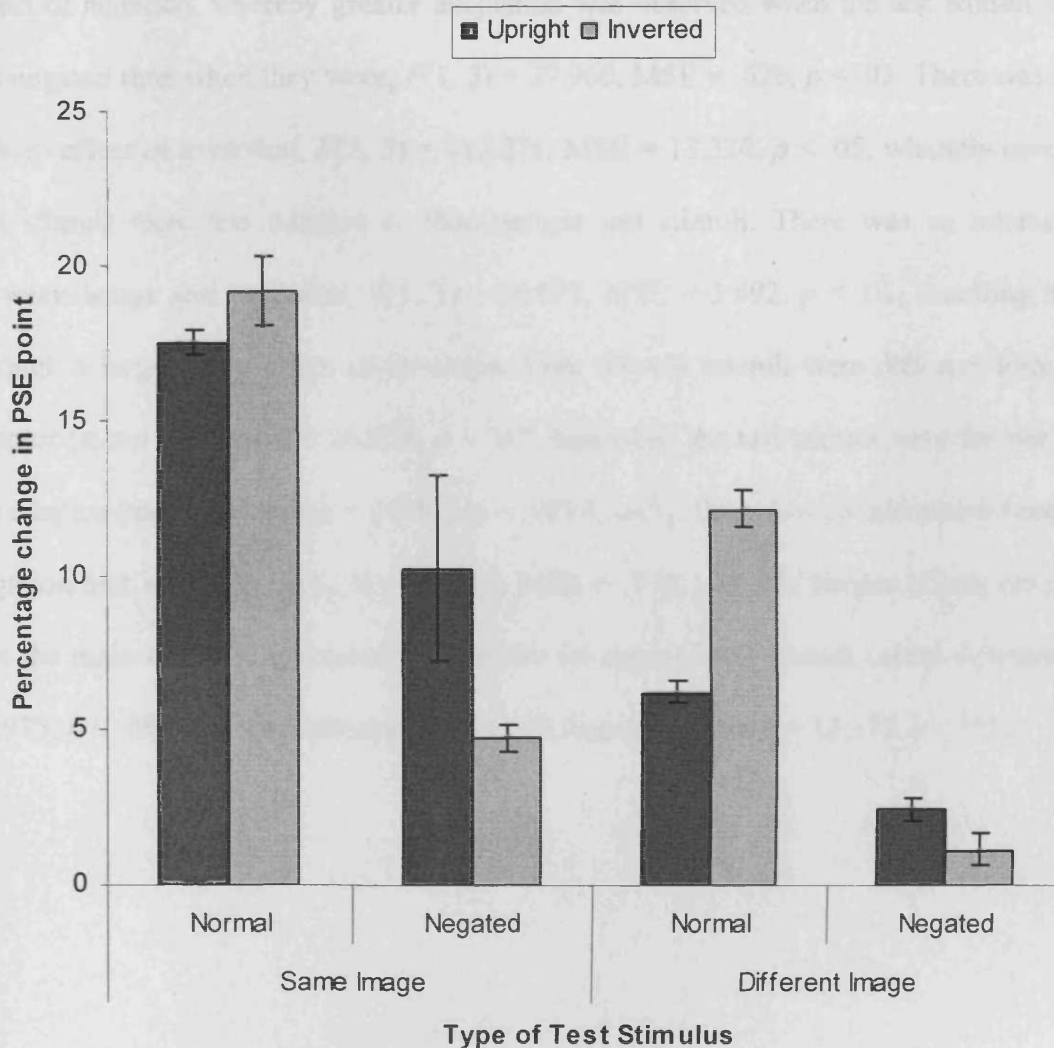
Figure 4.12. Mean perceptual midpoint shift of identity needed to respond with that identity post-adaptation. Darker bars represent upright test stimuli, lighter bars represent inverted test stimuli. Error bars represent standard error.

4.6.2.4. Negated and Inverted Adaptor

A fourth analysis was run on the data for when the adaptor was both inverted and negated, as shown in Figure 4.13. This revealed a significant effect of image, whereby greater adaptation was observed when the adaptation and test stimuli matched than when they were different, $F(1, 3) = 160.032$, MSE $= 1.019$, $p < .05$. There was a significant

effect of negation, whereby greater adaptation was observed when the test stimuli were not negated than when they were, $F(1, 3) = 27.960$, MSE = .626, $p < .05$. There was also a main effect of inversion, $F(1, 3) = 112.271$, MSE = 13.326, $p < .05$, whereby inverted test stimuli were less adapted to than upright test stimuli. There was an interaction between image and inversion, $F(1, 3) = 24.677$, MSE = 3.492, $p < .05$, revealing itself through a larger main effect of inversion when the test stimuli were different from the adaptor (mean difference = 16.958, $p < .05$) than when the test stimuli were the same as the adaptor (mean difference = 10.392, $p < .05$). Finally, there was an interaction between negation and inversion, $F(1, 3) = 16.915$, MSE = .799, $p < .05$. Simple effects revealed that the main effect of inversion was greater for negated test stimuli (mean difference = 14.975, $p < .05$) than for unaltered test stimuli (mean difference = 12.375, $p < .05$).
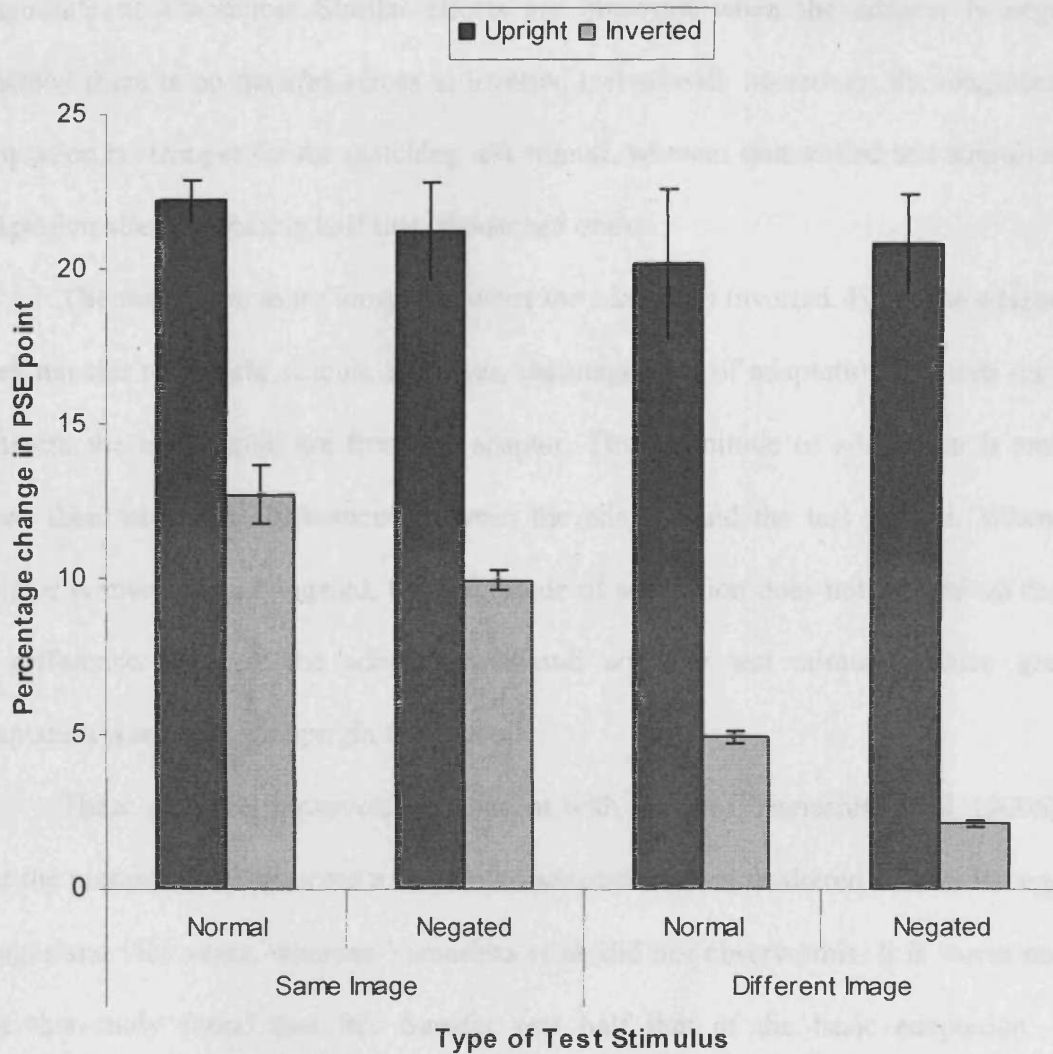
Figure 4.13. Mean perceptual midpoint shift of identity needed to respond with that identity post-adaptation. Darker bars represent upright test stimuli, lighter bars represent inverted test stimuli. Error bars represent standard error.

4.6.3. Discussion

Adaptation to an upright stimulus does not transfer to inverted test stimuli. Moreover, the magnitude of adaptation when the adaptor and test stimuli match is twice that of when they do not match, but negation and pose changes do not add to the detriment in the

210

magnitude of adaptation. Similar effects are observed when the adaptor is negated, whereby there is no transfer across to inverted test stimuli. Moreover, the magnitude of adaptation is stronger for the matching test stimuli, whereas unmatched test stimuli show adaptation strengths nearly half that of matched ones.

The results are more intriguing when the adaptor is inverted. Here, the adaptation does transfer to upright stimuli. However, the magnitude of adaptation depends on how different the test stimuli are from the adaptor. The magnitude of adaptation is smaller when there are more differences between the adaptor and the test stimuli. When the adaptor is inverted and negated, the magnitude of adaptation does not depend on degree of difference between the adaptation stimuli and the test stimulus, since greater adaptation was noted for upright test stimuli.

These data are, however, inconsistent with those of Yamashita et al. (2006), in that the present study observed a transfer of adaptation from unaltered stimuli to negated images and vice versa, whereas Yamashita et al. did not observe this. It is worth noting that this study found that this transfer was half that of the basic adaptation. The differences between these two studies are that the present study tested the FIAE in familiar faces, whereas Yamashita et al. observed the FDAE in unfamiliar faces. Thus, it is possible that the differences observed are due to the processing of familiar and unfamiliar faces. The one difference that is, at present, difficult to explain is that Yamashita et al. observed a greater transfer from an unaltered adaptor to negated test stimuli, whereas the present study observed the reversed pattern of results. Unfortunately, no explanation for this can be offered.

These data seem to suggest two important facets of the FIAE. First, there is some image-based adaptation that is occurring, which is lower-level and probably based in the striate cortex (c.f., Hurlbert, 2001). Some of the FIAEs will simply be due to this

allowing for differences between stimuli to be better detected. This conclusion is based on the fact that stimuli that are matched at adaptation and at test produce stronger FIAEs than unmatched stimuli. Second, there are also some facets of the FIAEs that are likely to be unique to face recognition, since the FIAE is based in part on expert processing. As such FIAEs represent a unique class of high-level shape after effect due to the expert use of configural processing involved in face processing.

4.7. Experiment 19 – The Duration of the FIAE

Rhodes et al. (2003) have suggested that the FDAE causes a recalibration of the face-space, such that incoming faces can be encoded according to recent experience. Recent evidence has indicated that face after-effects are rather more long-lived than Rhodes et al. suggest. Carbon and Leder (2006) tested the FDAE in the Mona Lisa and demonstrated that the after-effects were observed for 80 minutes after adaptation. Carbon and Leder (2005) demonstrated that the FIAE was still present in highly familiar faces (they tested Lady Dianna Spencer's face) after a 5 minute break between exposure and test. This suggests that the FIAE may cause permanent changes in the representation of highly familiar faces.

The study proposed here is one that aimed to test whether FIAE is short lasting or longer lasting. If this effect is a storage effect, then it is likely to last for a long time, whereas if it is a recalibration, it should not last more than a few minutes. Indeed, it is possible that through adaptation, the identity regions in face-space are shrunk in some way. Assuming a Voronoi model of face-space (c.f., Lewis and Johnston, 1998; 1999), adapting to one identity may warp the identity region in such a way as to make it smaller,

thus allowing for other faces to occupy part of the region that was once belonging to that face.

4.7.1 Method

4.7.1.1. Participants

Four participants volunteered for this experiment. All had normal or corrected vision and three were blind to the hypothesis of the Experiment. All were familiar with the identities used in this Experiment.

4.7.1.2. Materials

The images used in Experiment 13 were used here. They were morphed together using Smartmorph™ Software using 200 points, creating 12 images, increasing in steps of 8.3%. These were presented to the participants using Superlab Pro 2™ Research Software using an RM PC.

4.7.1.3. Design and Procedure

Participants undertook the standard adaptation procedure at Time 1 and then again at Time 2 (two months later). This procedure is described in Experiment 15. Adaptation was measured by the PSE in the morph continua pre- and post-adaptation. Each participant was adapted to a different identity and different image. Each participant saw the test morph continua containing images made up of the adaptor they had seen and the different image of the identity. Thus, a 2 by 2 by 2 within-subjects design was employed, whereby the independent variables were: type of adaptor face (same as and different to that used in the test morph continua); adaptation (pre- and post-adaptation); and Time (Time 1 and

Time 2 – 2 months later). Presentation order of the faces was randomised, and the participants were randomly allocated to be adapted to one of the adaptation stimuli, such that each participant was adapted to a different stimulus.

4.7.2. Results

Psychophysical functions were fitted for all participants, before and after adaptation at Time 1 and at Time 2. The PSE assessed pre- and post-adaptation at Time 1 and Time 2 are presented in Figure 4.14. This shows an increase in PSE over time. Interestingly, the data appear to show a linear progression when different images were used as the adaptor to those in the morph test continua, whereas this progression is not linear when the same image is used.

The data summarised in Figure 4.14 were subjected to a 2 by 2 by 2 within-subjects ANOVA with the factors: type of adaptor face (same as or different to that used in the test morphs), time (Time 1 or Time 2), and adaptation (pre- or post-adaptation). This revealed a significant effect of time, $F(1, 3) = 69.699$, MSE $= 22.689$, $p < .05$, whereby the PSE was higher at Time 2 than Time 1 (mean difference $= 14.060\%$). There was a main effect of adaptor, $F(1, 3) = 615.749$, MSE $= 3.763$, $p < .05$, whereby post-adaptation PSEs were greater than pre-adaptation PSEs (mean difference $= 17.018\%$). The interaction between face and time was significant, $F(1, 3) = 11.099$, MSE $= 3.378$, $p < .05$. Simple effects revealed that the difference between PSEs at Time 1 and Time 2 for same adaptor face as test morph continua faces was less than for different adaptor face to test faces (compare mean difference $= 11.895\%$ to 16.225%). The interaction between face and adaptation was significant, $F(1, 3) = 26.530$, MSE $= 8.203$, $p < .05$. Simple effects revealed that the change in PSE pre- to post-adaptation was greater for same

adaptor face as that used in the test morph continua than for different adaptor face as that used in the test morph continua (compare mean difference = 22.213% to 11.801%, respectively). Finally, the 3-way interaction was approaching significance, $F(1, 3) = 7.101$, MSE = 1.648, $p < .08$. This interaction was explored and revealed that the interaction was borne out due to the PSE being greater for different adaptor face than same adaptor face as that used in the test morph continua only at Time 2 pre-adaptation (mean difference = 7.819%, $p < .05$). At all other times, the PSE was greater for when the same image was used as the adaptor and in the test morphs (all ps $< .05$), except for at Time 1 pre-adaptation when the difference was not significant (mean difference = 1.070%, $p > .4$). The baseline PSE at time 2 was significantly greater than the baseline PSE at time 1 (mean difference = 14.993, $p < .05$). The interaction been time and adaptation was not significant, F (1, 3) = 1.313, MSE = 5.312, $p > .33$.
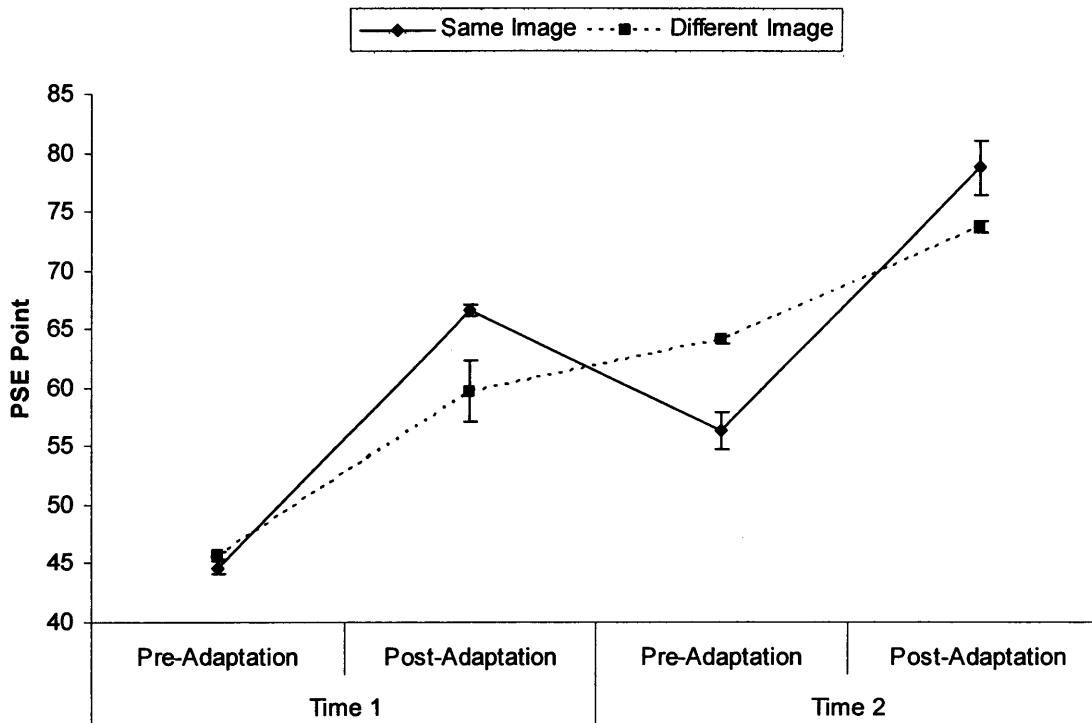
Figure 4.14. Mean PSE pre- and post-adaptation at Time 1 and Time 2 for adaptation to the same image as that used in the test morph continua and for adaptation to the different image as that used in the test morph continua.

4.7.3. Discussion

Two months after the initial adaptation experiment, two participants demonstrated an increase in their perceived midpoint of the continua between Pierce Brosnan and Harrison Ford. The baseline at Time 2 was not significantly different from the test phase at Time 1. This indicates some level of permanent change in the identity regions of the faces tested. The magnitude of adaptation was the same at both Time 1 and Time 2.

These results indicate that the adaptation effect is not an effect of learning to complete the task[17]. If it were, then an increase in the magnitude of adaptation at Time 2

---

[17] This is not to say that adaptation cannot be a mechanism of learning about faces.

would be expected. Instead, adaptation is of the same magnitude, indicating some permanent change in the perception of the facial identity in question. In other words, the identity region for a particular identity has been substantially reduced in size due to the adaptation effects. Alternatively, the identity region for Harrison Ford has been increased substantially to include what was once an unknown face.

It seems difficult to distinguish between whether the effect is one of shrinking the identity region or increasing the size of the opposing region. Both appear possible. In the 2-alternative forced-choice paradigm, participants are giving a name to a morphed face. In other words, they are attributing an identity to a morph. This may mean that subsequently they will always see that blend as the identity that they had given. This argument implies an increase in size of the identity region not being adapted to. This certainly seems most plausible currently.

There is also an issue with the possibly permanent change in the identity regions. It has been noted that simple visual after effects can last for hours (Blakemore, et al., 1970; Magnussen and Greenlee, 1985) after protracted adaptation, or even days (Harris, 1980). The theory of such permanent changes in perception is based on some type of long-term synaptic change. Long-term potentiation seems plausible to account for changes in neuronal substrates (T. Brown, Chapman, Kairiss, and Keenan, 1988) but may not reach permanence (Gustafsson and Wigstrom, 1988). Indeed, long-term potentiation has been found in the visual cortex of both cats (Komatsu, Fujii, Maeda, Sakaguchi, and Toyama, 1988) and rats (Artole, Brocher, and Singer, 1990; Berry, Teyler, and Taizhen, 1989).

The evidence on long-term potentiation gives a plausible account for the changes occurring in the identity region of the faces tested. The face-space is being recalibrated due to some degree of long-term potentiation. This can be short lived or longer depending

on the length of adaptation (e.g., Albrecht, Farrar, and Hamilton, 1984), and the type of methodology employed. Whether the change is permanent or just long-term is up for debate. Nevertheless, identity adaptation appears to last for some time.

4.8. General Discussion

Experiments 13 to 19 demonstrated that the FIAE could be brought about in participants following adaptation to non-visual identity-specific information in a forced choice paradigm (Experiment 13), free response procedure (Experiment 14), or a psychophysical procedure (Experiments 15 to 19). These results indicate that the way faces are encoded into face-space can be affected by non-visual information. In this way, the face-space can be recalibrated without visual stimulation, indicating the importance of cognitive mechanisms in face recognition.

The findings from Experiment 15 suggest that the FIAE is mediated and moderated by visualisation abilities. Participants better able to visualise show stronger effects of adaptation. Moreover, visualising a face is sufficient to cause adaptation (but only for familiar faces, Experiment 17). This implies that the face-space can be activated through visual imagery. The lack of adaptation to unfamiliar faces and to imagined distortions (c.f., Moradi et al., 2005), indicates that dimensions of face-space cannot be imagined, but identities within it can be. This indicates that the face-space is an unconscious mechanism for the encoding of faces.

Experiment 16 noted that adaptation to a caricatured image produced a greater magnitude in PSE shift than adaptation to veridical images. Interpreting this result in the face-space model indicates that the dimensions of face-space may not be veridical. Alternatively, the identity region that a face occupies covers the veridical and the

caricatured face. Indeed, the caricatured face is further away from the nearest exemplars than the veridical image. As such, adaptation to this causes greater change in the dimensions of the face-space associated with that particular face.

Experiment 18 indicated that the FIAE is mediated by expert processing, whereby the effects of adapting to an upright face do not transfer to inverted test stimuli. However, the converse does occur. Interpreting this within the face-space metaphor suggests that adaptation can transfer from some dimensions to others, but not all. Negation does not alter the dimensions that faces will be encoded on, thus they will act like upright faces in terms of adaptation (and show image-based adaptation as well as FIAE). This suggests that inverted faces are stored along other dimensions in face-space than upright faces. To explain why the transfer from inverted faces to upright faces does occur, it is likely that this is due to the non-visual identity adaptation that has been shown in previous experiments, since identity can be obtained from inverted faces after prolonged exposure.

Finally, Experiment 19 suggested that the FIAE for familiar faces is a relatively permanent change, lasting at least two months. This finding is important, for it may suggest how error regions surrounding faces in face-space are reduced with prolonged exposure. The neural tunings become increasingly finely tuned to a particular identity after adaptation. This tuning is possibly due to the neural system containing the minimal amount of "fatigued" neurons possible after adaptation. The system allows the neural populations less associated with the particular identity to become less activated by the identity and thus less likely to suffer adaptation. It may be that repeated exposure causes adaptation and in turn this aids learning of a particular face in a process that is akin to the idea of stimulus discrimination presented in Chapter 1.

The neural model can be used to account for learning about faces and the permanent changes in face perception following adaptation. Following adaptation to a

particular identity, the neural population will respond less as presented in Figure 1.3. It is implausible that the neural pools remain fatigued for two months. Thus, to explain this effect, one must return to the idea that to represent identity, many neural pools are involved. These pools will not remain fatigued for two months, for it would cause observable after effects in other faces. Instead one can look at stimulus discrimination and the neural responses associated with it. Figure 4.15 presents a possible change in neural tuning curves during stimulus presentation. This neural change occurs such that the neural pools will fire for a smaller range of stimuli values. An alternative presentation of this is shown in Figure 4.16, which shows that the region that the neural pool for a particular face will respond to a smaller range of stimuli (similar faces) following repeated presentation.
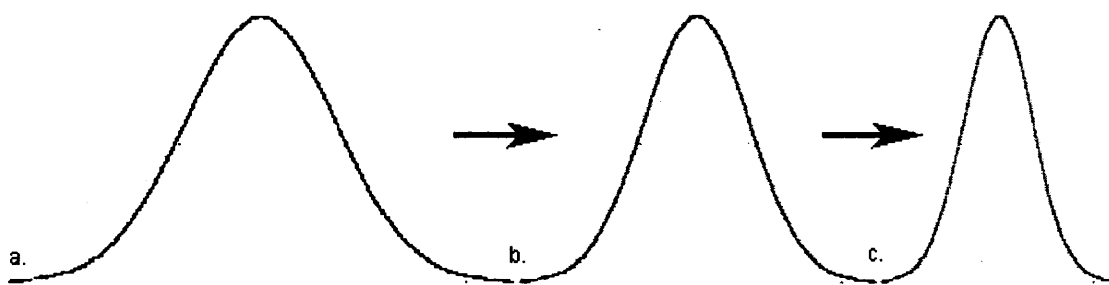
Figure 4.15. The possible development of neural tunings to particular facial identities. The neural tuning curve is initially broad as in a. With prolonged exposure, the range of stimuli that the neural pools that will respond to become narrower as in b. Eventually, only a very small range of stimuli will cause the neural pool to respond as in c.

There is one caveat with this discussion. A large alteration in the perception of faces is altered by a single 60 second exposure. However, the experience people have with such faces is likely to be much greater than this exposure duration. For example, in the average James Bond film, Pierce Brosnan will be on screen for roughly 30 minutes. The explanation offered would thus suggest that following watching a Bond film people's

perception of Pierce Brosnan will be altered significantly and to a greater degree than in the experimental sessions observed here. Thus, the tuning of the relevant tunings might already have reached a maximum before this experimental session. Perhaps the fixation during this experimental session causes neural tuning whereas in a film, Pierce Brosnan will be constantly moving and interacting with other people. Thus, other identities will be activating at the same time. The experimental sessions may thus cause an artificial state whereby people are forced to solely activate one identity rather than many.
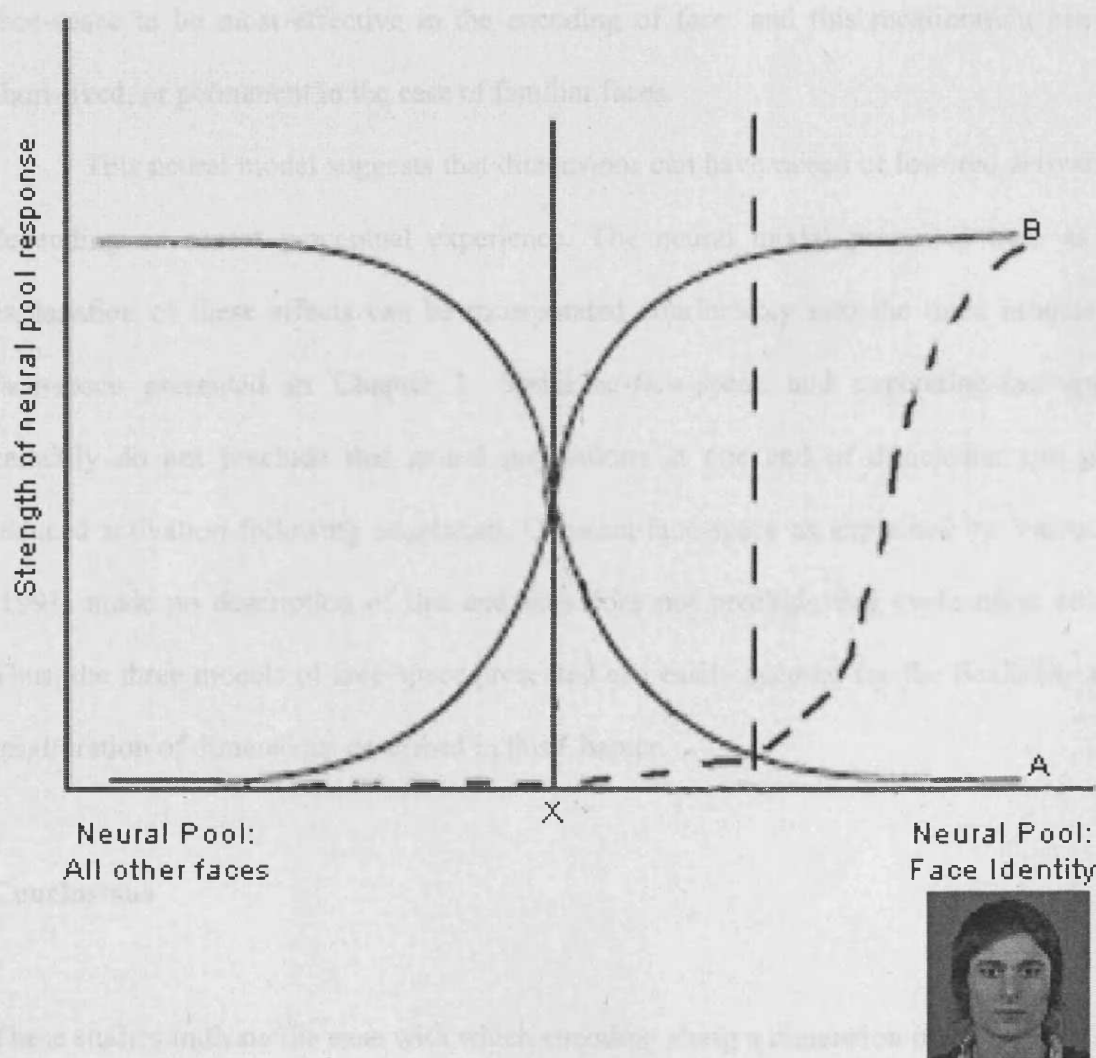
Figure 4.16. The solid lines represent the neural responses early in learning. With prolonged exposure, the neural response for facial identity B becomes the dashed line, and the new perceptual midpoint shifts to the right.

221

Interpreting these results in face-space, one can suggest that the neural responses distributions along each dimension in face-space gradually become tuned to selectively respond to the most encountered faces. Neural pools associated with a particular identity can be activated by non-visual information and this in turn activates the areas of the dimensions that are associated with that particular identity. Given time, the neural responses become selective to a particular identity. As such, adaptation recalibrates the face-space to be most effective in the encoding of face, and this recalibration can be short-lived, or permanent in the case of familiar faces.

This neural model suggests that dimensions can have raised or lowered activation depending on recent perceptual experience. The neural model presented here as an explanation of these effects can be incorporated comfortably into the three models of face-space presented in Chapter 1. Shrinking-face-space and expanding-face-space certainly do not preclude that neural populations at one end of dimension can give reduced activation following adaptation. Constant-face-space as explained by Valentine (1991) made no description of this and thus does not preclude this explanation either. Thus, the three models of face-space presented can easily account for the flexibility and recalibration of dimensions described in this Chapter.

Conclusions

These studies indicate the ease with which encoding along a dimension of face-space can be affected by recent perceptual experience. Perceptual experience recalibrates the face-space according the most recently activated neural pool – this can be at an extreme end of a particular dimension or could represent a particular facial identity. The results also

demonstrate the cross-modal nature of face perception. The final experimental chapter

aimed to explore the effect of Navon processing on face recognition.

Chapter 5 – Configural and Featural Processing

The preceding chapters have indicated that the face-space is a more flexible system than previously hypothesised. As an aside, but looking at alternative methods for adjusting the parameters of face recognition, the present chapter explores the effect Navon stimuli have on face recognition. In Experiment 20, the materials and basic effect that Navon stimuli have on face recognition are tested in a line-up paradigm. Experiment 21 tested this effect using a set of shape Navon stimuli. Experiment 22 explored an alternative to Navon letters in inducing the hypothesised shift in processing style. In Experiments 23 and 24, the effect Navon stimuli have on face recognition is assessed in a recognition paradigm. Experiments 25 and 26 assessed the effect adjusting Navon stimuli (blurring in Experiment 25 and reducing their contrast in Experiment 26) has on their effect on recognition of upright and inverted faces. Experiment 27 assessed performance in a face recognition task following processing different sized letters. Experiments 28 and 29 tested the effect Navon stimuli have on bandpass filtered faces. The findings from these experiments are consistent with a perceptual after effect account of the effect Navon stimuli have on face recognition.

5.1. Experiment 20 – Local Processing of Navon Letters Causes a Face Recognition Deficit

Processing the local features of Navon letters appears to cause a deficit in facial identification (e.g., Macrae & H. Lewis, 2002). While this effect is not always replicated (Brand, 2005; Lawson, 2006; Ryan, et al., 2006), it is theoretically important. Perfect (2003) has explained this effect in terms of a transfer-inappropriate processing shift as explained in Section 1.2.2. The research conducted on the effect processing Navon letters has on face recognition also suffers from a validity issue. Many researchers are using copies of copies of stimuli in their experiments. The video of the mock crime used to assess the effect of Navon stimuli on face recognition has been described as 'grainy' (Perfect, personal communication). As such, new stimuli are developed here to ensure generalisability of the results in addition to a refusal by many researchers to share their stimuli. Due to the unreliability of the detrimental effect of processing the local features of Navon stimuli on subsequent face recognition, the first task was to replicate this basic effect and create a baseline with which to compare further findings to.

5.1.1. Method

5.1.1.1. Participants

One-hundred-and-twenty undergraduates from the School of Psychology at Cardiff University between 18 and 30 years old with normal vision, as defined by self report, took part in this study for course credits. They were randomly divided into one of three conditions: local, control, global.

5.1.1.2. Materials

A 30 second video clip was made of a football team scoring a goal. The football team was a five-a-side team which was based in a nearby town, such that participants were unlikely to know those in the video. Each of the football team had the same frontal view photograph, all wearing the same kit, presented in a line-up. There were 10 photographs in the line-up all photographed in the same team strip. The photographs were each 70 mm by 50 mm with a resolution of 72 dpi. The line-up and video were pretested to assess functional size (number of mock witnesses divided by correct guesses, see Lindsay, Smith, and Pryke, 1999) of the line-up: Four postgraduate students gave a short description of the goal scorer with the photograph in front of them. Subsequently, 36 undergraduate participants were given all four verbal descriptions of the goal scorer and asked to pick him out. The functional size was 8, giving an accuracy of 12.5%. The video was presented on a PC using Windows™ Media Player and the line-up was presented using QuickTime™ Image Viewer. The video was presented in 720 pixel x 560 pixel dimensions and high resolution.

The set of 125 Navon letters produced by Brand (2005) were used[18]. The Navon stimuli were 91 mm by 47 mm. They have been used in several previous experiments and have been shown to produce reliable Navon effects. The Navon stimuli were presented on an RM PC using SuperlabPro 2™ Research Software.

5.1.1.3. Design

A three-level between-subjects design was employed, whereby participants were randomly divided into three groups and given the task of identifying the global Navon figure, the local Navon figure, or a control condition. The order of presentation of the

---

[18] I would would like to thank Andrew Brand for providing these stimuli.

Navon stimuli was randomised. The dependent measure was the accuracy (correct or incorrect) of the participants' response to the line-up of faces.

5.1.1.4. Procedure

The experiment had three phases: presentation, Navon, and identification. In the presentation phase, the participants were shown the video of the goal being scored. They were asked simply to watch the screen with no further instructions. This video lasted 30 seconds.

The participants were then introduced to Navon letters with three examples provided on screen. Participants were either instructed to identify the global Navon letter by saying it out loud, or instructed to identify the local letter, again by saying it out loud. In the control condition, participants were asked to read a book on cognitive psychology for the duration. After 5 minutes of processing either the local or global letters, participants in the experimental conditions swapped to the opposite Navon task and processed these critical letters for 5 further minutes. The latter 5 minutes designated which condition the participants fell in. The first five minutes controlled for difficulty (see Perfect, 2003) [19].

After five minutes of this Navon identification task, participants were asked to identify the person who scored the goal in the video by giving the number which corresponded to the face. All ten were presented at the same time. After this phase, participants were debriefed and thanked.

---

[19] Swapping from reading the global letters to the local letters is unlikely to completely control for task difficulty if reading local letters is more difficult than processing the global letters. Nevertheless, this is the procedure adopted by previous researchers.

5.1.2. Results

Figure 5.1 shows that local processing did indeed reduce performance, whereas global processing improved performance relative to the control condition in terms of percentage accuracy. The raw frequency data was examined using a series of chi-square tests: an overall chi-square followed by pairwise comparisons between each condition. The only effect to approach significance was the difference between the global and local conditions, χ^2 (1) = 3.737, p = .053. No other differences were significant: global and control, χ^2 (1) = 0.427, p > .51; control and local, χ^2 (1) = 1.843, p > .17. The effect size for this result was r = .20.
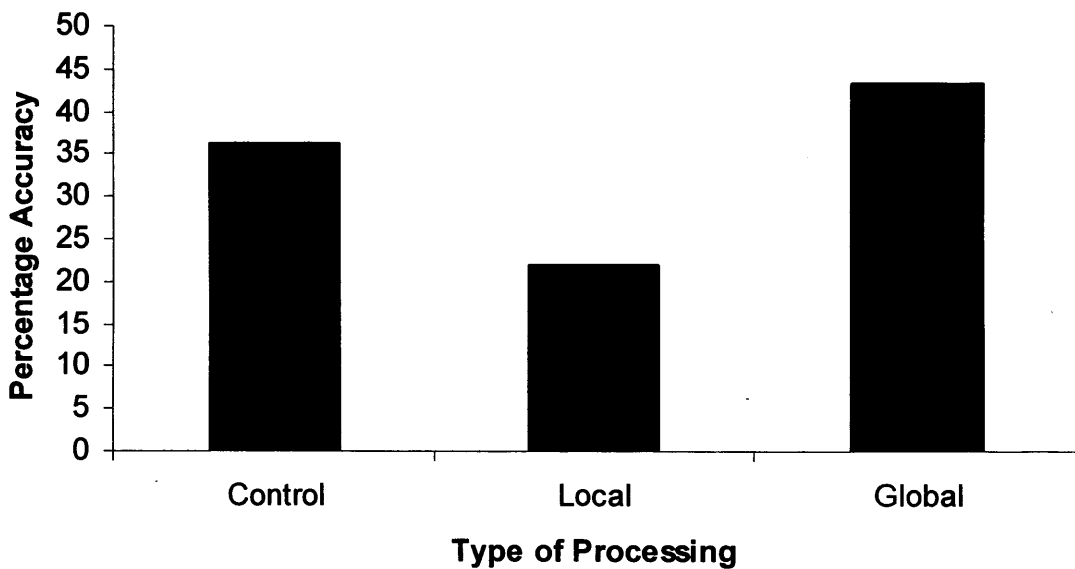
Figure 5.1. Percentage accuracy in a line-up task split by type of Navon processing undertaken.

5.1.3. Discussion

These results replicated the detrimental effect of local Navon processing on face recognition, albeit only approaching significance using the standard two-tailed test. As the hypothesis is unambiguously one-tailed (as it is a replication) then the results can be seen as supportive of the standard effect. More importantly, however, the data allow comparison with other manipulations to explore the necessary properties of the intervening task required to produce this effect.

5.2. Experiment 21 – Local Processing of Navon Shapes Causes a Face Recognition Deficit

Experiment 20 established that the detrimental effect of local Navon processing observed on subsequent face processing can be elicited using new face stimuli. Perfect (2003) hypothesised a link between the verbal overshadowing effect and the effect that Navon stimuli have on face processing. It is possible to look for similarities between these disparate tasks. They both involve an element of verbalisation in that participants are required to give a verbal response. One explanation for the verbal overshadowing effect is that of a mismatch between verbal and visual codes. Though this explanation has not been widely suggested to explain the effect of Navon stimuli processing observed here, it is possible since Navon letters contain a verbal component. By processing a series of Navon letters, a verbal code is induced which replaces the visual code. To explain why the local features cause a detriment, one could argue that since there are many more letters, the verbal code is exaggerated, whereas it would not be for the global figure.

Experiment 21, therefore, aimed to test whether the verbal component of the Navon task is important for the observed detriment in performance. Instead of using Navon letters, Navon shapes are used. Navon shapes have the benefit that they can be responded to without a verbal component, and although they have associated semantic and lexical information, by pointing to pictures of the shapes rather than naming them the verbal component is most likely be removed.

A second aim of Experiment 21 is to further extend the generalisability of the Navon effect by using new Navon stimuli. As such, should the verbal overshadowing effect be based upon a transfer-inappropriate processing shift, and should Navon shapes successfully induce this shift, then we would expect to replicate Experiment 20 with shapes rather than letters. However, if the explanation based on verbal and visual codes is more accurate, then we predict that there would be either a smaller effect or none at all.

5.2.1. Method

5.2.1.1. Participants

One-hundred-and-forty-two participants from Cardiff University's School of Psychology served as the participants for this experiment. All had normal vision and received course credits for taking part. The participants were randomly divided into one of three conditions: local, control, and global.

5.2.1.2. Materials

The same line-up and video were used as in Experiment 20. A set of 25 Navon shapes were created using Adobe Photoshop™ software. Five basic shapes were used: triangle, star, square, circle, and diamond. Each basic shape was combined to form five large

versions of the basic shapes. Images were presented at a size of 90 mm by 60 mm. An example of a shape Navon is presented in Figure 5.2.
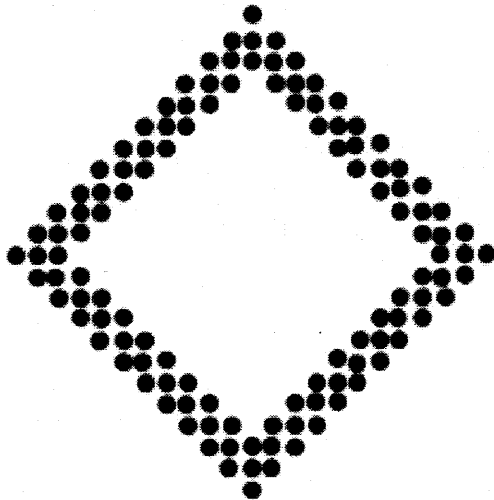
Figure 5.2. An example of a shape Navon used in Experiment 21 (A global diamond with local circles).

5.2.1.3. Design and Procedure

The design and procedure was identical to Experiment 20, except for the Navon task. For the Navon task, the participants were introduced to Navon shapes with three examples provided on screen. Participants were either instructed to identify the global Navon shape by pointing to a shape printed on the wall, or instructed to identify the local shape, again by pointing. The solid shapes were printed onto A4 paper, and were 150 mm by 150 mm positioned on the wall behind the computer screen one metre back and one metre higher than the computer screen.

5.2.2. Results

Percentage accuracy is shown in Figure 5.3, but the raw frequency data was tested using a series of chi-square tests. The overall chi-square was significant, χ^2 (2) = 7.645, $p < .05$. The difference between global and local processing conditions was significant χ^2 (1) = 7.604, $p < .05$, while the difference between control and local, and between control and global was not found to be significant. The effect size for the global-local difference was $r = .28$.
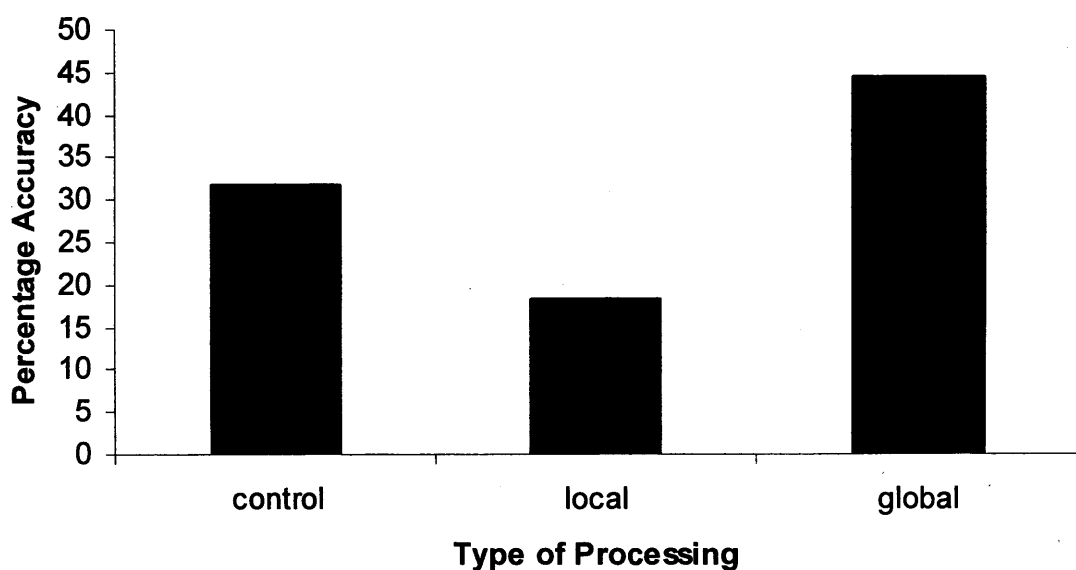
Figure 5.3. Percentage accuracy in a line-up task split by type of Navon processing.

The data from Experiments 20 and 21 were combined and analysed using a logistic regression (computed using Statview 5™, SAS Institute Inc.™). The dichotomous dependent variable was accuracy of face recognition. The independent variables were: the experiment (Navon letters or Navon shapes); and the Navon stimuli

processing type. This analysis works in a similar manner to multiple regression except that sigmoidal functions are used to allow dichotomous data to be analysed. Comparing local and global Navon stimuli processing revealed no significant effect of experiment, χ^2 (1) = 0.064, p > .88 and no significant experiment by condition interaction, χ^2 (1) = 0.157, p > .76. The same analyses were conducted for global-control and local-control comparisons. These all revealed no significant effect of experiment or interaction, all χ^2s < 0.1, all ps > .50.

5.2.3. Discussion

Participants processing the local Navon features were less accurate in a subsequent face recognition task. This was observed when using Navon shapes as well as Navon letters. Since it is most likely that the shape task did not involve a verbal component, as it did not require one, any explanation for this effect based on verbal and visual codes has not been supported. Granted shape stimuli can be easily verbalised, they can also be processed under conditions of verbal suppression. Thus, it is assumed that the shape Navon task relies less on verbal components. Despite this criticism, these results serve to generalise the Navon effect further. The logistic regression analysis demonstrates that these results are not significantly different from those in Experiment 20. Navon stimuli are very basic artificial visual stimuli and faces are much more complex natural stimuli. Experiment 22 aimed to extend this effect using somewhat more natural stimuli.

5.3. Experiment 22 – High- and Low-Pass Filtered Faces Do Not Cause a Transfer-Inappropriate Processing Shift

It has been commented that featural information within faces is contained in the higher spatial frequencies, while configural (possibly global) information is contained in the lower spatial frequencies (Costen, et al., 1996). If this assumption is accurate, then it is plausible to suggest that the processing of such filtered faces would cause a transfer-inappropriate processing shift similar to what Perfect (2003) hypothesises occurs when processing Navon stimuli. More specifically, if participants processed five minutes of faces that had low spatial frequencies filtered out, they would be forced to rely on featural processing. A subsequent line-up performance should, theoretically, be lower if the transfer-inappropriate processing shift explanation is accurate. Moreover, if the faces had the high spatial frequencies removed, the participants would have to rely on more configural processes which should be akin to processing the global Navon figure.

By processing faces instead of Navon letters, participants will be exposed to many more faces within the experimental paradigm. The problem associated with this is the possibility of interference from the adjusted faces onto the test faces (see e.g., Davies, Shepherd and Ellis, 1979; Deffenbacher, Carr, and Leu, 1981). This can be minimised by using different size stimuli in the test phase and the Navon phase. Nevertheless, since all participants will have faces to process in this paradigm, the interference will be the same across all conditions.

The same video and line-up procedure as Experiments 20 and 21 will be used. High spatial frequency filtered faces will contain only low spatial frequencies and should, therefore, induce global processing and enhance line-up performance. Low spatial frequency filtered faces will contain mostly high spatial frequencies and should,

therefore, induce local processing and thereby impair line-up performance. To ensure that participants actively process the bandpass filtered faces, famous faces shall be used and named. The experiment explored whether the processing bias could be produced using these stimuli and hence affect subsequent line-up performance.

5.3.1. Method

5.3.1.1. Participants

One-hundred Cardiff University Psychology undergraduates, all with self-reported normal vision, took part in this study as partial fulfilment of a course requirement.

5.3.1.2. Materials

The same video and line-up was used as in Experiments 20 and 21. Bandpass filtered faces were created using CorelDraw™ software. Two sets of 125 famous faces were created. One set was band pass filtered images with all frequencies below 32 cycles per face removed, while the other set was filtered to have all spatial frequencies above 32 cycles per face removed using a method similar to Näsänen (1999). This value was chosen because it is the middle range of useable spatial frequencies for face recognition making the split roughly in the middle of the useful information. Eighty percent of both high-pass- and low-pass-filtered faces were recognised during pretesting. Figure 5.4 shows an example of each. An additional set of unfiltered faces was also used. All faces were presented 100 mm by 100 mm. These faces were presented using SuperlabPro™ 2 Research Software on an RM PC.
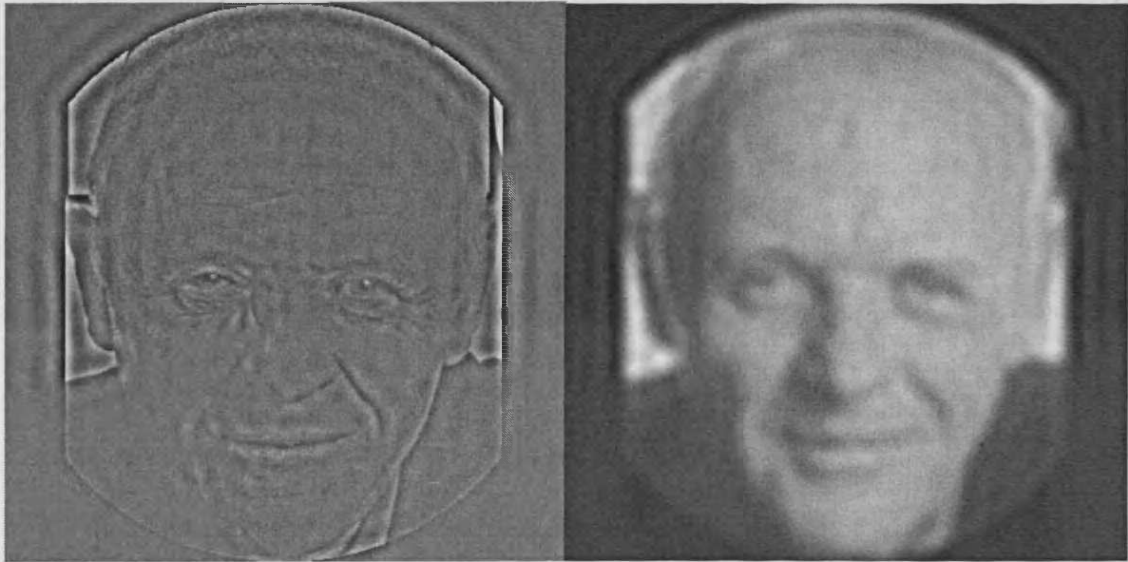
Figure 5.4. Anthony Hopkins containing only high spatial frequency information (left) and containing only low spatial frequency information (right).

5.3.1.3. Design

In this three-level between-subjects experiment, which followed the design of Experiments 20 and 21, the independent variable was the critical and final type of image that the participants were presented with during the intervening task. The dependent variable for this experiment was accuracy of face recognition at line-up. Participants were randomly selected to be in each group.

5.3.1.4. Procedure

As before, the video and line-up phases were identical to Experiment 20. Instead of the Navon task, participants were presented with one of the sets of bandpass filtered faces (faces containing mainly high or low spatial frequency), with 125 faces in all, and with each face viewed for 5 seconds. They were instructed to name the famous face. The instructions were the same for all three conditions (high spatial frequency filtered faces, low spatial frequency filtered faces, and unfiltered faces). After 5 minutes, the

participants in both filtered conditions swapped to the other filtered condition, as a replication of the previous work. Instead of reading in the control task, participants were presented with the unfiltered faces shown at the same rate as the filtered faces for 10 minutes.

5.3.2. Results

The accuracy rate was 32% for participants in the low spatial frequency filtered faces condition (potentially featural processors), 35% for participants in the high spatial frequency filtered faces (potentially configural processors), and 38% for the control participants. Chi-square tests were carried out in the same sequence as Experiments 19 and 20: global, followed by pairwise comparisons. There were no significant differences between any conditions (all χ^2s < 1, all ps > .80, all effect sizes r < .04).

The identification of faces preceded by either global or local Navon letter conditions (Experiment 20) or high and low spatial frequency filtered faces (Experiment 22) were compared in a similar logistic regression to that used in Experiment 21. This revealed a significant interaction between experiment and type of processing, χ^2 (1) = 4.085, p < .05, indicating that the results in this experiment were significantly different from those in Experiment 20. That is the effect of processing filtered faces was not just non-significant but it was significantly smaller than the effect observed in Experiments 20 and 21 where an effect was found.

5.3.3. Discussion

Costen et al. (1996) have suggested that featural or local aspects of faces are represented primarily by high spatial frequency information (albeit, the relative positions of the features remain in high-pass filtered faces). Conversely, configural or global information is primarily represented by low spatial frequency information. As such, processing high spatial frequencies for a considerable length of time (5 minutes in all) should induce featural processing, or at most processing first order relational information (cf., Rhodes, 1993) neither of which would be beneficial to face identification in this case. However, the results were statistically flat, indicating that this logic is flawed in some manner, and indeed the findings from this experiment were significantly different from the preceding experiments.

There are several possible reasons for the null results of Experiment 22. The first is that the Costen et al. paper is inaccurate in its assumption of spatial frequencies and their relation to types of processing. Even if their results are valid, it may be that studying faces is not able to induce configural or featural processing for other reasons, possibly due to the perceptual make-up of the stimuli set used. Each explanation will be elaborated on further.

Previous studies (see those summarised in Section 1.2.1.) have suggested that there is a correlation between the nature of facial encoding and spatial frequency information available from the face. Nevertheless, being presented with a series of filtered faces might not be sufficient to evoke a particular type of processing. It can be suggested that faces convey much more detailed semantic information than Navon stimuli. Familiar faces are associated with biographical histories, names, occupations and other semantic information (see e.g., Bruce and Humphreys, 1994; Burton and Bruce,

238

1993). This semantic information is not associated with Navon stimuli which only have an identity (the letter). Subsequently, the fact that filtered famous faces did not cause the detrimental effect on face recognition may suggest that the additional semantic information somehow blocks the effect. Alternatively, to obtain the semantic information from the faces, participants must attend to the whole stimulus, whereas to identify a Navon stimulus, participants need only attend to a small portion of the stimuli, especially in the local task. As such, it is possible that Navon stimuli direct attention to small components, rather than features of a whole.

It is also possible that semantic information has no influence on this effect. The detrimental effect of processing the local Navon letter on face processing may be much more low level. Experiments 24 to 29 directly assess this perceptual theory of this effect.

5.4. Experiment 23 – Local Processing of Navon Letters Causes a Face Recognition Deficit for 30 Seconds or 20 Trials

The number of participants required to test the effect of processing Navon stimuli on face recognition is vast. This experiment was conducted to replicate this basic effect of Navon processing on face recognition within a recognition paradigm. Verbal overshadowing effects last over 24 trials (Brown and Lloyd-Jones, 2002) or two days (Schooler and Engstler-Schooler, 1990). It was expected that those participants processing the global feature of a Navon letter would have better performance on a subsequent face recognition task than those processing the local features. The length of persistence of the effect upon recognition was also measured. Performance was measured using SDT measures d' and C.

5.4.1. Method

5.4.1.1. Participants

Twenty Cardiff University Psychology Undergraduates participated in this experiment as partial fulfilment of a course requirement. Three were male, 17 were female, and had a mean age of 19 years. All had normal or corrected vision and were selected by opportunity.

5.4.1.2. Materials

A collection of 40 photographs collected from the NimStim database was used (see Experiment 6). The Navon stimuli used in Experiment 20 were used here. All stimuli were presented using DirectRT™ Research Software (Empirisoft™) on an RM PC.

5.4.1.3. Design

A two-level one-way between-subjects design was employed, whereby participants were randomly divided into one of two groups and given the task of identifying the global Navon figure or the local Navon figures. The dependent variable was recognition performance. The presentation of faces was counterbalanced such that each face appeared as a target for as many participants as it appeared as a distracter. Moreover, the presentation order of faces in both the learning and the recognition phases was completely randomised.

5.4.1.4. Procedure

A standard recognition procedure was employed similar to Experiment 1, except that the Navon task replaced the distractor task. In the learning phase, the participants were

shown 20 target faces. They were asked to watch the screen. Each face was on the screen

for 3 seconds. As such, this phase lasted for 60 seconds. The Navon task was as described

in Experiment 20. The recognition phase followed immediately, in which participants

were asked to say whether each face was either 'old' or 'new' by pressing appropriate

keys on the keyboard. Faces were presented sequentially in a random order and all 40

faces were shown (20 targets and 20 distracters).

5.4.2. Results

Higher accuracy (d') was found for global Navon processors (mean d' = 1.53, standard

deviation = 0.19) than for local Navon processors (mean d' = 1.01, standard deviation =

0.28). This difference was found to be significant, $F(1, 18)$ = 17.658, MSE = 0.183, p <

.05. Its effect size was r = .55, Cohen's d = 1.33. No significant differences were found

for response bias, $F(1, 18)$ = 0.343, MSE = 0.017, p > .56 (mean C = 0.03 and 0.06

respectively).

During the recognition phase of the experiment, the length of time and number of

intervening stimuli since Navon processing increases with each face in the recognition

phase. As such, although the Navon effect may have been observed in full, further

analyses must be conducted to reveal if the Navon effect is consistent. To this end, the

data from the presentation of faces in the recognition paradigm was divided into four

blocks of ten. The accuracy measure d' was calculated for each block of ten faces for

each participant. The results are shown in Figure 5.5. The main effect of type of Navon

processing (reported above) was qualified by a significant interaction between trial block

and type of processing, $F(3, 54)$ = 5.785, MSE = 0.189, p < .05. This interaction was

explored using simple effects and showed that the global processors had a higher

accuracy than the local processors only for the first two blocks: mean difference = 0.71, p < .05 for trial block 1 and mean difference = 0.89, p < .05 for trial block 2. After face 20, the accuracy rates were similar, mean difference = 0.01, p > .88 for trial block 3 and mean difference = 0.01, p > .87 for trial block 4. The main effect of trial block was significant, $F(3, 54)$ = 2.979, MSE = 0.189, p < .05. Tukey HSD *post hoc* tests revealed that accuracy was greatest for trial block 4, but this was not significantly different from any of the other trial blocks, biggest mean difference = 0.405, p > .05, for trial block 1 and trial block 4. A parallel analysis was conducted on response criterion. No significant effects were observed, largest F = 2.018, p > .12.

Figure 5.5. Recognition Accuracy (d') split by type of processing over four trial blocks of 10 presentations. Error bars represent Standard Error.

5.4.3. Discussion

The results from Experiment 23 show the detrimental effect of local Navon processing being observed in a recognition paradigm, whereby participants processing the global figure of a Navon letter have higher accuracy in face recognition than those processing the local features. This effect is consistent with previous research (e.g., Perfect, 2003) but using this recognition paradigm the effect diminishes after 25 trials or 30 seconds. As such, it is important to develop a method to reinstate the effect during the recognition procedure.

5.5. Experiment 24 – The Effect of Local Navon Processing on Face Recognition Can be Reinstated

While Experiment 23 demonstrated the effect of processing Navon stimuli in a face recognition paradigm, it also observed that this effect is nullified by 25 faces being presented. Therefore, to fully explore such a transient effect in a recognition paradigm, Experiment 24 was devised. Between each face in the recognition test four additional Navon stimuli were processed to reinstate the effect. A control condition was added to assess whether this effect is due to local processors showing lower performance than they would at control or global processors showing an increased performance. The addition of the control task should also control for the inclusion of an interpolated task. Interpolated tasks are known to influence recognition tasks (c.f., the revelation effect, Verde & Rotello, 2003, 2004; Mulligan, & Lozito, 2006).

5.5.1. Method

5.5.1.1. Participants and Materials

All participants were Psychology undergraduates from Cardiff University with normal vision by self report. In total, 29 participants (9 males and 20 females), selected by convenience, took part for course credits and were divided randomly into one of three groups: global Navon processing, local Navon, or control. The same stimuli as used in Experiment 23 were used in Experiment 24.

5.5.1.3. Design and Procedure

A three-level one-way between-subjects design was employed, whereby participants were randomly divided into three groups and given the task of identifying the global Navon figure, the local Navon figure, or a control condition. The presentation of the Navon stimuli was randomised. The procedure for this Experiment was similar to that of Experiment 23, except that between each face in the test phase, four more Navon stimuli were presented and the participants were instructed to respond to them as they were previously. Control participants read from a book rather than processing any Navon stimuli. Between face stimuli, control participants continued to read from a passage in the book.

5.5.2. Results

Control participants had higher recognition accuracy (mean d' = 1.554) than global processors (mean = 1. 509) who had higher recognition accuracy than local Navon processors (mean = 0.596). These results were analysed using a univariate ANOVA. This

revealed a significant effect of processing, $F(2, 26) = 9.488$, MSE $= 0.287$, $p < .05$. The mean difference between local processors and control participants was significant (mean difference $= 0.958$, $p < .05$). The difference between local processors and global processors was also significant (mean difference $= 0.914$, $p < .05$). There were no significant effects of response bias, $F(2, 26) = 2.337$, MSE $= 0.109$, $p = .11$.

As in Experiment 23, the analysis was split by trial block. Figure 5.6. shows that the difference between global Navon processors and local Navon processors was consistent throughout. This main effect is described above. Tukey HSD *post hoc* tests revealed that local Navon processors had a significantly lower face recognition accuracy than global Navon processors (mean difference $= 0.914$, $p < .05$) and control participants (mean difference $= 0.958$, $p < .05$). The difference between control participants and global Navon processors was not significant (mean difference $= 0.044$, $p > .85$). There was also a main effect of trial block, $F(3, 78) = 5.063$, MSE $= 0.095$, $p < .05$, whereby accuracy was greater for trial block 4 than trial block 1 (mean difference $= 0.197$, $p < .05$), trial block 2 (mean difference $= 0.313$, $p < .05$), and trial block 3 (mean difference $= 0.184$, $p < .05$). No other differences were significant (largest mean difference $= 0.128$, $p > .12$). The interaction term was not significant, $F(6, 78) = 1.667$, MSE $= 0.095$, $p > .14$. The parallel analysis on response criterion revealed no significant effects, largest $F = 0.260$, $p > .85$.
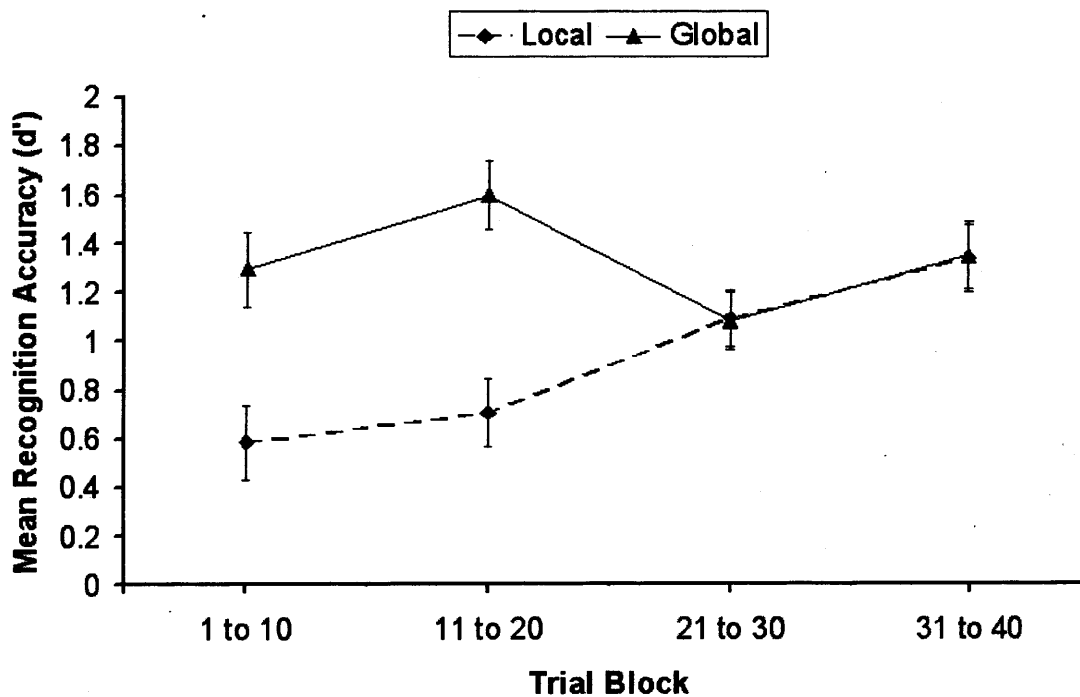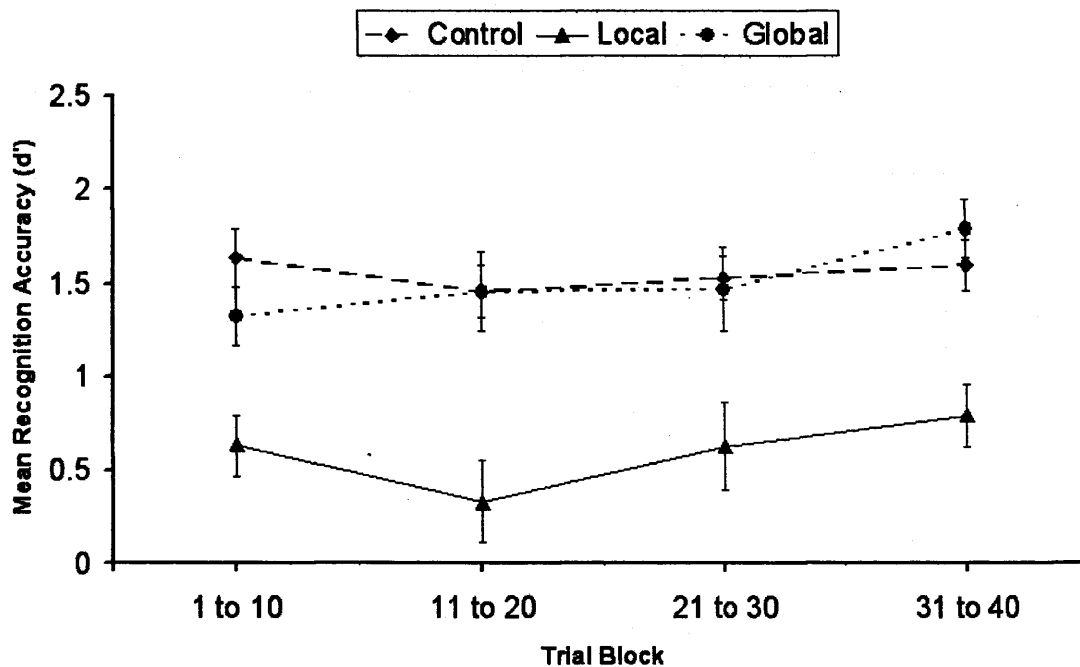
Figure 5.6. Recognition Accuracy (d') split by type of processing over four trial blocks of 10 presentations. Error bars represent standard error.

5.5.3. Discussion

Experiment 24 has demonstrated that it is possible to reinstate the effect that Navon stimuli have during the test phase of a recognition paradigm by presenting four further Navon stimuli. This process maintains the detrimental effect of processing the local features of the Navon letters over the duration of the 40 faces. Using this reinstatement method, it is possible to demonstrate the Navon effect in a recognition paradigm reliably. This development is important since fewer participants are required to obtain the same statistical power using this method.

Using this paradigm, significant differences were not found between the global Navon processors and the control participants, which is similar to findings by Perfect (2003). The results suggests that the key effect that the Navon stimuli have on face

recognition is the detrimental one of lowering accuracy. Since the verbal overshadowing effect has a prolonged duration (e.g., Schooler & Engstler-Schooler, 1990; Brown & Lloyd-Jones, 2002, 2003), the hypothesis that the detrimental effect of local Navon processing on face recognition is based on the same mechanism does not appear reliable. It becomes important to address what effects take five minutes of processing to occur that only last a minute. Elaborate perceptual after effects are a possibility, whereby processing high contrast high spatial frequency stimuli causes a detriment in discrimination performance on tasks that require high contrast and high spatial frequencies. Experiments 25 to 29 address this issue.

5.6. Experiment 25 – Processing the Local Features of Blurred Navon Stimuli Does Not Cause a Deficit in Face Recognition

The link between Navon stimuli and face recognition does not appear obvious. Section 1.2.1. introduced a possible correlation between local and global processing of Navon stimuli, local and global processing of faces, and the spatial frequency channels employed in perception. There is debate as to whether the link between local and global processing of Navon stimuli is the same as local and global processing in face processing (Weston and Perfect, 2005). The local features of a Navon stimulus are of high spatial frequency. These spatial frequencies are critically important for the identification of faces (Colin, et al., 2004; Schyns, et al., 2002). Since processing of Navon stimuli can cause perceptual after effects (e.g., Shulman and Wilson, 1987), it is entirely possible that the cause of the recognition deficit for faces following the processing of the local features of a Navon stimulus is due to adaptation to high spatial frequencies.

This discussion has therefore led to the suggestion that there are two possible explanations for the detrimental effect of processing the local features of a Navon letter on face recognition. An explanation based on transfer-inappropriate processing shift from expert configural coding to inexpert featural coding is the first. Second, a hypothesis based on contrast and spatial frequency adaptation interactions indicates a mismatch in the perceptual experience of the faces post adaptation to the Navon stimuli. These two hypotheses shall be referred to as the cognitive transfer-inappropriate processing shift and the perceptual processing mismatch respectively.

Since inverted faces are better processed using featural information, the cognitive processing shift explanation suggests local Navon processing will lead to improved recognition of inverted faces. However, the perceptual mismatch hypothesis suggests that inverted faces will be perceptually no different from upright faces and thus recognition will be poorer following processing of local Navon stimuli. By reducing the contrast and/or the spatial frequency of the Navon stimuli, the effect that the Navon stimuli have on face recognition can be removed. In this specific case, blurring removes the high spatial frequencies. Processing the local features of a low contrast (blurred) Navon stimulus should not cause a detriment in face recognition according to the perceptual mismatch hypothesis. Instead, it should cause an improvement in face recognition since there will be adaptation to less useful spatial frequency components. However, the standard detriment in face recognition will be observed if the cognitive processing shift explanation is accurate. The cognitive transfer-inappropriate shift explanation predicts that there will be an interaction between the Navon processing and the orientation of the faces, but no effect of blurring adjustment. The perceptual account, however, suggests that there will be an interaction between the adjustment of Navon stimuli and type of Navon processing, but no effect of orientation of the faces. The pattern of results

predicted by these two accounts for the detrimental effect local Navon processing has on

face recognition is presented graphically for clarity in Figure 5.7.

Predictions based on Cognitive Transfer

Predictions based on Perceptual After Effects
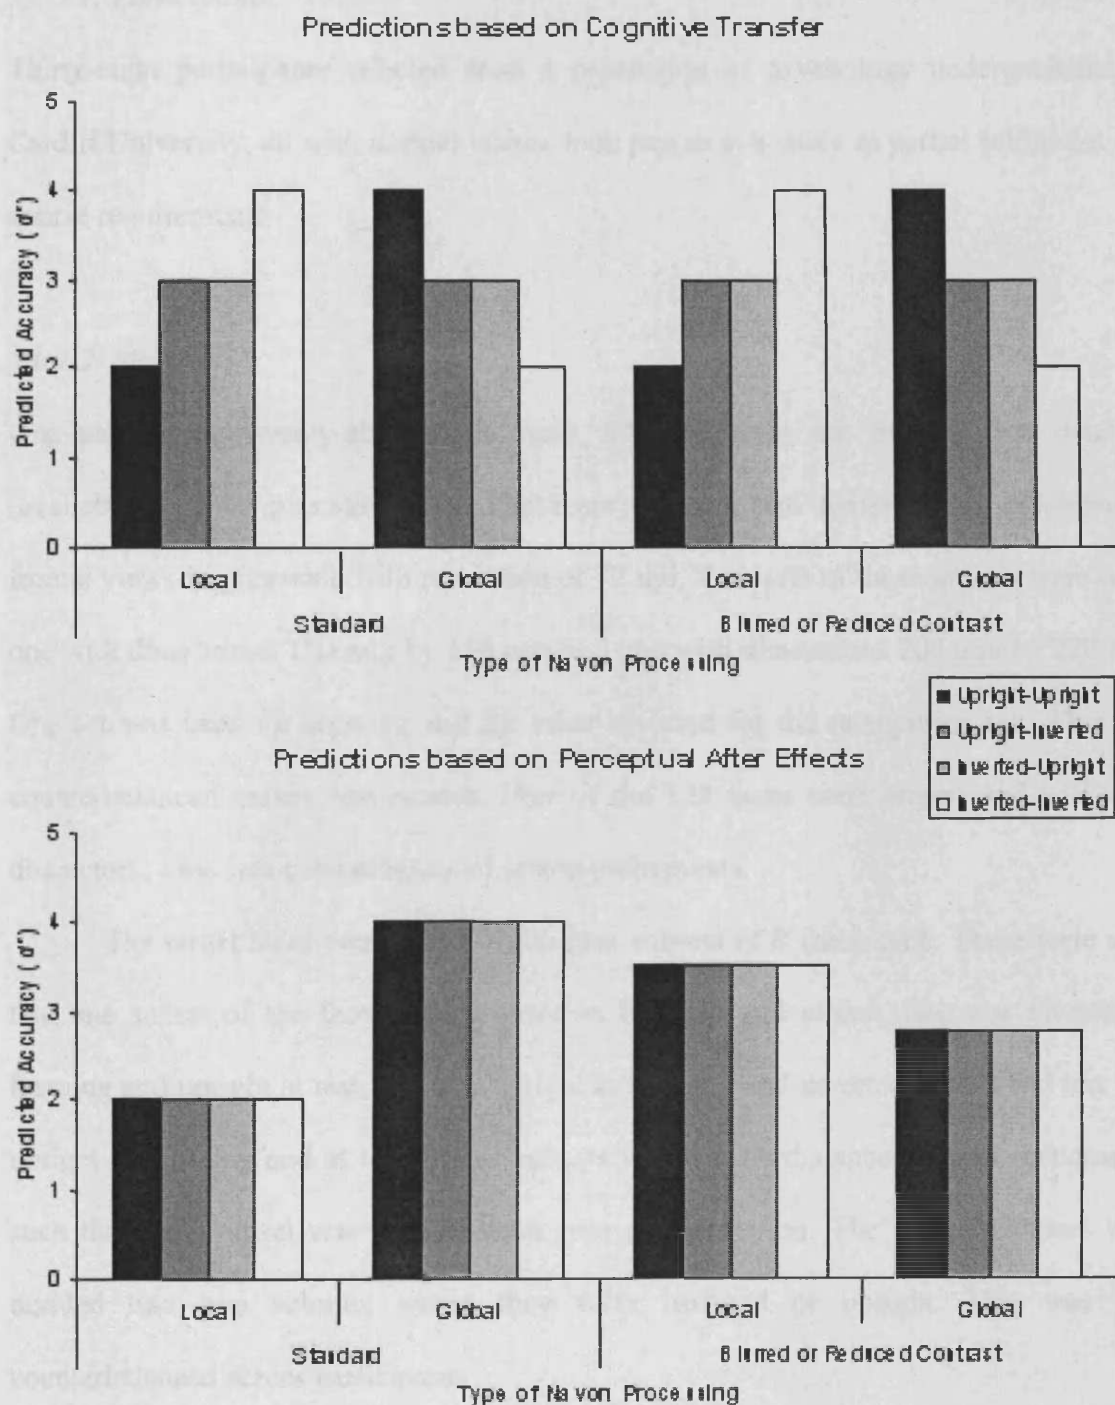
Figure 5.7. Predictions made by the cognitive transfer account (top panel) and the

perceptual after effect account (bottom panel) for Experiments 24 and 25.

249

5.6.1 Method

5.6.1.1. Participants

Thirty-eight participants selected from a population of psychology undergraduates at Cardiff University, all with normal vision, took part in this study as partial fulfilment of a course requirement.

5.6.1.2. Materials

One-hundred-and-twenty-eight faces were collected from the Stirling face database (available at www.pics.stir.ac.uk). These are of male and female faces presented in frontal views in greyscale with resolution of 72 dpi. Two sets of these images were used, one with dimensions 100 mm by 110 mm and one with dimensions 200 mm by 220 mm. One set was used for learning and the other set used for the recognition test. This was counterbalanced across participants. Half of the 128 faces were targets and half were distractors. This was counterbalanced across participants.

The target faces were divided into four subsets of 8 faces each. These were such that one subset of the faces was inverted at learning and at test, one was inverted at learning and upright at test, one was upright at learning and inverted at test, and one was upright at learning and at test. These subsets were counterbalanced across participants, such that each subset was used in each type of orientation. The distractor faces were divided into two subsets, where they were inverted or upright. This was also counterbalanced across participants.

In addition to the face stimuli, two sets of Navon stimuli were used. One was an unadjusted set used in Experiment 20. The second set of Navon stimuli was a blurred version of the first set. These were low-pass filtered using Adobe™ Photoshop™ with a

cut of 1.5 pixels (see Badcock et al., 1990, for a more detailed description of this procedure). The adjustments were made to the original Navon stimuli, as such all other variables were kept constant. An example of the Navon used in this Experiment and Experiment 25 are shown in Figure 5.8.

Figure 5.8. Examples of the Navon letters used in these experiments: a. An unaltered Navon stimulus (a global R made up of local Ns); b. A blurred Navon letter (a global U made up of local Os); c. A Navon stimuli with contrast reduced. (a global S made up of local Bs); compared to d. a face (PJH).

All the materials were present using an RM PC onto a high-resolution colour monitor using DirectRT™ Research Software (Empirisoft™). Participants sat 50 cm from the computer screen, and this was kept constant across all conditions.

5.6.1.3. Design

There were two between-subjects variables in this Experiment. These were the type of Navon processed (either global or local) and the type of Navon stimulus (unadjusted or blurred). There were two within-subjects variables in this Experiment. These were the orientation of the faces at learning (upright or inverted) and the orientation of the faces at test (upright or inverted). This led to a 2 x 2 x 2 x 2 mixed factorial design. The dependent variable in this Experiment was recognition accuracy measured in terms of hit rate, false alarm rate, and d'. Counterbalancing was implemented such that the faces appeared as a target or a distractor an equal number of times in each between-subjects condition. Further counterbalancing was implemented as described in the materials section. The faces were presented in a random order during the learning and the recognition phases of the Experimental procedure. Participants were randomly allocated to one of 4 experimental conditions (global unaltered Navon processing, local unaltered Navon processing, global blurred Navon processing, local blurred Navon processing), with the condition that there was roughly an equal number of participants in each condition (9 or 10).

5.6.1.4. Procedure

Participants were seated in a darkened laboratory, 50 cm from the computer screen. They were instructed not to move their head. A microphone was attached near to the participants' mouths to encourage them to respond accurately to the Navon stimuli. Participants were introduced to Navon stimuli. From then on, the Experiment had three phases: learning, Navon, and recognition. The Navon phase was the same as that reported

in Experiment 20. The recognition test was the same as that described in Experiment 24 but with 128 faces.

5.6.2. Results

A by-items analysis was conducted on these data. The recognition accuracy data in terms of d' are summarised in Figure 5.9. Figure 5.9 shows that when a face is learnt inverted it is subsequently harder to recognise. Moreover, Figure 5.9 shows poorer recognition for all types of facial orientation for local Navon processors than global processors when the Navon stimuli are unaltered. However, there is no recognition deficit for local processing when the Navon stimuli are blurred.

The data summarised in Figure 5.9 were subjected to a 2 x 2 x 2 x 2 mixed ANOVA with the factors: type of Navon stimuli (Unadjusted or Blurred); type of Navon processing (global or local); orientation of the face at learning (upright or inverted); and orientation of the face at test (upright or inverted). This revealed a significant interaction between the type of Navon stimuli and the type of Navon processing undertaken, $F(1, 34)$ = 21.327, MSE = 4.652, $p < .05$. An exploration of the simple effects revealed that the recognition accuracy was greater for global processors than local processors when the Navon stimuli were unaltered (mean difference = 1.828, $p < .05$), whereas recognition accuracy was greater for local processors than global processors when the Navon stimuli were blurred (mean difference = 1.409, $p < .05$). The only other main effect to reach significance was the main effect of facial orientation at learning, $F(1, 34)$ = 25.116, MSE = 4.689, $p < .05$. Recognition accuracy was greater when the face was learnt upright than when it was learnt inverted (mean difference = 1.763). No other main effects or interactions were significant, largest $F(1, 34)$ = 2.699, $p > .11$.
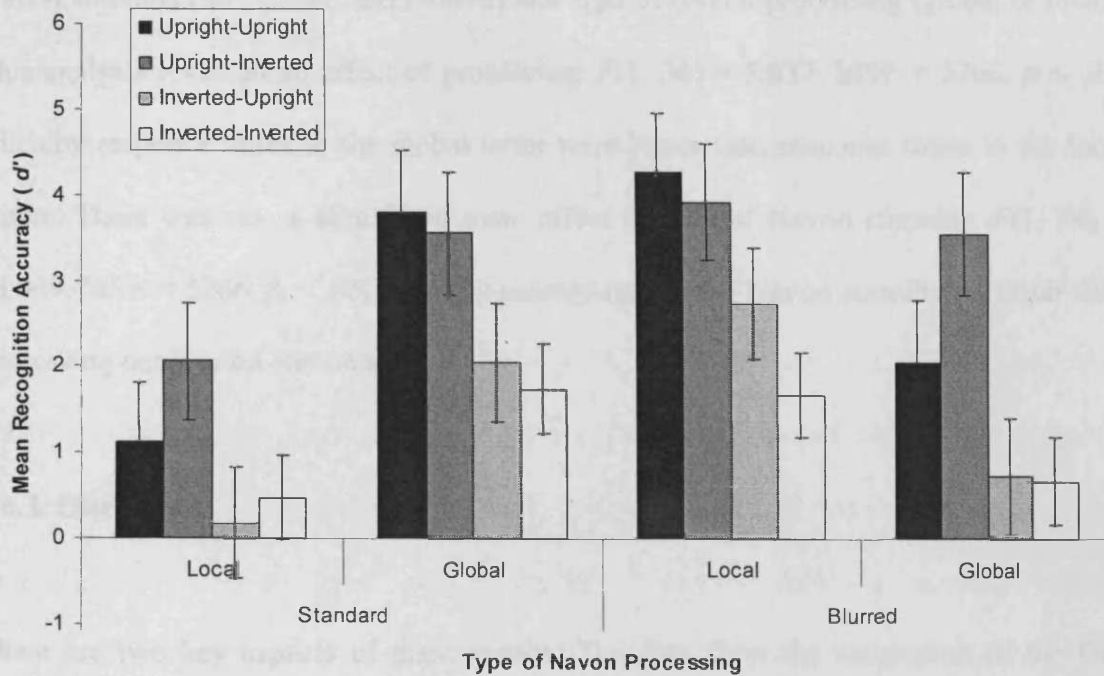
Figure 5.9. Recognition accuracy (d') split by type of Navon processed and orientation of faces at learning and at test. Error bars represent standard error.

One possible explanation of these effects is task difficulty, since the local Navon task may be more difficult than the global Navon task, and processing blurred Navon stimuli may be more difficult than processing unadjusted Navon stimuli. Such a suggestion is unlikely, since it is easy to detect the local features of an unadjusted Navon letter and less so for a blurred Navon letter. Nevertheless, the reaction time data for participants to respond to the Navon letters was assessed. The mean response time to identify the global letter in an unadjusted Navon stimulus was 754 ms compared to 796 ms (standard error = 23 ms) to identify the local letters in an unadjusted Navon stimulus. The mean response time was 637 ms to identify the global letter in a blurred Navon stimulus compared to 708 ms (standard error 24 ms) to identity the local letters in a blurred Navon stimulus.

The response time data was subjected to a two by two ANOVA with the factors: type of

Navon stimulus (unadjusted and blurred) and type of Navon processing (global or local).

This analysis revealed an effect of processing, $F(1, 34) = 5.637$, MSE = 5266, $p < .05$,

whereby response times to the global letter were lower than response times to the local

letters. There was also a significant main effect of type of Navon stimulus, $F(1, 34) =$

18.949, MSE = 5266, $p < .05$, whereby processing blurred Navon stimuli was faster than

processing unadjusted Navon stimuli.

5.6.3. Discussion

There are two key aspects of these results. The data from the orientation of the face

indicates that faces learnt inverted are harder to recognise irrespective of their orientation

at test. Indeed, this result is inconsistent with the suggestion that a mismatch in

orientation at presentation and recognition is detrimental to accuracy (e.g., Bruce, 1982;

Valentine, 1988). This result will be discussed in more detail in the general discussion.

The second aspect regards the recognition advantage after processing the global

letter in a Navon stimulus. This was observed for the unadjusted Navon stimuli but not

for the blurred Navon stimuli. Indeed, there was a recognition advantage for those who

processed the local letters. Thus, the standard effect of processing Navon stimuli before a

face recognition test is actually reversed by blurring the Navon stimuli. Removal of the

high spatial frequencies of the Navon letter prevented the detriment in subsequent face

recognition when processing the local letters of the Navon stimulus.

Crucially for the test of the mechanism behind this effect is the fact that the effect

of orientation was irrespective of the type of Navon processing undertaken. The

hypothesis that inverting a face would result in better recognition following local Navon

processing based on the cognitive shift explanation was not borne out by this Experiment. Local processing of unadjusted Navon stimuli led to lower recognition performance of all faces, irrespective of orientation. According to the cognitive shift explanation, blurring the Navon stimuli should not alter the effect they have on subsequent face recognition performance. The results from this Experiment clearly demonstrated that blurring Navon stimuli negates the detrimental effect of local processing in face recognition.

These three findings are inconsistent with a cognitive transfer explanation for the effect of Navon processing on face recognition. They are consistent with a perceptual account of the effect of Navon processing on face recognition, whereby processing the local features of Navon stimuli causes adaptation in high contrast high spatial frequency channels that are critical for the identification of faces. These findings demonstrate that removal of high spatial frequencies from the Navon stimuli is sufficient to remove the detrimental effect they have on face recognition.

There is a minor concern, however. The reaction time data for processing the Navon stimuli indicates that the blurred Navon stimuli are significantly easier to process than the unadjusted Navon stimuli. Due to the fact that the stimuli are easier to process, they may have less effect on subsequent face recognition (c.f., the revelation effect, e.g., Bornstein and Wilson, 2004). As such, it is possible that the ease of processing of the blurred Navon stimuli may account for the failure to find a detriment in recognition performance post local processing of blurred Navon stimuli.

5.7. Experiment 26 – Processing the Local Features of Contrast Adjusted Navon Stimuli Does Not Cause a Deficit in Face Recognition

The results from Experiment 25 indicate that a shift in cognitive processing from global to local does not explain why processing the local letters of a Navon letter should lead to a detriment in face recognition performance. An explanation based on perceptual after effects due to adaptation to specific spatial frequencies appears more appropriate. However, adaptation to specific spatial frequencies is dependent on contrast, such that adaptation to high spatial frequencies leads to great adaptation if the adaptor is of high contrast (Snowden and Hammett, 1996). The process of blurring a Navon stimulus removes the high spatial frequencies, but also reduces the contrast of the stimulus. Thus, the adjustments made to the Navon stimuli in Experiment 25 altered both contrast and spatial frequency. Experiment 26, therefore, aimed to replicate the findings of Experiment 25 using Navon stimuli that had their contrast adjusted directly. Experiment 26 also addressed the two outstanding issues of Experiment 25, described above.

A direct contrast manipulation will interfere with after-effects of spatial frequency. If the Navon stimuli are of lower contrast, adaptation to high spatial frequencies will be weaker than in a Navon stimulus of high contrast. The cognitive shift explanation maintains that adjusting the Navon stimuli in this way will not alter the detrimental effect on face recognition that processing the local letters of the Navon stimuli will have. Inverting the faces will reverse the Navon effect according to the cognitive shift explanation, whereas the perceptual account suggests that the inversion effect will be independent of the Navon processing style. The explanations again make two distinct predictions about the pattern of results, as shown in Figure 5.7.

5.7.1. Method

5.7.1.1. Participants

Thirty-nine participants selected from a population of psychology undergraduates at Cardiff University, all with normal vision, took part in this study as partial fulfilment of a course requirement.

5.7.1.2. Materials

The same face database and the same unadjusted set of Navon letters were used in Experiment 26 as in Experiment 25. In addition, an adjusted set of Navon stimuli were created using CorelDraw™. These had their contrast reduced by 75%. The resulting stimuli were darker than the original set. This was compensated by increasing the luminance by 25%. The adjustments were made to the original Navon stimuli and all other variables were kept constant. See Figure 5.7 for an example of the stimuli used in this Experiment.

5.7.1.3. Design and Procedure

All aspects of the design and procedure were identical to Experiment 25: participants saw a set of faces then went under the Navon manipulation, followed by the recognition test intermixed with additional Navon stimuli.

5.7.2. Results

As in Experiment 25, an analysis on d' stimulus discriminability was conducted. This Experiment also employed a by-items analysis. The data are summarised in Figure 5.10.

Figure 5.10 shows that when a face is learnt inverted it is subsequently harder to recognise. Moreover, Figure 5.10 shows poorer recognition for all types of facial orientation for local Navon processors than global processors when the Navon stimuli are unaltered. However, there is no recognition deficit for local processing when the Navon stimuli are reduced in contrast.

The data summarised in Figure 5.10 were subjected to a 2 x 2 x 2 x 2 mixed ANOVA with the factors: type of Navon stimuli (unadjusted or contrast-adjusted); type of Navon processing (global or local); orientation of the face at learning (upright or inverted); and orientation of the face at test (upright or inverted).. This revealed a significant effect of type of Navon processing, $F(1, 35) = 5.997$, MSE $= 2.830, p < .05$, whereby global Navon processors had a significantly greater accuracy than local Navon processors (mean difference $= 0.660$). This main effect was qualified by an interaction between the type of Navon stimuli and the type of Navon processing, $F(1, 35) = 11.958$, MSE $= 2.830, p < .05$. Simple effects revealed that participants who processed the global letters of unadjusted Navon stimuli were more accurate than those who processed the local letters of unadjusted Navon stimuli (mean difference $= 1.593, p < .05$). However, the difference in accuracy between those who processed the local and global letters in a reduced contrast Navon stimulus was in the opposite direction and not significantly different (mean difference $= 0.172$, ns). The main effect of face orientation at learning was also significant, $F(1, 35) = 31.226$, MSE $= 4.216, p < .05$, whereby faces that were learnt upright were better recognised than faces that were learnt inverted (mean difference $= 1.839$). No other main effects or interactions were significant.
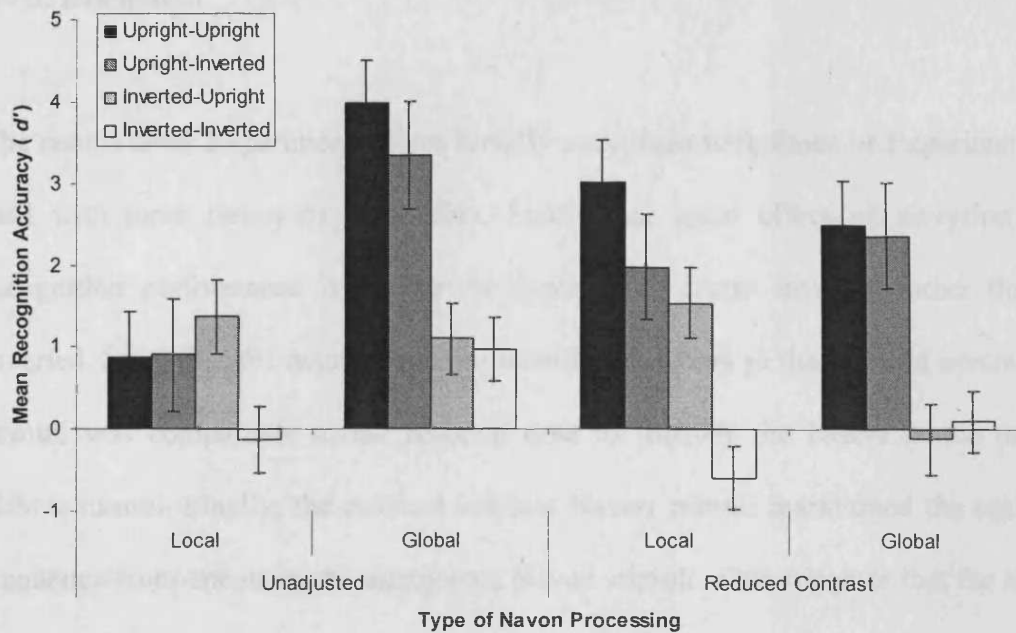
Figure 5.10. Recognition accuracy split by type of Navon processed and orientation of faces at learning and at test. Error bars represent standard error.

As in Experiment 25, the response time to identify the global and local letters in both types of Navon stimuli were compared to assess the effects of task difficulty. The mean response time to identify the global letter in an unadjusted Navon stimulus was 750 ms (standard error 26 ms) which was less than the response time to identify the local letters in unadjusted Navon stimuli (mean = 803 ms, standard error = 28 ms). The response time to identify a global letter in a reduced contrast Navon stimulus was 775 ms and was less than the response time to identify the local letter in a reduced contrast Navon stimulus (mean = 794 ms). The response time data was subjected to a univariate ANOVA, which failed to reveal any significant effects, largest $F(1, 35) = 1.853$, $p > .18$.

260

5.7.3. Discussion

The results from Experiment 26 are broadly consistent with those of Experiment 25 and deal with three issues of those data. Firstly, the main effect of inversion on face recognition performance is due to the faces being learnt inverted rather than tested inverted. Secondly, the reaction time to identify the letters in the reduced contrast Navon stimuli was comparable to the reaction time to identify the letters in the unadjusted Navon stimuli. Finally, the reduced contrast Navon stimuli maintained the same spatial frequency components as the unadjusted Navon stimuli. This suggests that the effect that Navon stimuli have on face recognition is due to the interaction between contrast and spatial frequency. The detriment in face recognition performance after processing the local letters of the Navon stimuli is due to the high contrast high spatial frequency components of these stimuli. If either is removed then the effect that Navon stimuli have on face recognition is removed.

There are some minor inconsistencies in the data between Experiments 25 and 26 that do need to be considered. Generally, the pattern of results is the same in these two Experiments, although the data are slightly untidy. However, for local Navon processors in Experiment 26, faces learnt inverted but tested upright were recognised more accurately than in Experiment 25, irrespective of the type of Navon processed. Since the task is identical in the two Experiments for standard Navon stimuli, there is no explanation for this apart from either random fluctuations or the possibility that the effect that Navon stimuli have on face recognition is vulnerable to even subtle changes in the testing environment.

Since the effect of processing the local letters in Navon stimuli is removed by reducing the contrast of the Navon stimuli and this effect remains irrespective of the

orientation of the faces, these results are incompatible with the cognitive shift explanation. These results are consistent with a perceptual after-effects account. The detriment in face recognition performance after processing the local features of Navon stimuli only occurs when they are of high contrast and high spatial frequency. Experiment 27 aims to directly test different types of high contrast and high spatial frequency stimuli on subsequent face recognition performance.

5.8. Experiment 27 – High Contrast High Spatial Frequency Letters Cause a Face Recognition Performance Deficit

While maintaining a letter identification task akin to the Navon task, Experiment 27 aimed to test whether other stimuli of similar perceptual make-up to Navon stimuli can cause detriment in facial recognition. Instead of identifying Navon stimuli, participants in Experiment 27 will be presented with letters. These letters will either be big (low spatial frequency) or small (high spatial frequency), with contrast adjusted to mimic that of the Navon stimuli. These conditions are an extension of a study conducted by Brand (2005). Brand found a marginal reduction in facial identification performance after processing small high contrast letters akin to processing the local letters in a Navon stimulus. If the effect of processing local letters in a Navon stimulus is due to the spatial frequency and contrast components of the letters, then simply reading small high contrast letters will cause a detriment in subsequent face recognition. Processing large low contrast letters will cause a slight improvement in face recognition, akin to global processing of Navon stimuli. Two additional conditions shall be implemented where low contrast high spatial frequency letters shall be used, which are akin to the local features in reduced contrast Navon stimuli used in Experiment 26. The final type of stimuli will be low spatial

frequency high contrast letters. This condition is conducted to complete the design (a 2 by 2: size by contrast). A large high contrast letter is not like any of the Navon type processing, whereas the low contrast large letter is akin to global Navon processing. As such, processing a low contrast large letter should improve face recognition, whereas processing a high contrast large letter should not improve nor cause a detriment to face recognition. Reading these letters should not have an effect on subsequent face recognition if the basis for the effect of local processing of Navon stimuli on face recognition is due to a cognitive shift. Examples of the stimuli and the predictions based on the cognitive shift and the perceptual after effect explanations are presented in Figure 5.11.

Predictions based on Cognitive Transfer

Predictions based on Perceptual After Effects
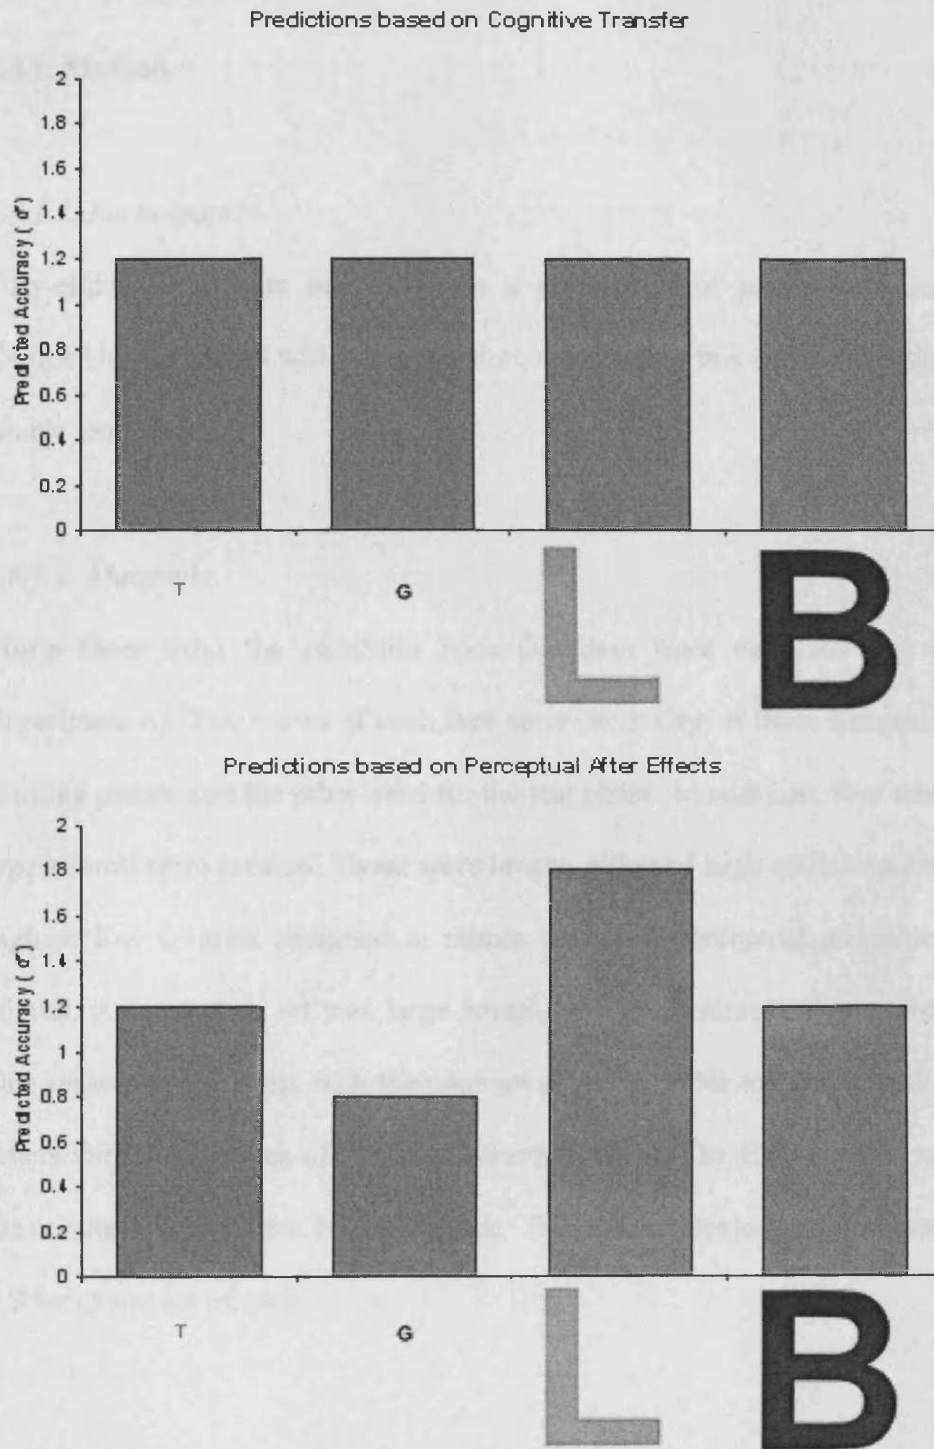
Figure 5.11. Predictions made by the cognitive transfer account (top panel) and the perceptual after effect account (bottom panel) for Experiment 3 and examples of the stimuli used in Experiment 3. From left to right: a small, low contrast letter; a small, high contrast letter; a large, low contrast letter; and a large, high contrast letter.

5.8.1. Method

5.8.1.1. Participants

Fifty-eight participants selected from a population of psychology undergraduates at Cardiff University, all with normal vision, took part in this study as partial fulfilment of a course requirement.

5.8.1.2. Materials

Thirty faces from the NimStim Face Database were used for this experiment (see Experiment 6). Two views of each face were used: One of these images was used for the learning phase; and the other used for the test phase. In addition, four new sets of Navon-type stimuli were created. These were letters, either of high or low spatial frequency and high or low contrast designed to mimic the basic perceptual properties of the Navon stimuli. As such, one set was large letters with the contrast of the global Navon figure. One set was large letters with the contrast of the local Navon features. One set was small letters with the contrast of the local Navon features. The final set was small letters with the contrast of the global Navon feature. These were created in CorelDraw™. See Figure 5.9 for examples of each.

5.8.1.3. Design and Procedure

A 2 by 2 between-subjects design was employed whereby the independent variables were the contrast (high or low) and the spatial frequency (high or low)[21] of the letters to be identified. All other aspects of the design and procedure were identical to Experiment 25.

5.8.2. Results

The recognition accuracy data is presented in Figure 5.12, and shows that recognition accuracy was highest when low contrast and low spatial frequency letters had been identified. Conversely, recognition accuracy was lowest when high contrast and high spatial frequency letters had been identified. A 2 x 2 between-subjects ANOVA was run on these d' recognition accuracy data. This revealed a non-significant interaction between spatial frequency and contrast, $F(1, 33) = 0.409$, MSE $= 0.360$, $p > .52$. However, the main effect of spatial frequency on recognition accuracy was significant, $F(1, 33) = 4.535$, MSE $= 0.36$, $p < .05$, whereby participants identifying letters of lower spatial frequency were more accurate than participants identifying letters of higher spatial frequency (mean difference $= 0.424$). The main effect of contrast was approaching significance, $F(1, 33) = 3.520$, MSE $= 0.360$, $p > .07$, whereby participants processing lower contrast stimuli were more accurate than those processing higher contrast stimuli (mean difference $= 0.374$).

---

[21] Granted spatial frequency was only indirectly manipulated by having two sizes of stimuli (big – low spatial frequency) and small (high spatial frequency).
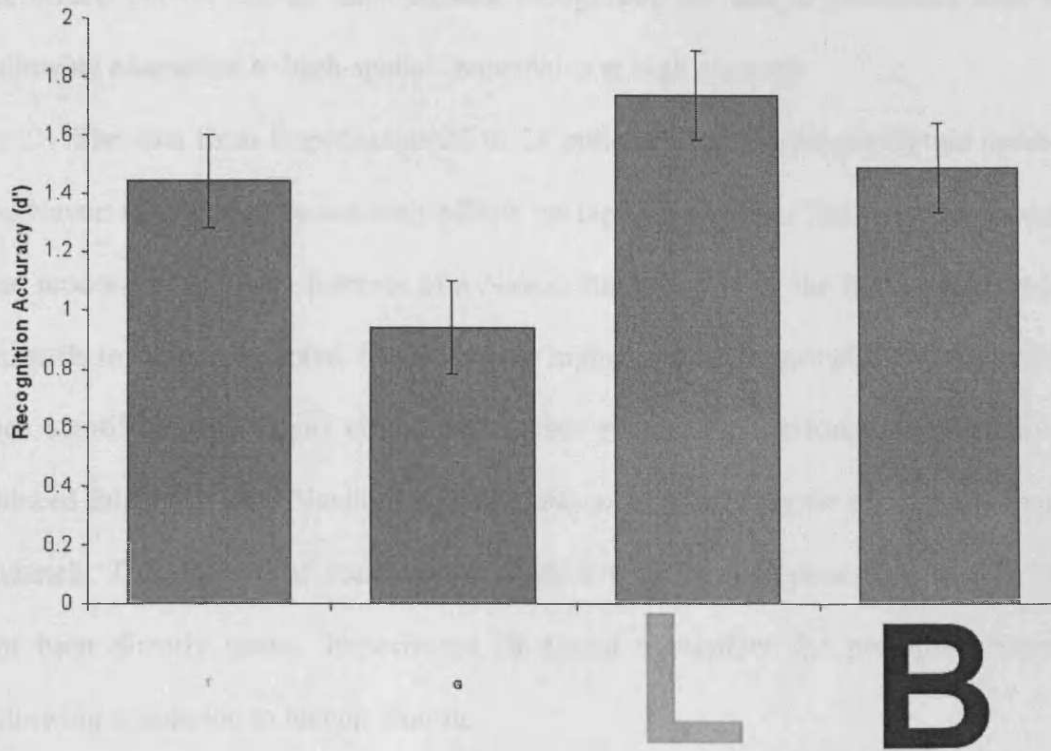
Figure 5.12. Recognition accuracy (d') split by type of stimulus. Error bars represent standard error.

5.8.3. Discussion

The results from Experiment 27 indicate that the detrimental effects of processing the Local features of a Navon letter on face recognition can be found in non-Navon letters that are of the same spatial frequency and contrast as the local features of a Navon letter. In other words, the spatial frequency and contrast components of the Navon stimuli are more crucial than the fact they contain local and global features in their effect on face recognition. Thus, it is not anything specific to the hierarchical nature of the Navon stimuli that causes their processing to influence face recognition, but simply the spatial frequency components that make up the stimuli. This theoretical analysis suggests that

the effects Navon stimuli have on face recognition are due to perceptual after-effects following adaptation to high-spatial frequencies at high contrast.

The data from Experiments 25 to 27 indicate that it is the perceptual make-up of the Navon stimuli that causes their effects on face recognition. The hypothesis posited is that processing the local features of a Navon stimulus causes the high spatial frequency channels to become adapted. Since it is the higher spatial frequencies that are crucial for face identification (Schyns et al., 2002), face recognition performance is significantly reduced following local Navon processing due to adaptation to the high spatial frequency channels. This theoretical analysis is consistent with the data presented thus far, but has not been directly tested. Experiment 28 aimed to explore the perceptual experience following adaptation to Navon stimuli.

5.9. Experiment 28 – The Effect of Processing Navon Stimuli on High- and Low-Pass Filtered Faces

Experiments 25 to 27 indicate that processing the local features of a Navon letter cause a detriment in face recognition performance due to the fact that the local features are of high contrast and high spatial frequency. Processing these high contrast high spatial frequency stimuli causes adaptation to the high contrast high spatial frequency perceptual channels. This adaptation causes high spatial frequencies to be less discriminable post adaptation. These high spatial frequencies are important for facial identification (Schyns et al., 2002). Removal of these spatial frequencies will thus make facial identification less accurate. Indeed, faces that have high spatial frequencies removed are recognised less readily than broadband faces (e.g., Gold et al., 1999).

Experiment 28 takes these lines of evidence and combines them, by testing this directly by presenting participants with a series of faces that were low-pass filtered, high-pass filtered, or broadband faces and asking them to identity either the local or the global components of a series of Navon stimuli. If a face is filtered such that it contains only high spatial frequencies, post adaptation to high spatial frequencies, these faces will be largely impossible to recognise. Figure 5.13. presents the hypothesised pattern of data for the cognitive shift and the perceptual after effect explanation for the detrimental effect processing the local features of a Navon stimulus has on subsequent face recognition.

Predictions based on Cognitive Transfer

Predictions based on Perceptual After Effects
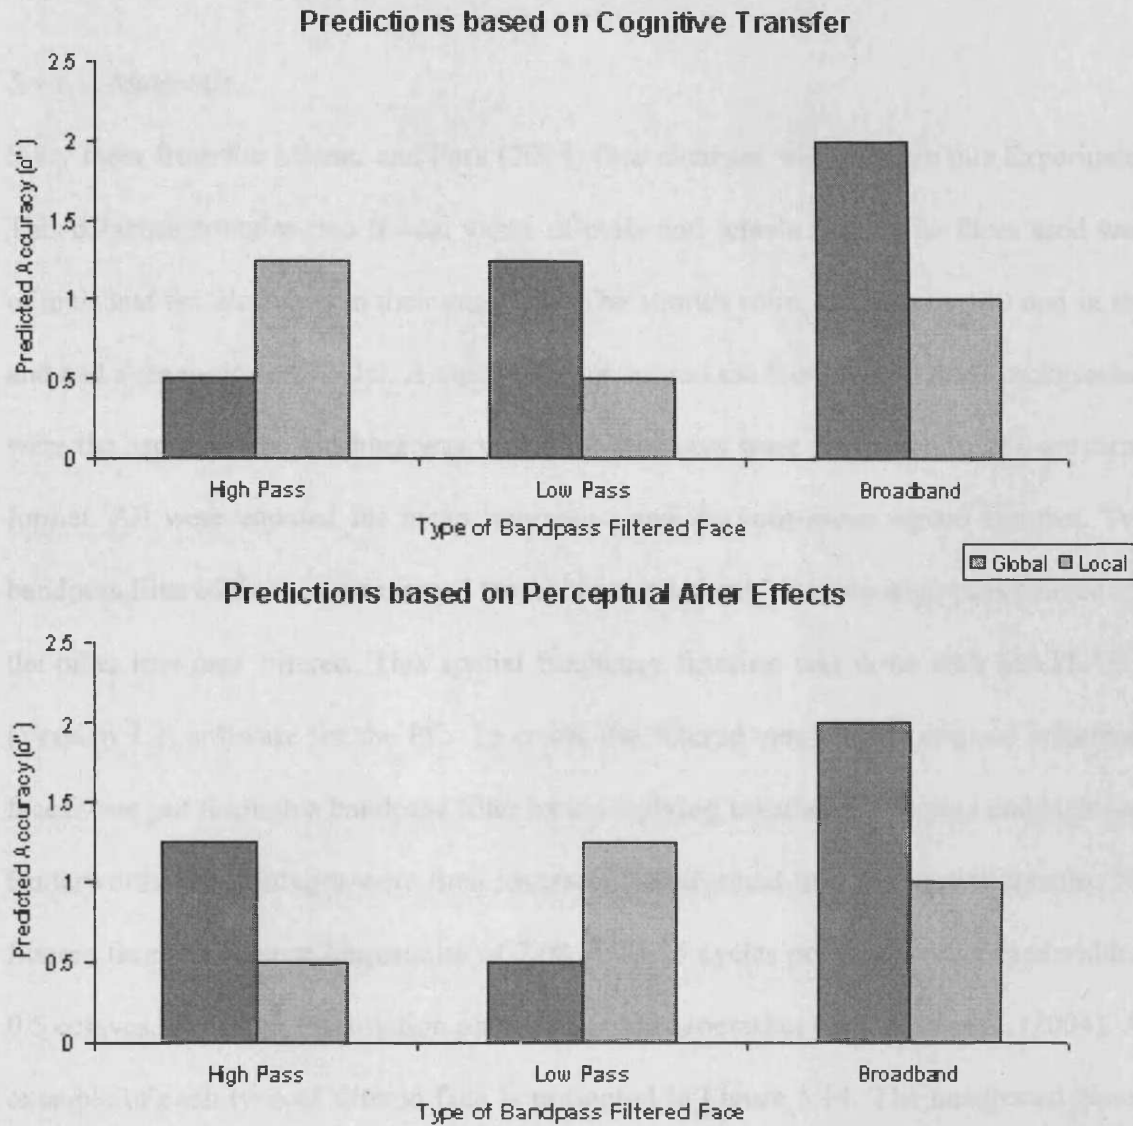
Figure 5.13. Predictions made by the cognitive transfer account (top panel) and the perceptual after effect account (bottom panel) for Experiments 28 and 29.

5.9.1. Method

5.9.1.1. Participants

Sixteen participants selected from a population of psychology undergraduates at Cardiff University, all with normal vision, took part in this study as partial fulfilment of a course requirement.

5.9.1.2. Materials

Sixty faces from the Minear and Park (2004) face database were used in this Experiment. This database contains two frontal views of male and female faces. The faces used were of male and females aged in their early 20s. The stimuli were 130 mm by 140 mm in size and had a resolution of 72 dpi. A mask was put around the face, such that all backgrounds were the same and no clothing was visible. All images were converted to 256 greyscale format. All were equated for mean luminance and the root-mean square contrast. Two bandpass filtered faces were created from this original set, one was high-pass filtered and the other low-pass filtered. This spatial frequency filtering was done with MATLAB™ (Version 7.2) software for the PC. To create the filtered version, the original broadband faces were put through a bandpass filter by multiplying together a low-pass and high-pass Butterworth filter. Images were then inversely transformed into the spatial domain. The filtered faces had centre frequencies of 7.08 or 14.15 cycles per face, with a bandwidth of 0.5 octaves. For more information on the stimulus generation see Collin et al. (2004). An example of each type of filtered face is presented in Figure 5.14. The unadjusted Navon stimuli from Experiment 24 were used in this study. All presentation software and equipment was the same as that used in the previous Experiments.
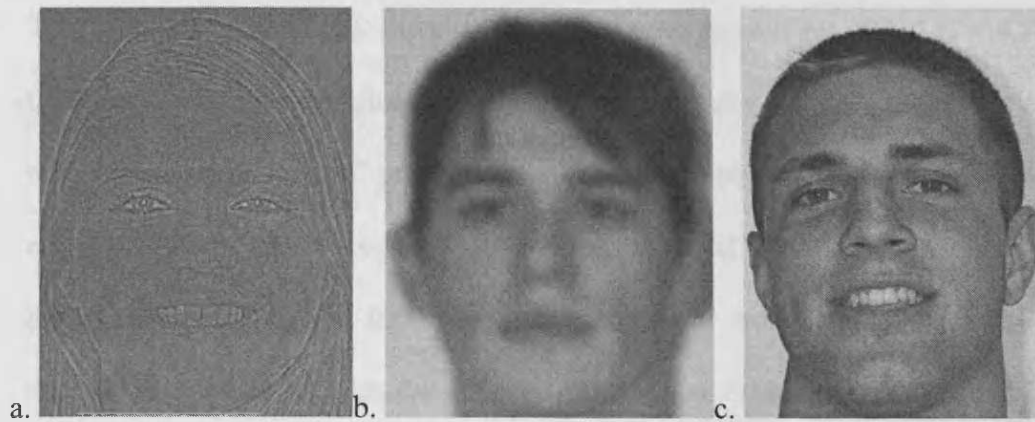
Figure 5.14. Examples of the bandpass filtered faces used for Experiments 28 and 29. a. High-pass filtered. b. Low-pass filtered. c. Broadband (unfiltered).

5.9.1.3. Design and Procedure

This Experiment employed a 2 by 3 Mixed design, whereby the type of Navon processing (Global or Local) was manipulated between-subjects and the bandpass filtered faces (high-pass, low-pass, or broadband) were manipulated within-subjects. The bandpass filtering was matched at learning and at test, such that the spatial frequency overlap was 100%. The faces were counterbalanced, such that each face appeared as a target and as a distractor an equal number of times. Moreover, the faces were counterbalanced, such that they appeared in each class of bandpass filtering a roughly equal number of times. The presentation order of the faces was randomised. The procedure for this Experiment was identical to that in Experiment 25.

5.9.2. Results

These recognition accuracy data scores are summarised in Figure 5.15 and were subjected to a 2 x 3 mixed ANOVA with the factors: type of Navon processing (global or local) and type of bandpass filtering of the faces (high-pass, low-pass, or broadband).

This revealed a significant effect of type of Navon processing, $F(1, 14) = 4.787$, MSE = 0.557, $p < .05$, whereby global processors had a greater recognition accuracy in the face recognition test than local processors (mean difference = 0.471). There was also a main effect of type of bandpass filtering on the recognition accuracy of the faces, $F(1, 14) = 8.021$, MSE = 0.361, $p < .05$, whereby recognition of broadband faces was more accurate than the recognition of high-pass filtered faces (mean difference = 0.463, $p < .05$) and the recognition of low-pass filtered faces (mean difference = 0.850, $p < .05$). Recognition of high-pass and low-pass faces were not significantly different from each other (mean difference = 0.387). These main effects were qualified by a significant interaction between the type of Navon processing and the type of bandpass filtering of the faces, $F(1, 14) = 9.167$, MSE = 0.361, $p < .05$. Simple effects showed that the recognition accuracy of broadband faces was significantly greater for global than local Navon processors (mean difference = 0.898, $p < .05$). Recognition accuracy of high-pass filtered faces was significantly greater for global than local Navon processors (mean difference = 1.079, $p < .05$). Recognition accuracy of low-pass filtered faces was significantly *less* for global than local Navon processors (mean difference = 0.574, $p < .05$).
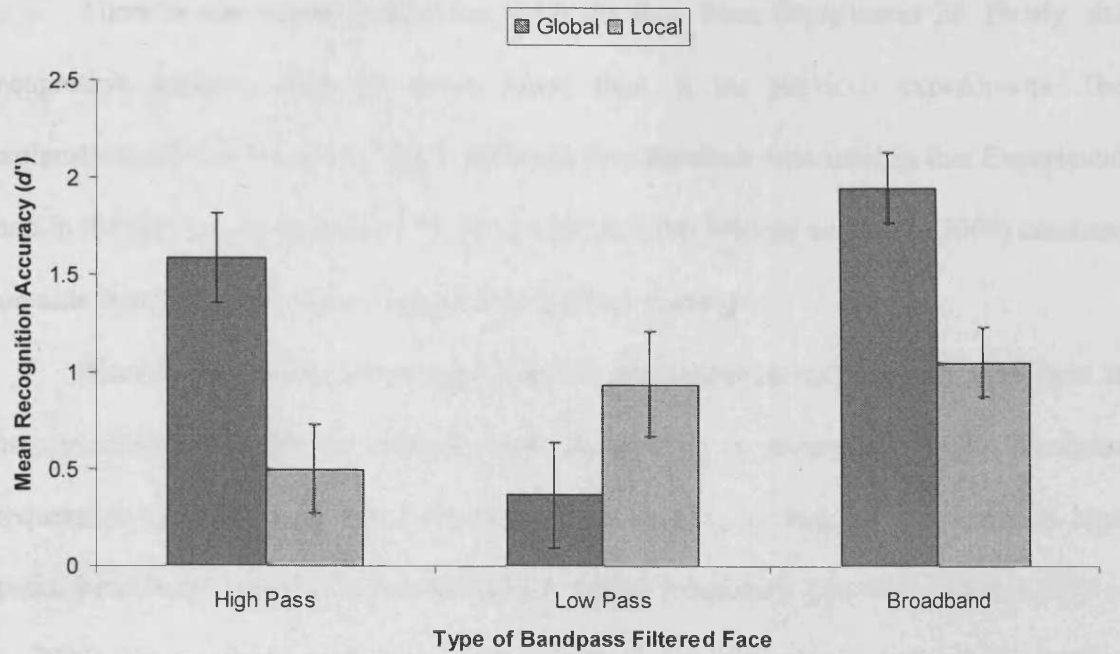
Figure 5.15. Recognition accuracy split by type of Navon processing undertaken and type of bandpass filtered faces for Experiment 28. Error bars represent standard error.

5.9.3. Discussion

These results are broadly consistent with the hypothesis that processing the Navon stimuli causes spatial frequency adaptation. Processing the local components of a Navon letter causes adaptation in the high spatial frequency channels. Thus, recognition of faces containing only high spatial frequencies (high-pass faces), will be virtually impossible following processing of the local components of the Navon stimuli. Processing of the global component in Navon stimuli will cause adaptation in the low spatial frequency channels that are less useful for face identification. A low-pass face contains only these spatial frequencies, and thus after adaptation to the low spatial frequencies, these faces will become virtually impossible to recognise. This is indeed what is observed.

There is one concern, however, with the data from Experiment 28. Firstly, the recognition accuracy data are much lower than in the previous experiments. The explanation offered for this is that a different face database was used in this Experiment than in the previous experiments. It is possible that the Minear and Park (2004) database contains less distinctive faces than the Stirling face database.

Another interesting point regarding the experimental method employed here is that processing the Navon stimuli may be causing a mismatch in the bandpass frequencies available to process the faces from learning to test. A face learnt at high spatial frequency is better recognised at high spatial frequency (Liu et al., 2000; Collin et al., 2004). Thus, a face learnt at low spatial frequency will be more accurately recognised in a broadband face following local Navon processing than global Navon processing because the local Navon processing removes the high spatial frequencies. Since these, in this example, were not used for learning the face then local Navon processing will benefit subsequent recognition of a broadband face. Experiment 29 aimed to test this hypothesis.

5.10. Experiment 29 – Local Processing of Navon Stimuli Causes a Deficit in Face Recognition Due to Adaptation to High Contrast High Spatial Frequencies

Experiment 28 demonstrated that faces containing only high spatial frequencies are virtually impossible to recognise following the processing of the local features of the Navon stimuli. It was hypothesised that the processing of Navon stimuli selectively removes part of the spatial frequency channels available for recognition. In the case described, processing the local features of a Navon stimulus removes high spatial frequencies, leading to virtually no available spatial frequency channels available to recognise the high spatial frequency face. Thus, it was hypothesised that local processing

of a Navon stimulus may actually benefit the recognition of broadband faces if the face had been originally learnt without any low spatial frequencies, since the local processing removes the high spatial frequency information in the broadband face forcing the participant to rely on the low spatial frequencies which they had already used to learn the face. Experiment 28 tested this by presenting participants with a series of high-pass, low-pass, or broadband faces to learn followed by the standard Navon identification task. During the test phase, the faces were always broadband.

5.10.1. Method

5.10.1.1. Participants and Materials

Twenty participants were selected from the participation panel at Cardiff University, all with normal vision. Participants were paid for their time. The same materials used in Experiment 28 were used in Experiment 29.

5.10.1.2. Design and Procedure

A 2 by 3 Mixed design was employed whereby participants were either instructed to identify the global Navon letter or the local Navon letter, and all learnt faces that were either high-pass, low-pass filtered, or broadband. Faces were counterbalanced, such that each appeared as a target the same number of times as it appeared as a distractor. Moreover, faces were counterbalanced such that each was learnt high-pass, low-pass, and broadband a similar amount of times. The presentation order of the faces was randomised. The experimental procedure was identical to that in Experiment 28.

5.10.2. Results

Figure 5.16 shows the mean recognition accuracy data in terms of d'. This shows that recognition accuracy was higher for global Navon processors than local Navon processors for high-pass filtered faces and broadband faces, but was lower for low-pass filtered faces. These data were subjected to a 2 by 3 mixed ANOVA with the factors: type of Navon processing (global or local) and type of bandpass filtering of the faces at learning (high-pass, low-pass, or broadband). This revealed a significant main effect of type of Navon processing on recognition accuracy, $F(1, 18) = 8.228$, MSE $= 0.307$, $p <$ .05, whereby global Navon processors had a higher recognition accuracy in the subsequent face recognition test than local Navon processors (mean difference $= 0.410$). There was also a main effect of type of bandpass filtering of the face, $F(1, 18) = 4.751$, MSE $= 0.534$, $p < .05$, that revealed itself through higher recognition accuracy for broadband faces than high-pass filtered faces (mean difference $= 0.474, p < .05$) and low-pass faces (mean difference $= 0.698, p < .05$). These main effects were qualified by a significant interaction, $F(1, 18) = 13.959$, MSE $= 0.534, p < .05$. Simple effects revealed that global Navon processors had a significantly greater recognition accuracy than local Navon processors for broadband faces (mean difference $= 1.279, p < .05$) and for high-pass filtered faces (mean difference $= 0.938, p < .05$), whereas local Navon processors had a significantly greater recognition accuracy than global Navon processors for low-pass filtered faces (mean difference $= 0.987, p < .05$).
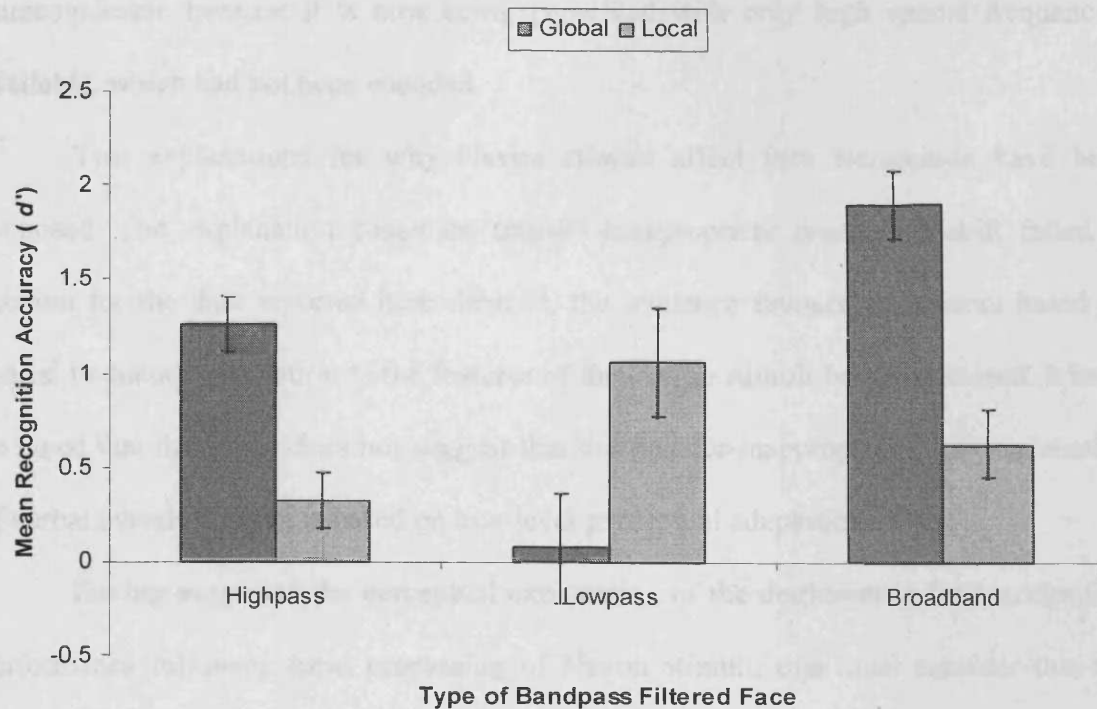
Figure 5.16. Recognition accuracy split by type of Navon processing undertaken and type of bandpass filtered faces at learning and type of bandpass filtered faces at test for Experiment 29. Error bars represent standard error.

5.10.3. Discussion

The results from Experiment 29 are consistent with those of Experiment 28. A high-pass filtered face is one that contains only high spatial frequencies. When it is learnt, its subsequent recognition is unaffected by global Navon processing, whereas local Navon processing brings its recognition down to near chance levels. Conversely, when a low-pass face is learnt, processing the global figure of a Navon stimulus brings its recognition down to chance levels. The explanation for this is that processing the global figure in a Navon stimulus lowers the discrimination of low spatial frequencies from subsequent perception. As such, a face learnt with only low spatial frequencies becomes virtually

278

unrecognisable because it is now being perceived with only high spatial frequencies available, which had not been encoded.

Two explanations for why Navon stimuli affect face recognition have been proposed. The explanation based on transfer-inappropriate processing shift failed to account for the data reported here. Instead, the evidence favours an account based on spatial frequency adaptation to the features of the Navon stimuli being processed. It must be noted that this study does not suggest that the transfer-inappropriate shift explanation of verbal overshadowing is based on low-level perceptual adaptation.

Having suggested the perceptual explanation of the detriment in face recognition performance following local processing of Navon stimuli, one must consider that the perceptual explanation does not allow for any form of inversion effect. An upright face and an inverted face contain the same spatial frequency components. Nevertheless, inverted faces were less well recognised than upright faces in all conditions in Experiments 28 and 29. These results are indicative of the face recognition system involving higher-level visual areas as well as low-level spatial frequency components. The effect Navon stimuli have on face recognition is, thus, unrelated to the inversion effect, suggesting that the mechanisms behind the inversion effect on face recognition and the spatial frequencies involved in face processing are independent of each other to a certain extent (c.f., Boutet, Collin, and Faubert, 2003; Rondan and Deruelle, 2004).

One other aspect of inversion data reported in this study is the fact that orientation at learning is more crucial than orientation at test. Indeed, a mismatch in orientations from learning to test is not as detrimental to face recognition as suggested by Bruce (1982) and Valentine (1988). Thus, these data indicate that if a face is learnt upright, it can be recognised inverted, whereas if a face is learnt inverted it is more difficult to recognise inverted. The face inversion effect is theorised to be due to a disruption of

configural information (Bartlett and Searcy, 1993; Leder and Bruce, 2000; Valentine, 1988). The present data suggests that if a face is learnt with configural information intact, it can be recognised with this information disrupted. Thus, featural information can be extracted from a face if it has been stored with configural information. However, if a face has been stored with only featural information intact, the configural information cannot be extracted from it. This original finding could be explored in more detail in a separate research project, but is independent of the effect of Navon processing on face recognition and thus is not the main focus of this paper.

The argument put forward here is that an effect believed to be cognitive in origin turns out to have a perceptual explanation, albeit one that required directed attention to inform the perceptual processes. Perception of the world is the starting point of all cognitive processes. As such, the perceptual explanations should be ruled out before applying high-level cognition explanations. In the present example, perceptual after-effects are able to explain the detrimental effect of processing Navon stimuli on subsequent face recognition, rather than relying on higher-level under-specified cognitive explanations.

5.11. General Discussion

The experiments in this Chapter explored the detrimental effect of processing the local features of a Navon stimulus on face recognition. Experiments 20 to 22 demonstrate that the effect is not based on the verbal components of Navon stimuli. Experiments 23 and 24 demonstrate that the temporal nature of the effect is such that it decreases rapidly over time or over multiple presentations of face stimuli. Experiments 25 to 29 have shown that

the best explanation for the effect that Navon stimuli have on face recognition is due to spatial frequency adaptation.

Face-space as a metaphor for explaining the encoding of faces can account neatly for the findings presented in this Chapter. Schyns et al. (2002) suggested the spatial frequency channels used for face recognition are modulated by the task in hand. This indicates that the task demands determine the spatial frequency channels to be attended to. Fourier analysis suggests that all visual scenes can be represented by a collection of spatial frequencies. As such, dimensions in face-space can be represented by a subset of spatial frequencies. The data presented by Schyns, et al. (2002) indicated that attention modulates the dimensions in face-space depending on the task in hand. Local processing of Navon stimuli thus can cause adaptation in the high spatial frequency channels that make up the most diagnostic dimensions of face-space for facial identification.

It is assumed that the high spatial frequencies are useful for featural processing and that recognition of inverted faces involves featural processing. The data from Experiments 25 and 26 indicate that the inversion effect is not based on the same spatial frequency mechanisms as the effect Navon stimuli have on face perception. This has been reported previously (e.g., Boutet, et al., 2003; Rondan and Deruelle, 2004). As such, there is a problem with the logic presented here. Either spatial frequencies are unrelated to configural or featural processing or configural and featural processing are unrelated to the inversion effect. If the definition of configural and featural processing of Lewis and Glenister (1999) is to be believed then the former posit must be false, and spatial frequencies are not related to configural or featural processing.

In Chapter 3, it was proposed that the inversion effect was, at least partially, due to attention being paid to less diagnostic visual features in the first glance. As such, an inverted face is being encoded along less diagnostic dimensions of the face-space. This,

then, acts as a different definition of configural and featural processing. The inversion effect is based on this rather than the rather simplistic terms configural and featural processing. Indeed, this allows for configural and featural processing to be related to spatial frequencies.

The data from Experiments 25 and 26 also indicate that the inversion effect is primarily one of encoding (c.f., Valentine, 1988). If a face is learnt inverted, it is coded along less diagnostic dimensions than if it was learnt upright. It seems then, if a face is learnt along less diagnostic dimensions it is difficult to extract the more appropriate dimensions from the face. However, if a face is learnt upright and thus along the more appropriate dimensions, it is possible for face-space to extract less diagnostic dimensions for recognition purposes. This again suggests that the face-space is far more flexible than originally presupposed.

There is one caveat regarding the present discussion. If the explanation for the effect that Navon stimuli have on face recognition is based upon spatial frequency adaptation then one would assume other types of stimuli made up of only high spatial frequencies would cause similar effects. In Experiment 22, it was demonstrated that bandpass filtered faces containing only high spatial frequencies did not cause a detriment in subsequent face recognition performance. This does not seem compatible with a spatial frequency after effect explanation of the effect Navon stimuli have of face recognition. However, there are several possible explanations for this. It may be that the semantic information contained within the faces bypasses low-level spatial frequency channels. Alternatively, since the filtered faces in Experiment 22 are famous faces and the test faces are unfamiliar faces, this may be the cause for the null effect. Famous faces are recognised using a different set of spatial frequencies than unfamiliar faces.

Unfortunately, these explanations appear weak and thus may simply reflect the relatively unstable nature of this effect.

Although the experiments presented in this Chapter are an aside from the other experiments presented in this thesis, it is worth noting that the three models of face-space presented in Chapter 1 can account for these findings based on two possible assumptions. There may be dimensions of face-space that are configural dimensions and may be dimensions of face-space that are featural. The configural dimensions are those used by default. Thus, as described in Section 3.9., constant-face-space (as described by Valentine, 1991), expanding-face-space, and shrinking-face-space can be applied to account for these findings. Alternatively, the dimensions of face-space may represent particular spatial frequencies. Adaptation may, thus, affect particular dimensions of face-space, or the whole space in a manner described in Section 4.8. Again, the three models of face-space can thus explain the flexibility of the face recognition system in terms of how the dimensions of the multidimensional space are used.

Conclusions

Ten experiments reported in this Chapter demonstrate that the detrimental effect of processing the local features of Navon stimuli on subsequent face recognition performance is based on spatial frequency adaptation. This novel finding also indicates that the inversion effect in face recognition is not related to the same types of configural and featural processing that spatial frequencies are. As such, face-space has been used to offer a new explanation of the inversion effect. Moreover, it has been suggested that dimensions of face-space are related to spatial frequency analysis.

Chapter 6 - Conclusions

In this Chapter, these apparently disparate lines of research will be drawn together and summarised. Key results will be elaborated upon and the data will be interpreted within face-space. This will inevitably lead to a detailed discussion and laying the assumptions and foundations of the developmental model, shrinking-face-space and the resulting adult version "flexible-face-space". Finally, several unanswered questions will be posed and lines of future research proposed.

6.1. Research Summary

In total, 29 separate experiments were conducted exploring the nature of face-space. The experiments presented in Chapter 2 examined the nature of face recognition in children. Children were shown to have a smaller inversion effect and ORB to adults. Moreover, children were adaptable to and able to learn "unnatural" facial configurations to a far greater degree than adults. The evidence collected indicated that the immature face-space contains many more dimensions than the adult face-space, allowing children to process faces that adults would find more difficult to process. It was hypothesised that these dimensions could be used by adults under certain circumstances.

The experiments reported in Chapter 3 indicate that priming, extensive training, and attentional cuing are methods that can alter which dimensions used to encode faces in adults. Stereotype labels presented during learning of the faces improved face recognition of stereotype-congruent faces. The ORB was reduced by training participants to focus on

facial features more diagnostic in the coding of other-race faces by using a perceptual learning regime and by fixation crosses preceding the lower half of the face. The effect of the fixation cross was removed if the participants were forced to delay their responses by more than 1 second. Fixation crosses preceding the forehead allowed White participants' recognition of White faces to be highly accurate irrespective of the orientation of the face. The effect of fixation crosses was only observed during initial learning of the faces.

Chapter 4 used a different methodology to explore the flexibility of face-space in terms of how the dimensions of face-space are used. Adaptation was used to recalibrate the attentional weighting to particular dimensions of face-space. It was demonstrated that these visual dimensions could be recalibrated using non-visual adaptation – the presentation of a name, voice, or imagination caused adaptation to face stimuli for familiar faces, but only visually induced adaptation was observed for unfamiliar faces. This adaptation was shown to last at least two months, which is indicative of permanent changes in the face recognition system due to adaptation. Moreover, the after-effects would transfer from inverted adaptors to upright test stimuli, but not the other way round.

The final experimental chapter, Chapter 5, looked the detrimental effect of processing the local features of Navon stimuli on face recognition. It was shown that this effect occurred for Navon letters and shapes, and would last over few trials before the effect faded away. The final experiments explored what happens to this effect when the Navon stimuli or the face stimuli had their spatial frequency or contrast altered. When low contrast or blurred Navon stimuli were used, the detrimental effect of local processing was not observed. Moreover, the detrimental effect of local processing was found for processing of single high-contrast letters. Local processing aided the recognition of faces learnt containing only low spatial frequency.

These basic findings will be detailed in relation to the major effects of face recognition: the inversion effect, the own-race bias, and adaptation. The data reported in this thesis has major implications for all three of these effects, and many of these findings are novel findings. As such, an extensive analysis of some of the critical findings shall be presented here.

6.1.1. The Inversion Effect as Shift in Processing Style

Experiment 1 tested the inversion effect in children aging from 5-years-old to adults. The findings were broadly consistent with those of Fagan (1979) and Flin (1983), in which the magnitude of the inversion effect increases with age. However, these data were unable to rule out a shift in processing at age 11 to 13 as posited by Carey and Diamond (1977) in contrast to a gradual increase in the magnitude of the ORB.

However, consistent with research by Itier and Taylor (2004), the adult magnitude of the inversion effect appears before the adult levels of absolute recognition abilities (approximately ages 11- and 15-years old respectively). This indicates that experience causes face recognition abilities to improve. This same experience causes the face recognition system to become less flexible before adult levels of performance are reached. Thus, there is some evidence for a change in face-space at age 11 to 13 that may be related to configural and featural processing.

The evidence presented in this thesis indicates that children have a more flexible face recognition system with more dimensions than adults have. The developmental shift may appear to be from featural to configural processing, but may actually be a reduction in the number of dimensions of face-space that are being attended to. In this way, featural dimensions are being discarded and more configural dimensions have more attentional

weighting paid to them. This would explain why there is an apparent change in processing at age 11 to 13, while maintaining the increase in face recognition abilities, as more dimensions can be discarded or changed, and general cognitive abilities increase.

Chapter 3 also demonstrated effects of inversion and how they could be reduced by the presentation of a fixation cross preceding the face. The location of the fixation cross was critical for the recognition of White faces, whereby recognition accuracy was higher when the fixation cross preceded the forehead than when it preceded the chin. This pattern of results was replicated in three separate experiments, using three stimuli sets, and three sets of participants, with slight changes to the methodologies. Thus, this effect appears to be reliable.

The presentation of a fixation cross prior to a stimulus provides the participant with a cue as to where they should first look. Eye-tracking studies have demonstrated a standard pattern of eye movements when looking at a face. These eye patterns start with glances to the forehead, around the eyes, with few glances to the nose and mouth (Althoff and N. Cohen, 1999; Cook, 1978; Stacey, et al., 2005). As such, a fixation cross preceding the forehead will be consistent with the standard first glance at a face; whereas a fixation cross preceding the chin will be inconsistent with the standard first glance at a face.

Since this cuing paradigm was able to reduce the magnitude of the inversion effect, it can be hypothesised that part of the inversion effect in face recognition is due to the attention paid during the first glance at a face. This is a novel hypothesis for the inversion effect and suggests an alternative to the configural/featural debate in the processing of inverted faces. This attentional mechanism is not dependent on types of processing for it to work. Instead, it suggests the first feature glanced at in a face affects

accuracy. If this is diagnostic (forehead in White faces), accuracy is higher than if it is not diagnostic (chin in White faces).

There is a major caveat with this explanation and that relates to the fact that a similar result for the ORB was not observed if the participants were forced to delay their responses. This means that though the fixation cross guides where participants first look, if they have time, they will revert to their usual scan. This will over-ride the benefit of the first look. Thus, the importance of the first glance is qualified by subsequent glances that are in addition to the first glance or counteract its benefit.

The effect of this cuing on the inversion effect was only observed when the fixation cross was presented during the original encoding of the face. This is a second aspect of the inversion effect that has not previously been reported. The attention paid during the original encoding of a face is more important than the attention during the recognition of the face. However, as mentioned above, delaying the responses removes the effect of the fixation cross and the test phase was self-paced. Thus, participants may have been delaying their own responses allowing time for their usual scanpath to take over. This seems unlikely since all responses in the test phase were under 1200 ms and the effect of the fixation cross is present up to delays of 1000 ms to 2000 ms. Thus, this effect highlights the importance of encoding conditions on face recognition (e.g., Coin and Tiberghein, 1997) and in particular the inversion effect.

Experiments 25 and 26 also tested aspects of the inversion effect and found the importance of inversion at learning to be more important than the effects of inversion at test. This result is consistent with the findings of the importance of the attention paid during the initial encoding of an inverted face discussed above. Thus, rather than an orientation mismatch causing a recognition deficit, the deficit is an inversion effect at face learning, which indicates that the inversion effect is an encoding phenomenon (c.f.,

Valentine, 1988). A face that is encoded inverted has a poorer storage in memory, possibly due to the lack of appropriate dimensions in face-space to code onto. Thus, its subsequent recognition will be more difficult. However, if a face is already stored in the face-space then dimensions can be extracted from it that fit with an inverted face. This is akin to a mental rotation within the face-space for an already stored face (c.f., Valentine and Bruce, 1988).

Evidence from Experiments 25 and 26 is suggestive of a further problem with the definitions of configural and featural processing. Since these types of processing are sometimes linked to particular spatial frequency channels, if the inversion effect is based on the same mechanism it would also be linked to spatial frequency channels. It was found that the inversion effect was unrelated to spatial frequency (see also, Boutet, et al., 2003; Nagayama, Yoshida, and Toshima, 1995; Rondon and Deurelle, 2004) suggesting that inversion does not disrupt the same configural processing that is processed by low spatial frequencies.

Face-space can be used to explain why spatial frequency is unrelated to the inversion effect but is related to configural and featural processing. The suggestion is that an inverted face is encoded onto the face-space using a different set of dimensions than an upright face and that these are less appropriate for the recognition of faces. These dimensions are unrelated to distinct spatial frequency channels, whereas there are other dimensions that are related to distinct spatial frequency channels.

In summary, though the inversion effect was not a critical part of this thesis, it has proven to reveal some highly interesting effects. Primarily, the importance of attention paid at encoding indicates that inversion is mainly detrimental due to the first glance and resulting scanpath over a face. Indeed, there is not an interaction between orientation of a face at learning and at test on the magnitude of the inversion effect, suggesting that

encoding is the most critical period for the inversion effect. Moreover, the independence of spatial frequencies and inversion (both are said to be linked to configural and featural processing) has led to new interpretations of the inversion effect based upon attentional encoding onto less diagnostic dimensions of face-space.

6.1.2. The Own-Group Biases

Several results observed in the experiments reported in this thesis further our understanding of the own-group biases in face recognition. Sporer (2001) posited an in-group/out-group model to explain the own-race bias, in which humans recognise faces of the in-group better than the out-group. This difference is likely to be due to individuating processing employed when looking at own-group faces (Levin, 2000). The data in the present experiments offer some suggestions as to what the individuating and race processing features may be.

Experiment 6 tested the ORB and indicated the possible dimensions within face-space that are appropriate to recognise Black faces. The supposition, based on a carry-over from a perceptual learning task, was that face encoding using the nose, mouth, and chin would lead to improved recognition in Black faces and a detriment in recognition performance in White faces. This suggests that the ORB is due to a failure to attend to and encode the dimensions of faces that relate to the nose, mouth, and chin. In other words, the dimensions useful for individuating Black and White faces are different and participants tend to use the own-group dimensions causing the other-race faces to be processed according to group.

Experiments 10 to 12 explored the attentional effects of the ORB using cuing paradigms. It was shown that the ORB could be reduced and even reversed if White

participants' first glance was to the lower portion of a face (namely the chin). This effect was one based on the encoding of the face, indicating the importance of encoding on face recognition. Moreover, the results indicated that the first glance at a face was crucial for determining how well it would be recognised and required participants to attend to the most diagnostic features of a face for most accurate face encoding. For the recognition of White faces this first glance is in the upper portion of the face, namely the forehead and eyes. For the recognition of Black faces, this first glance is in the lower portion of the face, namely the chin and mouth.

This cuing effect is not easily accounted for in most explanations of the ORB, but can be explained simply in the face-space. Attention directs which dimensions are to be used for the encoding of the face. If these dimensions are diagnostic, then face recognition will be accurate. If, however, these dimensions are not diagnostic, then face recognition will be less accurate. As such, attention directs which dimensions of face-space are to be used and this directly relates to the accuracy with which faces will be recognised. Attentional weighting (c.f., Nosofsky, 1986) is thus an important and vital role in face-space and the recognition of faces.

Like the discussion of the inversion effect, there is a major caveat with the theoretical descriptions offered here. The location of the preceding fixation cross to improve the recognition of Black faces is at the chin area and the inversion effect observed in White people may in part be due to the first glance not being to the most diagnostic feature for White faces. In fact, when White people are presented with an inverted face, their first glance is to the chin (as it is at the top of the screen). Thus, if the inverted face is a Black one then the results of the experiments presented here suggests that the recognition accuracy of inverted Black faces by White participants should be greater than the recognition accuracy of upright Black faces. This is not the case - as

Valentine and Bruce (1987) noted that the inversion effect is found for Black and White faces in White participants. To explain this apparent contradiction, one must return to the data presented in Experiment 9 in which the effect of the preceding fixation cross was removed if the participants were forced to delay their responses. This suggests that the first glance is overridden by the default scanpath. Thus, it might be that the scanpath over a face has substantial differences for Black and White inverted faces which alters the dimensions that are used for encoding. Further work is required on this topic before any firm conclusions can be drawn.

6.1.3. Adaptation

Several experiments were conducted exploring the nature of face after effects following adaptation. Both FDAEs and FIAEs were looked at in this thesis and some interesting findings were observed. These findings indicate that face after effects are not necessarily caused solely by visual input and are based on some kind of expertise and experiential components.

Firstly, Experiment 4 replicated findings from Robbins, et al. (2007) that some more "natural" facial configurations are more adaptable in adults than "unnatural" facial configurations. This indicates that the recalibration along dimensions of face-space (Hurlbert, 2001) occurs for more experienced facial configurations rather than the less experienced configurations. In children, however, adaptation to "unnatural" facial configurations is possible. This implies that the neural tunings for facial processing change with age (Nelson, 2001). Furthermore, it suggests that the neural tunings get less flexible with development, similar to neuronal pruning.

Related to this are results from the final experiment in Chapter 4, Experiment 19, in which it was shown that FIAEs for familiar faces cause permanent (at least lasting 2 months) changes in the representation of the face. Adaptation, therefore, may be a possible mechanism for learning about faces and in particular the expert nature of face processing. Repeated presentation of faces may cause the neural representations to become more precise and specific to a set configuration - possibly due to neuronal pruning. To explain this neural pruning, a neural account of adaptation needs to be explored.

Robbins et al. (2007) provided a neuronal account of the FDAEs, which is presented in Figure 1.3. This two-pool model shows how adaptation to a particular class of face causes a lower response in the neural pool that codes that class, shifting the perceived midpoint. Such an account explains the adaptation for dimensions in face-space that have two polar ends. This explanation can be developed to explain the results of Experiments 3 and those in Chapter 4 in different ways.

To explain the developmental plasticity of the neural responses to "unnatural" facial configurations (Experiment 3), one can assume that there may be neural pools similar to those presented in Figure 1.3., but for left eye high to left eye low and for right eye high to right eye low. The explanation of adaptation presented in Figure 1.3 can thus explain adaptation to "unnatural" facial configurations observed in children. The reason why adults do not show adaptation to "unnatural" facial configurations is due to the neural pools to "one eye high to low" becoming unused. Unused neural pools become "pruned" and their use is altered, they may be removed or other critical change to them. This ensures that the brain does not get cluttered by pools of neurons that have no use in perceptual experience.

To explain the effects presented in Chapter 4, whereby non-visual information can affect the perception of faces, one has to slightly modify the neural presentation shown in Figure 1.3. It is difficult to envisage a dimension from one facial identity to an opposite as presented in Figure 1.3. However, it is possible to envisage a particular facial identity as being made up of several neural pools, which each pool having an opposite value. Added together, a face identity will have a neural pool and "all other faces" similar to that presented in Figure 1.3. These neural pools combine to form one identity and are made up of visual dimensions and non-visual neural pools. Thus, the neural pool for a particular facial identity can be fatigued by visual stimulation and by non-visual stimulation for those people who have good cross-modal neural connexions. This idea is presented in Figure 6.1.

Taken together our understanding of these three key effects has been greatly increased by the findings presented in this thesis. These findings have demonstrated how face-space is unable to account for some of these effects in its present form. As posited in the introduction, an alternative version of face-space, theorised in the introduction, can be applied. The next section describes key assumptions and differences between the newly devised "flexible-face-space" and previous versions of face-space.
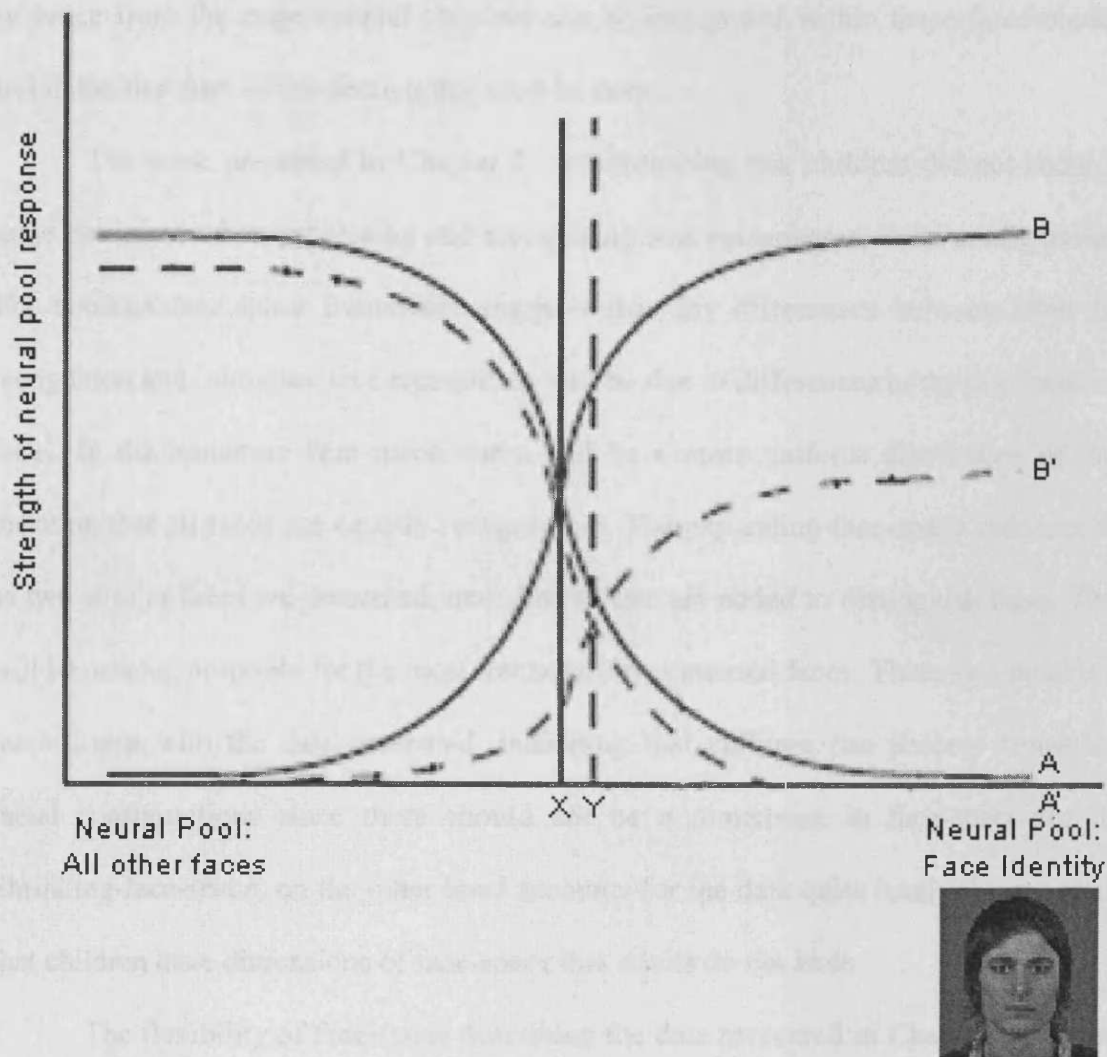
Figure 6.1. Adaptation to a particular face caused the midpoint shift between that face and any other face to be moved from X to Y. The neural response from pool to the facial identity reduces from line B to B'. The neural pool for the face identity can be activated by non-visual identity-specific information.

6.2. Flexible-face-space

In the introduction, three possible models of the immature face-space were hypothesised: the constant-face-space, the expanding-face-space, and the shrinking-face-space. The

295

evidence from the experimental chapters can be interpreted within these three models, and in the first part of this section this shall be done.

The work presented in Chapter 2, demonstrating that children did not show the same detriment when processing and recognising less encountered facial configurations. The constant-face-space framework suggests that any differences between adult face recognition and immature face recognition will be due to differences in the distribution of faces. In the immature face-space, there will be a more uniform distribution of faces meaning that all faces are equally recognisable. The expanding-face-space indicates that as two similar faces are presented, new dimensions are added to distinguish them. These will be more appropriate for the most frequently encountered faces. These two models are inconsistent with the data presented, indicating that children can process "unnatural" facial configurations since there should not be a dimension in face-space for this. Shrinking-face-space, on the other hand accounts for the data quite neatly, by suggesting that children have dimensions of face-space that adults do not have.

The flexibility of face-space describing the data presented in Chapters 3, 4, and 5 is different from the flexibility described above for the data in Chapter 2. Chapter 2 described flexibility in the nature and number of dimensions, whereas Chapters 3, 4, and 5 described flexibility in terms of how and when the dimensions of face-space are actually used. This type of flexibility is not explicitly discussed in previous models of face-space (e.g., face-space-r, Lewis, 2001), however can be easily interpreted within all three models of face-space that were considered here. Attentional weighting (c.f., Nosofsky, 1986, Valentine and Endo, 1992) can alter which dimensions are used for encoding for a particular task. Thus, stereotype priming, training, and fixation crosses can simply alter when and in which order dimensions are used for encoding of faces. This has

not been made explicit in previous models of constant-face-space and expanding-face-space but certainly can be incorporated within these models.

There are some subtle differences in the description of dimension flexibility for the data presented in Chapter 4. Here neural models of dimensions were presented explaining the adaptation results. These neural models were not explicitly described in previous models of face-space but are quite consistent with the architecture of constant-face-space, expanding-face-space, and shrinking-face-space. This flexibility is in terms of how each particular dimension of face-space is used.

The three models of face-space presented in Chapter 1 can thus be applied to explain all of the data in this thesis. Shrinking-face-space accounts for the data in all the experimental Chapters, whereas constant-face-space and expanding-face-space presents a less convincing explanation for the developmental data presented in Chapter 2. As such, the architecture of shrinking-face-space is considered to be the have the most explanatory power of the three models and is thus the basis for "flexible-face-space" to be described in more detail here. This model makes novel predictions regarding the number and types of dimensions contained within face-space.

6.2.1. The Number of Dimensions in Shrinking-Face-Space

The evidence presented in Chapter 2 indicates that the number of dimensions that the immature face-space contains is many more than that of the adult face-space. Lewis (2004) suggested that the adult face-space contains between 15 and 22 dimensions and figures in this region have been reported by other authors (e.g., Burton and Vokey, 1998). Since the immature face-space has a greater number of dimensions than the adults', all that can be said is that the immature face-space has more than 22 dimensions. No limit on

the number of dimensions can be indicated here. Indeed, it is possible that children can use an almost infinite number of dimensions to recognise faces.

Of crucial importance to the theory of shrinking-face-space is what happens to these dimensions. It is suggested that with development, the number of dimensions attended to decreases. There are several possibilities about what happens to the dimensions that are no longer diagnostic. The first possibility is that the dimensions are lost from face-space entirely. These dimensions can never be recovered for they have proven to be completely useless. The second possibility is that the dimensions remain in face-space but are not used for the encoding of most faces. However, under certain circumstances, they can be used.

Thus, there are three types of dimensions that shrinking-face-space contains. Dimensions in shrinking-face-space can be: active and used to encode all types of faces; dormant, whereby they remain in the face-space, but are largely unattended to; and finally, extinct, in which the dimensions that existed in childhood are no longer present. It can not be hypothesised why or how a dimension becomes extinct rather than simply dormant. However, it can be hypothesised why a dimension becomes dormant or extinct rather than remaining active. This is covered in the next subsection.

The basic tenet of shrinking-face-space is that the immature face-space contains a virtually infinite number of dimensions. With development these become active ones, dormant ones, or extinct ones. Dormant ones can be reactivated for particular tasks. Extinct ones can never be recovered, and active ones are the default subset used to encode the majority of faces and there is likely to be 15 to 22 of these dimensions (Lewis, 2004).

6.2.2. Learned Attention and Perceptual Narrowing/Learning

One plausible explanation for how the dimensions in the immature face-space become dormant is through perceptual narrowing and learned attention. In the introduction a brief summary was provided concerning these topics. As such, they shall not be repeated here. The basic premise of these effects is that with development comes the ability to properly shift attention and guide it to the most diagnostic parts of the visual field. Perceptual narrowing indicates that with experience, people tend to become experts in ignoring irrelevant features and thus discriminating objects and stimuli they have much experience with by focusing on more diagnostic features.

Applying these tenets to face-space, it is indicative that it is unfocused attention in children that causes them to attend to many more dimensions than is actually required to process faces. Due to this overload in attention, children are less good at all forms of face recognition tasks. However, this defocused attention means that they are more likely to be able to process "unnatural" facial configurations and faces less frequently encountered. This unfocused attention characteristically gradually becomes more focused not reaching an asymptote until roughly age 13. This is further, albeit correlational, evidence that attentional focus is important for the development of face recognition. Nevertheless, it suggests that attentional direction is the cause for dimensions to become dormant or extinct. It must be noted that defocused attention, as an explanation of why children have lower face recognition performance than adults, is not unequivocal. It may be possible to test this using the fixation cross technique (Experiments 7 to 12) to focus children's attention artificially. Further research is thus required before the defocused attention explanation offered here is uncritically accepted.

Perceptual narrowing suggests that with experience, people will only encode the most relevant perceptual cues. They will virtually lose the ability to discriminate less relevant perceptual cues. Experience of the visual environment directly relates to the way the world is perceived. As such, as children process upright, own-race faces, of a particular configuration, they will learn to attend to the most diagnostic visual features. This attention will cause them to ignore and inhibit dimensions that are less diagnostic. Hence, this is indicative of the importance of attention on the coding into shrinking-face-space. The next subsection explores this in a little more detail.

6.2.3. Attention

Several studies presented in this thesis have demonstrated that the initial glance paid to a face is most crucial for its subsequent recognition providing that there is not sufficient time to override this. Furthermore, these studies have indicated that attention paid to particular regions will cause faces to be recognised to varying degrees of accuracy: for example, attention paid to the lower facial features causes Black faces to be better recognised than White faces. Attention paid to the lower facial features in a White face causes it to be less well recognised. However, given long enough, participants will override this beneficially attentional track with a more stereotyped one. These results indicate another important facet about flexible-face-space that needs to be addressed. That is, the order with which features in a face are attended to affects their encoding and subsequent recognition.

There are two possible explanations and ways of incorporating attentional order into flexible-face-space. Firstly, there may be a temporal effect in the processing of dimensions of face-space, where by default dimension-A is always processed before

dimension-B. This fits with the standard eye-tracking data (Stacey et al., 2004). If

attention is drawn to dimensions other than dimension-A in the first instance, the face

recognition system may be confused and cause increased error (see Matrix 6.1). Such an

attentional order effect has been observed for non-facial stimuli (Frey, Konig, and

Einhauser, 2007). One must also include an override mechanism to explain why

participants given longer to respond will show no such benefit of initial glance.

Matrix 6.1.

Accuracy of encoding is dependent on the correlation between the attentional order of the

incoming face and the default attentional order of dimensions.

	Incoming Face Attentional Order		
Default Attentional Order	Good Encoding	Poor Encoding	Very Poor Encoding
Dimension A	A	B	D
Dimension B	B	A	C
Dimension C	D	D	B
Dimension D	C	C	A

One alternative is that rather than there being a temporal effect, there may be an

inherent mathematical feature of the flexible-face-space whereby the first dimension

attended to is given more weighting than the second dimension attended to. If this first

dimension is one of those that is critical for recognition then the accuracy with which the

face will be encoded will be higher than if the first dimension is not critical for

recognition. This suggestion requires a greater elaboration on the attentional weighting

mechanism (c.f., Nosofsky, 1986) first posited in face-space by Valentine and Endo

(1992) and is presented in Matrix 6.2.

Certain dimensions in face-space have a greater default attentional weighting. These are the dimensions that are most relevant to the face processing task at hand. Using the face scanning evidence (Stacey et al., 2004), two possible dimensions that are likely to have higher attentional weighting are the eye colour and the distance between the eyes and the hair line, whereas the distance between the mouth and the chin is likely to be a dimension that has a lower default attentional weighting in White participants. These default attentional weightings can be changed (shown in Matrix 6.2.) depending on the task.

These attentional effects in flexible-face-space make this model of face-space inherently more flexible in terms of number and nature of dimensions than previous models of face-space and offers an account of flexibility of how and when dimensions are used. Dimensions can be attended to, ignored, or even partially attended to. If the dimensions to be attended to are diagnostic and are attended to first then recognition will be most accurate. However, if the dimensions attended to are not diagnostic, or are not attended to first, then recognition will be less accurate.

<u>Matrix 6.2.</u>

Top left panel: Default attentional weighting given to four dimensions for two different tasks. Top right panel: The attentional weighting given to the dimensions attended to depends on the order in which they are attended to. Bottom panels: The quality of encoding the face depends on the correlation between the order of dimensions encoded and the default attentional weighting for Task 1 (left panel) and Task 2 (right panel).

	Default Attentional Weighting	
Dimension	Task 1	Task 2
A	1	.5
B	.75	1
C	.5	.75
D	.25	.25

Attention Order	Weighting
1st Feature	1
2nd Feature	.75
3rd Feature	.5
4th Feature	.25

Task 1:

Presentation Order (Quality of Encoding)		
Good	Poor	Very Poor
A	B	D
B	A	C
C	D	B
D	C	A

Task 2:

Presentation Order (Quality of Encoding)		
Good	Poor	Very Poor
B	C	D
C	B	A
A	D	C
D	A	B

6.2.4. Types of Dimensions

Though it is difficult to actually define what the dimensions of face-space are, and even more for shrinking-face-space, there can be some suggestions made. One suggestion that was made by McKone et al. (2005) is that the dimensions of face-space are related to the

population variability in facial characteristics. As such, the dimensions in the adult face-space relating to the eyes will include both eyes shifted up rather than one eye shifted up. This has received empirical support here.

Another class of dimension that can be hypothesised from data in this study is that some dimensions may be classifying dimensions. For example, gender may be a dimension of the face-space used to classify faces. This may differ from the subsequent dimensions that may reflect physiognomic characteristics such as the distance between the eyes and the hairline.

Finally, the visual basis for the dimensions of face-space may be based on spatial frequency analysis and in particular Fourier analysis. This suggestion is entirely plausible, given the findings from the experiments presented in Chapter 5. Removal of certain spatial frequencies changes the accuracy with which faces will be encoded. High-pass faces are likely to be encoded along a set of dimensions that differ from low-pass faces due to the difference in information that they contain. Thus, a high-pass face encoded into face-space will be more difficult to recognise when being compared to a low-pass face due to the fact that the stored face has not got values on the dimensions that the low-pass face is being encoded on.

In summary, the tenets of flexible-face-space are that the dimensions in the adult face-space are not equally used. The attentional weighting to each dimension of face-space is different from other dimensions and can be altered depending on the task at hand, and the cognitive and perceptual state of the participant. These dimensions are related to spatial frequencies in a way that cannot be fully explored in the present thesis.

6.3. Future Directions

The basic principals of shrinking-face-space have been established in the previous section. However, there are some unanswered questions remaining. Theoretically, the most interesting question is what the earliest form of face-space looks like and whether it comes from a face specific mechanism or develops from more general perception. To explore these questions, a theoretical piece regarding pre-immature face-space is presented here, followed by a discussion of how face-space may relate to object recognition. Finally, the possibility of computational and neurological modelling will be put forward.

6.3.1. Pre-immature Face-Space

E. de Haan (2001) was one of the few authors to suggest when face-space develops. According to him, face-space does not develop until the age of 3, however de Haan does not indicate why face-space develops or from where it develops. Prior to age 3, infants are thus encoding faces according to some non-face-space face coding mechanism, or faces are encoded as objects. However, there is sufficient evidence that faces are coded differently from objects prior to this age. To describe this, face recognition abilities in newborns up to 3-year-old infants shall be described and how this may relate to the development of face-space shall be theorised.

6.3.1.1. Neonatal Visual Preference Research

Perhaps the most commonly used method for examining face processing in neonates is the use of the preferential tracking procedure. Fantz (1963) presented newborns with a

non distorted picture of Joan Crawford's face, an internally scrambled version, or a totally scrambled version of the same face. Though all these face-like stimuli were preferred to coloured discs or other patterns, there were no differences between the levels of distortion. Goren, Sarty and Wu (1975) presented newborn babies (median age 9 minutes) a series of patterns. The newborns tracked the face-like pattern more than the scrambled pattern more than the blank pattern (see also Fitzgerald, 1968; Stechler, 1964).

Dziurawiec and Ellis (1985), Easterbrook, Kisilevsky, Hains, and Muir (1999), Haaf (1974), Hershenson (1964), and Thomas (1965) failed to replicate this preference in newborns as young as the first hour of birth. However, Johnson, Dziurawiec, Ellis, and Morton (1991) demonstrated that newborns within the first hour of birth would track a schematic face stimulus more than a scrambled schematic face and a blank template with their eyes and head. A second experiment by Johnson, et al. (1991) demonstrated that there was a significant difference for eye turning for schematic face over inverted configuration, and linear scramble. However, their analysis was based on the Wilcoxon ordinal test on degree of eye turning, which is interval level data. Using the more appropriate parametric statistic would not have revealed some of their effects. Moreover, for head turning, there was no significant difference between any of the stimuli.

Kleiner (1987) used four types of stimuli to test newborns' preference for face-like patterns: a standard schematic face; a standard lattice; a face with the amplitude spectrum of the lattice, but the phase spectrum of the face; and a face with the amplitude spectrum of the face and the phase spectrum of the lattice. The newborn babies looked more at the face than the lattice, and looked more at the amplitude spectrum of the face than the phase spectrum of the face. However, the newborns did look at the face more than the amplitude spectrum of the face.

Salapatek and Kessen (1966) demonstrated that newborns tended to direct their attention to the vertices of a triangle rather than the base or corners, indicating early edge detectors rather than shape knowledge. Salapatek (1968) reported that newborns' horizontal scanning was more dispersed than vertical scanning, which may indicate the presence of innate neurons for vertical edge detection (Kessen, Salapatek, and Haith, 1972). With more horizontal scanning, vertical edges attract more attention in the newborn (Kessen et al., 1972).

The visual acuity of neonates is much lower than that of adults. Indeed, Banks and Salapatek (1981) have shown that the contrast sensitivity function (CSF) of newborns is predictive of pattern preference. While at lower ages, low spatial frequencies occupy the peak sensitivity of the infant CSF, older infants have peak sensitivity at higher spatial frequencies (Atkinson and Braddick, 1980; Atkinson, Braddick, and Moar, 1977). Morison and Slater (1985) demonstrated that neonates' preference for high-amplitude spatial-frequency close to the peak sensitivity of the CSF was mediated by contrast, in that as contrast declined so did preferential looking.

In relation to the face tracking designs, a note must be made concerning the newborns' visual acuity. A newborn's peripheral visual system is more mature than its foveal system (Teller, McDonald, Preston, and Sebris, 1986), meaning that in the fovea, contrast discrimination of newborns is such that they can distinguish grey stripes of 30% contrast, whereas adults can detect 1% (black to white is 100%, all grey is 0%). Newborns look at stimuli with high contrast regions and have low spatial frequencies (Morison and Slater, 1985).

Visual acuity is at newborn levels immediately post removal of the cataracts. However, one hour after visual input, the visual acuity of infants was almost twice as good as it was prior to cataract treatment (Maurer, et al., 1999). This suggests that visual

acuity improvement requires visual input. Moreover, the 5 fold increase in visual acuity in the first six months of life (Mayer and Dobson, 1982) is due to visual input irrespective of maturation.

In terms of eye-tracking when looking at faces, it has been noted that the primary feature that newborns look at is the mouth. This changes quickly with development. The eyes are the most scanned visual feature in the second month of life (Maurer, 1983). Indeed, infants follow adults' gaze direction by 3- or 4-months of age (Farroni, Johnson, Brockbank, and Simion, 2000; Hood, Willen, and Driver, 1998). Emery (2000) noted that the eyes are particularly important for face processing. Blass and Camp (2004) conducted a rather elaborate study whereby they familiarised 8- to 21-week old infants to an experimenter's face either wearing a long or short wig. The experimenter left the room and returned a short while later. The infants' reactions were measured toward the experimenter who was either the same or different person wearing the same or the different wig. The infants' looking times were greatly increased when the different experimenter returned with the different wig. Such a result was interpreted as the infants requiring both internal and external features to match their representation for a face to be familiar.

6.3.1.2. Accounts of Face Tracking

Three accounts exist for the neonatal face tracking effects observed. The first account for very early tracking abilities in neonates is perhaps the most parsimonious by being the one able to account for the majority of effects. It was put forward by Banks and Salapatek (1981), Kleiner and Banks (1987) and Kleiner (1993) and is termed the sensory hypothesis based upon the linear systems model of infant perception. Without going into too much detail about this model, the model implies that infant perception is driven by

spatial frequency and contrast information relating to the CSF of newborns. All visual percepts contain information concerning the amplitude and phase spectrum of the spatial frequencies. Due to the immature nature of the newborns' visual systems, images that have high contrast and lower spatial frequencies are more visible to newborns. Faces are stimuli that have high contrast and low spatial frequencies and thus are more visible to newborns.

A development of the basic sensory model was put forward by Acerra, Burnod, and de Schonen (2002) in which it has been suggested that the newborn responds to top-heavy images more than bottom heavy images. This explains the findings from tracking designs that indicate a preference for face-like stimuli over the same features scrambled. These two models are loosely based on the visual system and identifiable structures within the visual cortex. Two further models have no real basis in the visual system but have been cited far too frequently.

A structural account of early tracking abilities was put forward by Morton and Johnson (1991) in which infants are born with the knowledge of the configurations of faces. In other words, there is inborn first-order relational knowledge within the child. Morton and Johnson call this CONSPEC. They indicate that this is a subcortical process, since the subcortices are much more developed than the cortex, which is required for subsequent learning one or two months after birth. This innate perceptual knowledge exists prior to any perceptual experience.

The internally generated activity model (cited in Bednar and Miikkulainen, 2003) is based on spontaneous neural activity in the visual centres of the brain (Feller, Wellis, Stellwagen, Weblin, and Shatz, 1996; Khazipov, Sirota, Leinekugel, Holmes, Ben-Arf, and Buzsaki, 2004; Leinekugel, Khazipov, Cannon, Hirase, Ben-Arf, and Buzsaki, 2002; Lippe, 1994; O'Donovan, 1999; Wong, 1999a, b; Wong, Meister, and Shatz, 1993) that

are correlated with dreams during REM sleep (Callaway, Lydic, Baghdoyan, and Hobson, 1987) and imagery (Marks and Isaacs, 1995) and are occurring in newborns and pre-natally (Shatz, 1990, 1993, 1996a, b). The implication of this model is that internally generated images pre-natally are of faces. As such the visual experiences of infants are predisposed due to dreams which are genetically influenced. This explains how the CONSPEC of Morton and Johnson (1991) can occur pre-birth. However, it requires that dreams are predetermined by genes.

There is a prolonged debate on the viability of an innate model of face preference (e.g., Morton and Johnson, 1991) compared to face preference being based on visual abilities in the newborn (e.g., Banks and Salapatek, 1987). These visual abilities are based on visual acuity, visible contrasts and spatial frequencies, and ocular muscle control in newborns. Either way, these models give a basis for the development of face-space.

6.3.1.3. Before Face-Space is CONSPEC?

Before face-space is CONSPEC, Morton and Johnson (1991) would argue. This innate structural representation of face is a knowledge of first-order relational information. Thus, first-order relational information could be the first stage of face-space. In other words, the first dimension of face-space is that eyes appear above a nose and mouth. This means that face-space does not have to be qualitatively different from CONSPEC. The idea here is that, as face-space develops, the reliance on this simple structural knowledge can become less and more dimensions can be added.

If CONSPEC is not an accurate assessment of what is occurring before face-space, and simple sensory explanations are the best way of accounting for the early face-tracking results, then before face-space is vision. In other words, face-space comes as a

result of the perceptual experience and is thus not unique to faces. Face-space begins with edge detectors and brightness detectors. As colour perception becomes more elaborate, colour detectors appear. Subsequently, faces become the most important visual feature, and so the face-space becomes narrowed to differentiate between faces, and so the edge detectors become replaced with feature detectors. This explanation actually mimics the neurological development, such that to begin with V1 is crucially well developed, whereas IT is not. Through visual development and experience, IT becomes more developed and more specialised. The plasticity of the brain indicates that the early visual systems are a precursor to later more complex visual discriminations. Indeed, the early visual systems are still vital in complex vision, such feature detectors in IT rely on input from spatial frequency analysers in V2. This is a much more parsimonious suggestion.

These two accounts actually indicate that there is a precursor to face-space. One is an innate first-order relational dimension that exists and then this drives how the visual system develops and the visual experience that is processed. The second suggests no such innate information and that visual experience drives the visual system and this acts as a precursor to face-space, which is still maintained in adulthood. The crucial difference between these two is what drives the visual system primarily: innate knowledge or visual experience. Either way, some form of face-space exists at birth, but both require experience to develop into the adult form of face-space. One requires added dimensions based on the visual input; the other requires the precise combination of earlier visual inputs in the most diagnostic features to be used for discrimination.

The above two explanations indicate that face-space appears when cortical processes begin to function within the brain. Prior to this, perceptual precursors exist which cause the face-space to form based on one dimension (CONSPEC) or any number

311

of sensory dimensions relating to contrast and spatial frequency information (Fourier analysis).

The sensory dimensions explanation has more explanatory power as there are a virtually unlimited number of spatial frequency combinations. As visual acuity increases, so to does the amount of information the newborn can attend to. As visual experience increases, the way these spatial frequencies are combined and form meaningful features improves. As such, the precursor to face-space and the creation of dimensions is likely to directly relate to the sensory information and the attention paid to faces. Very likely, the first dimension will relate to the eyes given the evidence from Maurer (1985) and a second dimension relating to hair (Blass and Camp, 2004).

6.3.2. Comparison to Object Recognition

The previous section suggests that the pre-requisites for face-space could be face specific or could be due to general perception. One possible method for exploring this would be to establish whether an "object-space" for encoding other objects exists in a similar manner to faces. Such a concept is less easy to envisage since there are few objects that are as homogenous as faces. Nevertheless, if face-space develops from perception and prominent visual features, then an object-space could also result in the same manner.

This line of research would broaden the argument on whether face recognition is special. If it could be established that face-space develops from a starting course that is the same as objects (i.e., perception), then it is plausible to consider why a similar object-space does not get created. A simple explanation is that faces occupy a primary position in the visual environment that most objects do not. As such, the neuronal populations become more responsive to particular facial features over features of objects.

312

An issue related to this is which dimensions appear in the developing face-space and which ones are lost. The research presented in Experiment 6 indicates that dormant dimensions in adults are those relating the lower facial features and the data presented by Eimer (2000) and Maurer (1985) is indicative of the eyes as forming an early dimension. However, such studies are only indicative of what the dimensions may represent. For example, the data in Experiment 6 may suggest that dormant dimensions are "the mouth size," "the width of the jaw," "the relative size of the lips compared to the size of the nose," or an almost infinite number of other possible combinations. Furthermore, if the dimensions are considered to be based primarily on spatial frequencies and Fourier analysis (as argued in Chapter 5), then the dimensions of flexible-face-space are likely to be virtually indescribable.

Not only are the precise definitions of the dimensions likely to remain unknown, the use of dimensions will be based on individual experience. As such, the dimensions of one person's face-space are likely to be different from another person's face-space. Moreover, the attentional weighting to each dimension may well be different for two separate people. These individual differences have not been extensively researched in face recognition, and could prove useful in exploring idiosyncrasies in face processing.

6.3.3. Computational and Neurological Modelling

Computational models are often applied to face processing abilities (c.f., the IAC, Burton et al., 1990; and face-space, Lewis, 2004). The flexible adult face-space proposed here may prove useful in such research. In particular, three areas of the flexible-face-space could be modelled: attentional order effects, attentional weighting, and semantic linking to dimensions. These shall be briefly described here.

In Matrix 6.1., it was suggested that not all dimensions of face-space are attended to at the same time. Instead, there is a default attentional order in the processing of the dimensions face-space. When the actual attentional order matches the default attentional order, the encoding and subsequent recognition of that face will be more accurate than if the actual attentional order does not match the default attentional order. This process could be computationally modelled given a set of assumptions relating to the number of dimensions and what these dimensions may actually be.

Given that different dimensions can be employed depending on the task at hand, as indicated in Matrix 6.2., this attentional weighting mechanism could also be computationally modelled. Subgroups of faces are recognised more accurately using different dimensions. In terms of an automated face recognition system, such a system may work better if a different set of features of a face is used for one subgroup of faces from another. Thus, such a system may require a first set of categorisation dimensions that determine the age, gender, or race of the face, as indicated in Section 2.8. Following this, the most relevant features can be used to examine the face.

The final aspect of the proposed model of flexible-face-space is the interaction between stored semantic knowledge and the visual input as reported in Experiment 8 and in Section 3.9. This could also be modelled providing more concrete links between the IAC and dimensions of face-space. Such a computational model would combine aspects of neural network architecture with the face-space and also neural models.

The flexible-face-space proposed here could also be modelled neurologically. The involvement of spatial frequency channels located in areas V1 and V2 in face recognition could be studied. Due to the observed effects that different task demands involve different dimensions of face-space (Experiments 6 and 7) and different spatial frequency channels (e.g., Schyns et al., 2002), it is likely that the projections from V1 and V2 to the

fusiform gyrus are different depending on the task. Thus far, only part of this issue has been explored due to limitations in scanning technology. Alternatively, these attentional effects may not be based in the visual cortex, but be an effect of neural acceleration in the attentional systems (Noguchi, Tanabe, Sadato, Hoshiyama, and Kakigi, 2007). Neuroimaging studies would be most appropriate for disentangling the low-level visual cortex effects and higher level attentional effects.

Another important aspect of the studies that needs to be investigated using imaging techniques is the cross-modal FIAE that has been shown in Experiments 16 to 20. This effect is further evidence for the distributed nature of the human face recognition system, whereby non-visual systems are involved in face processing. These cross-modal effects suggest the involvement of the auditory cortex and semantic cortices in face processing. Such involvement can be explored using imaging techniques to explore alternative routes through the face recognition system.

One final aspect of the results presented in this thesis that can be explored using neuroimaging is the possible involvement of attention in face recognition. Attentional guidance and the ability to shift eye gaze is controlled by systems in the prefrontal cortex, with inefficient searching directed by the inferior, middle, and superior prefrontal cortex, whereas efficient searching is only directed by the superior (Anderson, et al., 2007). Having established the importance of attention in face recognition, it may prove interesting to explore the correlates of prefrontal activity when processing inverted and upright faces. An inefficient search strategy is employed when scanning an inverted face (Barton, et al., 2006) and this may well be caused by inefficient attentional strategies. These could be explored in neuroimaging studies.

These three main lines of future research would considerably add to our understanding of face recognition and in particular the face-space model presented here. Some of these lines of further research directly stem from the developmental shrinking-face-space and the adult flexible-face-space. A more complete understanding of face recognition from birth until adulthood would be achieved from these lines of future work.

6.4. Conclusion

Three possible models of the immature face-space were presented. Evidence collected in this thesis favours a shrinking-face-space metaphor, whereby the number of dimensions in the immature face-space is greater than in adults. This greater number of dimensions allows children to process less experienced subtypes of faces and facial configurations relatively better than adults. Throughout development, the number of dimensions that children use to encode faces decreases until a default subset is attended to, with the exclusion of other dimensions. These dimensions are attended to in a particular order, and if they are attended to in the wrong order, encoding errors are likely to occur. The dormant dimensions can be re-activated and attended to given the right task demands. Thus, the adult face-space is more flexible than previously considered (by Valentine, 1991 or Lewis, 2004) due to the development of face-space being characterised by a reduction in the number of dimensions. The developmental shrinking-face-space thus becomes the adult flexible-face-space.

References

Abbonizio, G., Langley, K., & Clifford, C. W. G. (2002). Contrast adaptation may enhance contrast discrimination. *Spatial Vision, 16,* 45 – 58.

Acerra, F., Burnod, Y., & de Schonen, S. (2002). Modelling aspects of face processing in early infancy. *Developmental Science, 5,* 98 – 117.

Adams-Price, C. (1992). Eyewitness memory and aging: predictors of accuracy in recall and person recognition. *Psychology and Aging, 7,* 602 – 608.

Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). An area within human ventral cortex sensitive to "building" stimuli: evidence and implications. *Neuron, 21,* 373 – 383.

Akhtar, N. & Enns, J. T. (1989). Relations between covert orienting and filtering in the development of visual attention. *Journal of Experimental Child Psychology, 48,* 315 – 334.

Albrecht, D. G. Farrar, S. B. & Hamilton, D. B. (1984). Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *Journal of Physiology, 347,* 713 – 739.

Althoff, R. R. & Cohen, N. J. (1999). Eye-movement-based memory effect: a reprocessing effect in face perception. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 997 – 1010.

Ambady, N., Hallahan, M., & Conner, B. (1999). Accuracy of judgements of sexual orientation from thin slices of behaviour. *Journal of Personality and Social Psychology, 77,* 538 – 547.

Amedi, A., Malach, R., & Pascual-Leone, A. (2005). Negative BOLD differentiates visual imagery and perception. *Neuron, 5,* 859 – 872.

Anastasi, J. S. & Rhodes, M. G. (2005). An own-age bias in face recognition for children and older adults. *Psychonomic Bulletin and Review, 12,* 1043 – 1047.

Anastasi, J. S. & Rhodes, M. G. (2006). Evidence for an own-age bias in face recognition. *North American Journal of Psychology, 8,* 237 – 253.

Anastasi, J. S., Hodges, T. M. W., & Rhodes, M. G. (2006). Memory for faces: the interaction of race and age. Poster presented at the Psychonomic Society Conference, Houston, TX, 16th – 19th November.

Anderson, E. J., Mannan, S. K., Husain, M., Rees, G., Sumner, P., Mort, D. J., McRobbie, D., & Kennard, C. (2007). Involvement of prefrontal cortex in visual search. *Experimental Brain Research, 180,* 289 – 302.

Anderson, N. D. & Wilson, H. R. (2005). The nature of synthetic face adaptation. *Vision Research, 45,* 1815 – 1828.

Antes, J. R. & Penland, J. (1981). Picture context effects in eye movement patterns. In D. Fisher, R. Monty, & J. Senders (Eds.), *Eye Movements: Cognition and Visual Perception.* Hillsdale, NJ: Erlbaum.

Anthony, T., Copper, C., and Mullen, B. (1992). Cross-racial facial identification: A social cognitive integration. *Personality and Social Psychology Bulletin, 18,* 296 – 301.

Atkinson, J. & Braddick, O. (1980). Infant astigmatism: its disappearance with age. *Vision Research, 20,* 891 – 893.

Atkinson, J., Braddick, O., & Moar, K. (1977). Development of contrast sensitivity over the first 3 months of life in the human infant. *Vision Research, 17,* 1037 – 1044.

Artole, A., Brocher, S., & Singer, W. (1990). Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature, 347,* 69 – 72.

Bachmann, T. (1991). Identification of spatially quantised tachistoscopic images of faces: how many pixels does it take to carry identity? *European Journal of Cognitive Psychology, 3,* 85 – 103.

Badcock, J. C., Whitworth, F. A., Badcock, R. D., & Lovegrove, W. J. (1990). Low-frequency filtering and the processing of local – global stimuli. *Perception, 19,* 617 – 629.

Bahrick, H. P., Bahrick, P. O., & Wittlinger, R. O. (1975). *Journal of Experimental Psychology: General, 104,* 54 – 75.

Balakrishnan, J. D. (1998). Measures and interpretations of vigilance performance: evidence against the detection criterion. *Human Factors, 40,* 601 – 628.

Balakrishnan, J. D. (1999). Decision processes in discrimination: fundamental misrepresentations of signal detection theory. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1189 – 1206.

Banks, M. S. & Salapatek, P. (1981). Infant pattern vision: a new approach based on the contrast sensitivity function. *Journal of Experimental Child Psychology, 31,* 1 – 45.

Barkowitz, P. & Brigham, J. C. (1982). Recognition of faces: own-race bias, incentive, and time delay. *Journal of Applied Social Psychology, 12,* 255 – 268.

Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after effects. In C. Blakemore (ed.), *Vision: Coding and Efficiency.* Cambridge: Cambridge University Press.

Bartlett, J. C., Hurry, S., & Thorley, W. (1984). Typicality and familiarity of faces. *Memory and Cognition, 12,* 219 – 228.

Bartlett, J. C. & Searcy, J. (1993). Inversion and configuration of faces. *Cognitive Psychology, 25,* 281 – 316.

Barton, J. S., Deepak, S., & Malik, N. (2003). Attending to faces: change detection, familiarization, and inversion effects. *Perception, 32,* 15 – 28.

Barton J. J. S., Radcliffe, N., Cherkasova, M. V., Edelman, J., & Intriligator, J. M. (2006). Information processing during face recognition: the effects of familiarity, inversion, and morphing on scanning fixations. *Perception, 35,* 1089 – 1105.

Bednar, J. A. & Miikkulainen, R. (2003). Learning innate face preferences. *Neural Computation, 15,* 1525 – 1557.

Beale, J. M. & Keil, F. C. (1995). Categorical effects in the perception of faces. *Cognition, 57,* 217 – 239.

Benton, C. P., Jannings, S. J., & Chatting, D. J. (2006). Viewpoint dependence in adaptation to facial identity. *Vision Research, 46,* 3313 – 3325.

Berry, R. L., Teyler, T. J., & Taizhen, H. (1989). Induction of LTP in rat primary visual cortex: tetanus parameter. *Brain Research, 481,* 221 – 227.

Blakemore, C. (1973). The baffled brain. In R. L. Gregory & E. H. Gombrich (Eds.), *Illusion in Nature and Art.* London: Duckworth.

Blakemore, C., Nachmias, J., & Sutton, P. (1970). The perceived spatial frequency shift: evidence for frequency selective neurones in the human brain. *Journal of Physiology, 210,* 727 – 750.

Blaney, R. L. & Winograd, E. (1978). Developmental differences in children's recognition memory for faces. *Developmental Psychology, 14,* 441 – 442.

Blass, E. M. & Camp, C. A. (2004). The ontogeny of face identity 1. eight- to 21-week-old infants use internal and external face features in identity. *Cognition, 92,* 305 – 327.

Bloom, L. C. & Mudd, S. A. (1991). Depth of processing approach to face recognition: a test of two theories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 556 – 565.

Boeschoten, M. A., Kemner, C., Kenemans, J. L., & van Engeland, H. (2005). The relationship between local and global processing and the processing of high and low spatial frequencies studied by event-related potentials and source modeling. *Cognitive Brain Research, 24,* 228 – 236.

Bornstein, B. H. & Wilson, J. R. (2004). Extending the revelation effect to faces: haven't we met before? *Memory, 12,* 140 – 146.

Bothwell, R. K., Brigham, J. C., & Malpass, R. S. (1989). Cross-racial identification. *Personality and Social Psychology Bulletin, 15,* 19 – 25.

Boutet, I., Collin, C., & Faubert, J. (2003). Configural face encoding and spatial frequency information. *Perception and Psychophysics, 65,* 1078 – 1093.

Brand, A. (2005). Explaining the Verbal Overshadowing Effect. Unpublished PhD Thesis, Cardiff University.

Brigham, J. C. & Barkowitz, P. (1978). Do "They all look alike"? the effect of race, sex, experience, and attitudes on the ability to recognize faces. *Journal of Applied Social Psychology, 8,* 306 – 318.

Brigham, J. C. & Malpass, R. S. (1985). The role of experience and contact in the recognition of faces of own- and other-race persons. *Journal of Social Issues, 41,* 139 – 155.

Brigham, J. C., Maass, A., Snyder, L. D., & Spaulding, K. (1982). Accuracy of eyewitness identification in a field setting. *Journal of Personality and Social Psychology, 42,* 673 – 681.

Brigham, J. C. & Ready, D. J. (1985). Own-race bias in lineup construction. *Law and Human Behavior, 9,* 415 – 424.

Brigham, J. C. & Williamson, N. L. (1979). Cross-racial recognition and age: when you're over 60, do they still "all look alike?" *Personality and Social Psychology Bulletin, 5,* 218 – 222.

Brodeur, D. A. & Enns, J. T. (199). Lifespan differences in covert visual orienting. *Canadian Journal of Experimental Psychology, 51,* 20 – 35.

Brown, C. & Lloyd-Jones, T. J. (2002). Verbal overshadowing in a multiple face presentation paradigm: effects of description instruction. *Applied Cognitive Psychology, 16,* 873 – 885.

Brown, C. & Lloyd-Jones, T. J. (2003). Verbal overshadowing of multiple face and car recognition: Effects of within- versus across-category verbal descriptions. *Applied Cognitive Psychology, 17,* 183 – 201.

Brown, T. H., Chapman, P. F., Kairiss, E. W., and Keenan, C. L. (1988). Long-term synaptic potentiation. *Science, 242,* 724 – 728.

Bruce, C., Desimone, R., & Gross, C. G. (1981). Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. *Journal of Neurophysiology, 46,* 369 – 384.

Bruce, V. (1982). Changing faces: Visual and non-visual coding processes in face recognition. *British Journal of Psychology, 73,* 105 – 116.

Bruce, V. (1983). Recognising faces. *Philosophical Transactions of the Royal Society, London, B302,* 423 – 436.

Bruce, V. (1986). Influences of familiarity on the processing of faces. *Perception, 15,* 387 – 397.

Bruce, V. (1988). *Recognizing Faces.* Hove: Lawrence Erlbaum Associates.

Bruce, V. (1998). Fleeting images of shade: identifying people caught on video. *The Psychologist, 11,* 331 – 337.

Bruce, V., Burton, A. M., Hanna, E., Healey, P., Mason, O., Coombes, A., Fright, R., & Linney, A. (1993). Sex discrimination: how do we tell the difference between male and female faces? *Perception, 22,* 131 – 152.

Bruce, V., Campbell, R. N., Doherty-Sneddon, G., Import, A., Langton, S., McAuely, S., & Wright, R. (2000). Testing face processing skills in children. *British Journal of Developmental Psychology, 18,* 319 – 333.

Bruce, V., Doyle, T., Dench, N., & Burton, A. M. (1991). Remembering facial configurations. *Cognition, 38,* 109 – 144.

Bruce, V. & Humphreys, G. W. (1994). Recognizing objects and faces. *Visual Cognition, 1,* 141 – 180.

Bruce, V. & Valentine, T. (1985). Identity priming in the recognition of familiar faces. *British Journal of Psychology, 76,* 363 – 383.

Bruce, V. & Valentine, T. (1986). Semantic priming of familiar faces. *Quarterly Journal of Experimental Psychology, 35a,* 125 – 150.

Bruce, V., Valentine, T., & Baddeley, A. (1987). The basis of the ¾ advantage in face recognition. *Applied Cognitive Psychology, 1,* 109 – 120.

Bruce, V. & Young, A. W. (1986) Understanding face recognition. *British Journal of Psychology, 77,* 305 – 327.

Bruyer, R. & Velge, V. (1981). Lesions cérébrale unilaterale et trouble de la perception des visages: specificité du deficit? *Acta Neurological Belgica, 81,* 321 – 332.

Bull, R., & Green, J. (1980). The relationship between physical appearance and criminality. *Medicine, Science and the Law, 20,* 79 - 83.

Burton, A. M. (1994). Learning new faces in an interactive activation and competition model. *Visual Cognition, 1,* 313 - 348.

Burton, A. M. & Bruce, V. (1992). I recognize your face but I can't remember your name: a simple explanation? *British Journal of Psychology, 83,* 45 – 60.

Burton, A. M., & Bruce, V. (1993). Naming faces and naming names: Exploring an interactive activation model of person recognition. *Memory, 1,* 457 - 480.

Burton, A. M., Bruce, V., & Dench, N. (1993). What's the difference between men and women? Evidence from facial measurement. *Perception, 22,* 153 – 176.

Burton, A. M., Bruce, V., & Hancock, P. J. B. (1999). From pixels to people: a model of familiar face recognition. *Cognitive science, 23,* 1 – 31.

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

Burton, A. M. & Vokey, J. R. (1998). The face-space typicality paradox: understanding the face-space metaphor. *Quarterly Journal of Experimental Psychology, 51A,* 475 – 483.

Byatt, G. & Rhodes, G. (1998). Recognition of own-race and other-race caricatures: implications for models of face recognition. *Vision Research, 38,* 2455 – 2468.

Bywaters, M., Andrade, A., & Turpin, G. (2004). Determinants of the vividness of visual imagery: the effects of delayed recall, stimulus affect and individual differences. *Memory, 12,* 479 – 488.

Callaway, C. W., Lydic, R., Baghdoyan, H. A., & Hobson, J. A. (1987). Pontogeniculooccipital waves – spontaneous visual-system activity during rapid eye-movement sleep. *Cellular and Molecular Neurobiology, 7,* 105 – 149.

Carbon, C. C. & Leder, H. (2005). Face adaptation: changing stable representations of familiar faces within minutes? *Advances in Experimental Psychology, 1,* 1 – 7.

Carbon, C. C. & Leder, H. (2006). Last but not least. *Perception, 35,* 411 – 414.

Cardiff Council (2007). Numbers on roll 2007. Available at http://www.cardiff.gov.uk. Date accessed 16th July 2007.

Carey, S. (1978). A case study: face recognition. In E. Walker (Ed.), *Explorations in the biology of language.* Montgomery, VT: Bradford Books.

Carey, S. (1981). The development of face perception. In G. Davies, H. D. Ellis, J. Shepherd (Eds.), *Perceiving and Remembering Faces.* London: Academic Press.

Carey, S. & Diamond, R. (1977). From piecemeal to configurational representations of faces. *Science, 195,* 312 – 314.

Carey, S. & Diamond, R. (1994). Are faces perceived as configurations more by adults than by children? *Visual Cognition, 1,* 253 – 274.

Carey, S., Diamond, R., & Woods, B. (1980). The development of face recognition – a maturational component? *Developmental Psychology, 16,* 257 – 269.

Carson, D. R. & Burton, A. M. (2001). Semantic priming of person recognition: categorical priming may be a weaker form of the associative priming effect. *Quarterly Journal of Experimental Psychology, 54A,* 1155 – 1179.

Chance, J. E. & Goldstein, A. G. (1984). Face-recognition memory: implications for children's eyewitness testimony. *Journal of Social Issues, 40,* 69 – 85.

Chance, J. E., Goldstein, A. G., & Andersen, B. (1986). Recognition memory for infant faces: an analog of the other-race effect. *Bulletin of the Psychonomic Society, 24,* 257 – 260.

Chance, J. E., Goldstein, A. G., & McBride, L. (1975). Differential experience and recognition memory for faces. *Journal of Social Psychology, 97,* 243 – 253.

Chance, J. E., Turner, A. L., & Goldstein, A. G. (1982). Development of differential recognition for own- and other-race faces. *The Journal of Psychology, 112,* 29 – 37.

Charness, N. (1981). Search in chess: age and skill differences. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 467 – 476.

Chase, W. G. & Simon, H. A. (1973a). Perception in chess. *Cognitive Psychology, 4,* 55 – 81.

Chase, W. G. & Simon, H. A. (1973b). The mind's eye in chess. In W. G. Chase (Ed.), *Visual Information Processing.* London: Academic Press.

Cheour, M., Ceponiene, R., Lehtokoski, A., Luuk, A., Allik, J., Alho, K., & Naatanen, R. (1998). Development of language-specific phoneme representations in the infant brain. *Nature Neuroscience, 1,* 351 – 353.

Christman, S., Kitterle, F. L., & Niebauer, C. L. (1997). Hemispheric asymmetries in the identification of band-pass filtered letters. *Psychonomic Bulletin and Review, 4,* 277 – 284.

Chung, M. S. (1997). Face recognition: effects of age of subjects and age of stimulus faces. *Korean Journal of Developmental Psychology, 10,* 167 – 176.

Clifford, B. R. (1980). Voice identification by human listeners: on earwitness reliability. *Law and Human Behavior, 4,* 373 – 394.

Clifford, C. W. G. & Weston, E. (2005). Aftereffect of adaptation to Glass patterns. *Vision Research, 45,* 1355 – 1363.

Clifford, C. W. G., Wenderoth, P., & Spehar, B. (2000). A functional angle on adaptation in cortical vision. *Proceedings of the Royal Society of London, Series B, 267,* 1705 – 1710.

Cohen, K. (1981). The development of strategies of visual search. In D. Fisher, R. Monty, & J. Senders (Eds.), *Eye Movements: Cognition and Visual Perception.* Hillsdale, NJ: Erlbaum.

Coin, C. & Tiberghein, G. (1997). Encoding activity and face recognition. *Memory, 5,* 545 – 568.

Collin, C. A., Liu, C. H., Troje, N. F., McMullen, P. A., & Chaudhuri, A. (2004). Face recognition is affected by similarity in spatial frequency range to a greater degree than within-category object recognition. *Journal of Experimental Psychology: Human Perception and Performance, 30,* 975 – 987.

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

Cook, M. (1978). Eye movements during recognition of faces. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds), *Practical Aspects of Memory.* New York: Academic Press.

Costen, N. P., Parker, D. M., & Craw, I. (1994). Spatial content and spatial quantisation effects in face recognition. *Perception, 23,* 129 – 146.

Costen, N. P., Parker, D. M., & Craw, I. (1996). Effects of high –pass and low-pass spatial filtering on face identification. *Perception and Psychophysics, 58,* 602 – 612.

Courtois, M. R. & Mueller, J. H. (1979). Processing multiple physical features in facial recognition. *Bulletin of the Psychonomic Society, 14,* 74 – 76.

Craik, F. I. M. & Lockhart, R. S. (1972). Levels of processing: a framework for memory research. *Journal of Verbal Learning & Verbal Behavior, 11,* 671 – 684.

Craw, I. (1995). A manifold model of face and object recognition. In T. Valentine (ed.), *Cognitive and Computational Aspects of Face Recognition.* London: Routeledge.

Cross, J. F., Cross, J., & Daly, J. (1971). Sex, race, age, and beauty as factors in recognition of face. *Perception & Psychophysics, 10,* 393 – 396.

Cui, X., Jeter, C. B., Yang, D., Montague, P. R., & Eagleman, D. M. (2007). Vividness of mental imagery: individual variability can be measured objectively. *Vision Research, 47*, 474 – 478.

Davidoff, J. B. (1986). The specificity of face perception: evidence from psychological investigations. In R. Bruyer (Ed.), *The Neuropsychology of Face Perception and Facial Expression*. Hillsdale, NJ: Erlbaum.

Davies, G. M., Shepherd, J. W., & Ellis, H. D. (1978). Remembering faces – acknowledging our limitations. *Journal of the Forensic Science Society, 18*, 19 – 24.

Davies, G. M., Shepherd, J. W., & Ellis, H. D. (1979). Similarity effects in face recognition. *American Journal of Psychology, 92*, 507 – 523.

Davies, T. N. & Hoffman, D. (2002). Attention to faces: a change-blindness study. *Perception, 31*, 1123 – 1146.

Daw, P. S. & Parkin, A. J. (1981). Observations on the efficiency of two different processing strategies for remembering faces. *Canadian Journal of Psychology, 35*, 351 – 355.

Deffenbacher, K. A., Carr, T. H., & Leu, J. R. (1981). Memory for words, pictures, and faces: retroactive interference, forgetting, and reminiscience. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 299 – 305.

Deffenbacher, K. A., Leu, J. R., & Brown, E. L. (1981). Memory for faces: testing method, encoding strategy, and confidence. *American Journal of Psychology, 94*, 13 – 26.

DeGroot, A. D. (1965). *Thought and Choice in Chess*. The Hague: Mouton.

DeGroot, A. D. (1966). Perception and memory versus thought. In B. Kleinmuntz (Ed.), *Problem Solving*. New York: Wiley.

DeGroot, A. D. & Gobet, F. (1996). *Perception and Memory in Chess: Heuristics of the Processional Eye.* Assen: Van Gorcum.

de Haan, E. H. F. (2001). Face perception and recognition. In B. Rapp (Ed.), *The Handbook of Cognitive Neuropsychology: What deficits Reveal About the Human Mind.* New York: Psychology Press.

de Haan, M., Johnson, M. H., Maurer, D., & Perrett, D. I. (2001). Recognition of individual faces and average face prototypes by 1- and 3-month-old infants. *Cognitive Development, 16,* 659 – 678.

Delis, D. C., Robertson, L. C., Efron, R. (1986). Hemispheric specialization of memory for visual hierarchical stimuli. *Neuropsychologia, 24,* 205 – 214.

Deregowksi, J. B., Ellis, H. D., & Shepherd, J. W. (1975). Descriptions of White and Black faces by White and Black subjects. *International Journal of Psychology, 10,* 119 – 123.

de Valois, K. K. (1977). Spatial frequency adaptation can enhance contrast sensitivity. *Vision Research, 17,* 1057 – 1065.

de Valois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research, 22,* 531 – 544.

Devine, P. G. & Malpass, R. S. (1985). Orienting strategies in differential face recognition. *Personality and Social Psychology Bulletin, 11,* 33 – 40.

Diamond, R. & Carey, S. (1977). Developmental changes in the representation of faces. *Journal of Experimental Child Psychology, 23,* 1 – 22.

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

Dodson, C. S., Johnson, M. K., & Schooler, J. W. (1997). The verbal overshadowing effect: Why descriptions impair face recognition. *Memory and Cognition, 25,* 129 – 139.

Dunbar, R. (1995). Are you lonesome tonight? *New Scientist, 145,* 26 – 31.

Dziurawiec, S. & Ellis, H. D. (1986). Neonates attention to face-like stimuli – Goren, Sarty, and Wu revisited. *Bulletin of the British Psychological Society, 39,* A137.

Easterbrook, M. A., Kisilevsky, B. S., Hains, S. M. J., & Muir, D. W. (1999). Faceness or complexity: evidence from newborn visual tracking of facelike stimuli. *Infant Behavior and Development, 22,* 17 – 35.

Edmunds, A. J. & Lewis, M. B. (2007). The effect of rotation of configural encoding in a face-matching task. *Perception, 36,* 446 – 460.

Eger, E., Schweinberger, S. R., Dolan, R. J., & Henson, R. N. (2005). Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *NeuroImage, 26,* 1128 – 1139.

Eimer, M. (2000). Effects of face inversion on the structural encoding and recognition of faces – evidence from event-related brain potentials. *Cognitive Brain Research, 10,* 145 – 158.

Ellis, A. W. Young, B. M. Flude and D. C. Hay (1987). Repetition Priming of Face Recognition. *The Quarterly Journal of Experimental Psychology, 39A,* 193 – 210.

Ellis, H. D. (1981). Theoretical aspects of face recognition. In G.M. Davies, H. D. Ellis, & J. W. Shepherd (Eds.), *Perceiving and Remembering Faces.* London: Academic Press.

Ellis, H. D. (1986). Processes underlying face recognition. In R. Bruyer (Ed.), *The Neuropsychology of face perception and facial expression.* Hillsdale, NJ: Erlbaum.

Ellis, H. D. (1992a). The development of face processing skills. *Philosophical Transactions of the Royal Society B, 335,* 105 – 111.

Ellis, H. D. (1992b). The development of face processing skills. In V. Bruce, A. Cowey, A. W. Ellis, & D. I. Perrett (Eds.), *Processing the Facial Image.* Oxford: Clarendon Press.

Ellis, H. D., Deregowski, J. B., & Shepherd, J. W. (1975). Descriptions of White and Black faces by White and Black subjects. *International Journal of Psychology, 10,* 119 – 123.

Ellis, H. D., Ellis, D. M., & Hosie, J. A. (1993). Priming effects in children's' face recognition. *British Journal of Psychology, 84,* 101 – 110.

Ellis, H. D. & Flin, R. (1990). Encoding and storage effects in 7-year-olds' and 10-year-olds' memory for faces. *British Journal of Developmental Psychology, 8,* 77 – 92.

Ellis, H. D., Jones, D. M., & Mosdell, N. (1997). Intra- and inter-modal repetition priming of familiar faces and voices. *British Journal of Psychology, 88,* 143 – 156.

Ellis, H. D., Shepherd, J., & Bruce, A. (1973). The effects of age and sex upon adolescent's recognition of faces. *Journal of Genetic Psychology, 123,* 173 – 174.

Ellis, H. D., Shepherd, J. W., Gibling, F., & Shepherd, J. (1988). Stimulus factors in face learning. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical Aspects of Memory: Current Research and Issues.* Chichester: Wiley.

Emery, N. J. (2000). The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience and Biobehavioral Reviews, 24,* 581 – 604.

Endo, M. (1982). Cue saliency in upside-down faces. *Tohoku Psychologica Folia, 41,* 116 – 122.

Endo, M. (1983). Cue saliency in briefly presented faces: a comparison with cue saliency in upside-down faces. *Tohoku Psychologica Folia, 42,* 85 – 91.

Endo, M. (1986). Perception of upside-down faces: an analysis from the viewpoint of cue saliency. In H. D. Ellis, M. A. Jeeves, F. Newcombe, & A. W. Young. (Eds.), *Aspects of Face Processing*. Dordrecht: Martinus Nijhoff.

Enns, J. T. & Brodeur, D. A. (1989). A developmental study of covert orientation to peripheral visual cues. *Journal of Experimental Child Psychology, 48,* 171 – 189.

Eysenck, M. W. (1993). *Principals of Cognitive Psychology.* Hove: Erlbaum.

Fagan, J. F. (1972). Infants' recognition memory for faces. *Journal of Experimental Child Psychology, 14,* 453 – 476.

Fagan, J. F. (1979). The origins of facial pattern recognition. In M. H. Bornstein & W. Kessen (Eds.), *Psychological Development From Infancy: Image to Intention.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Fallshore, M. & Schooler, J. W. (1995). The verbal vulnerability of perceptual expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition, 21,* 1608 – 1623.

Fang, J., Blanz, V., & O'Toole, A. J. (2007). The role of familiarity in three-dimensional view-transferability of face identity adaptation. *Vision Research, 47,* 525 – 531.

Fantz, R. L. (1963). Pattern vision in newborn infants. *Science, 140,* 296 – 297.

Fantz, R. L. (1965). Visual perception from birth as shown by pattern selectivity. *Annals of the New York Academy of Sciences, 118,* 793 – 814.

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

Farroni, T., Johnson, M. H., Brockbank, M., & Simion, F. (2000). Infants' use of gaze direction to cue attention: the importance of perceived motion. *Visual Cognition, 7,* 705 – 718.

Feinman, S. & Entwisle, D. R. (1976). Children's ability to recognize other children's faces. *Child Development, 47,* 506 – 510.

Feller, M. B., Wellis, D. P., Stellwagen, D., Weblin, F. S., & Shatz, C. J. (1996). Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves. *Science, 272,* 1182 – 1187.

Findlay, J. (1981). Local and global influences on saccadic eye movements. In D. Fisher, R. Monty, & J. Senders (Eds.), *Eye Movements: Cognition and Visual Perception.* Hillsdale, NJ: Erlbaum.

Fiorentini, A., Maffei, L., & Sandini, G. (1983). The role of high spatial frequencies in face perception. *Perception, 12,* 195 – 201.

Fitzgerald, H. E. (1968). Autonomic papillary reflex activity during early infancy and its relation to social and non-social visual stimuli. *Journal of Experimental Child Psychology, 6,* 470 – 482.

Flin, R. H. (1980). Age effects in children's memory for unfamiliar faces. *Developmental Psychology, 16,* 373 – 374.

Flin, R. H. (1983). The development of face recognition. Unpublished PhD Thesis, University of Aberdeen.

Flin, R. H. (1985). Development of face recognition: an encoding switch? *British Journal of Psychology, 76,* 123 – 134.

Flin, R. H. & Dziurawiec, S. (1989). Developmental factors in face recognition. In A. Young & H. D. Ellis (Eds.), *Handbook of Research of Face Processing.* Amsterdam: Elsevier.

Fox, C. J. & Barton, J. J. S. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research, 1127,* 80 – 89.

Freire, A. & Lee, K. (2001). Face recognition in 4- to 7-year-olds: processing of configural, featural, and paraphernalia information. *Journal of Experimental Child Psychology, 80,* 347 – 371.

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

Frey, H. P., Konig, P., & Einhauser, W. (2007). The role of first- and second-order stimulus features for human overt attention. *Perception and Psychophysics, 69,* 153 – 161.

Frowd, C. D., Hancock, P. J. D., & Carson, D. (2004). EvoFIT: a holistic, evolutionary facial imaging technique for creating composites. *ACM Transactions on Applied Psychology, 1,* 1 – 21.

Fulton, A. & Bartlett, J. C. (1991). Young and old faces in young and old heads: the factor of age in face recognition. *Psychology and Aging, 6,* 623 – 630.

Gallant, J. L., Connor, C. E., Rakshit, S., Lewis, J. W., & van Essen, D. C. (1996). Neural responses to polar, hyperbolic, and Cartesian gratings in area V4 of the macaque monkey. *Journal of Neirophysiology, 76,* 2718 – 2739.

Galper, R. E. (1970). Recognition of faces in photographic negative. *Psychonomic Science, 19,* 207 – 208.

Galper, R. E. (1973). "Functional race membership" and recognition of faces. *Perceptual and Motor Skills, 37,* 455 – 462.

Ganis, G., Thompson, W. L., & Kosslyn, S. M. (2004). Brain areas underlying visual mental imagery and visual perception: an fMRI study. *Cognitive Brain Research, 2,* 226 – 241.

Gauthier, I., Anderson, A. W., Skudlarski, P., & Gore, J. C. (2000). Expertise for cars and birds recruits right hemisphere face areas. *Journal of Cognitive Neuroscience, 19,* 69.

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

Gauthier, I., Tarr, M. J., Moylan, J., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). The fusiform "face area" is part of a network that processes faces at the individual level. *Journal of Cognitive Neuroscience, 12,* 495 – 504.

Gauthier, I., Williams, P., Tarr, M. J., & Tanaka, J. (1998). Training "Greeble" experts: a framework for studying expert object recognition processes. *Vision Research, 38,* 2401 – 2428.

George, N., Dolan, R. J., Fink, G. R., Baylis, G. C., Russell, C., & Driver, J. (1999). Contrast polarity and face recognition in the human fusiform gyrus. *Nature Neuroscience, 2,* 574 – 580.

Georgeson, M. A. & Harris, M. G. (1984). Spatial selectivity of contrast adaptation: Models and data. *Vision Research, 24,* 729 – 741.

Gibson, J. J. & Walk, R. D. (1956). The effect of prolonged exposure to visually presented patterns on learning to discriminate them. *Journal of Comparative and Physiological Psychology, 49,* 239 – 242.

Ginsburg, A. P. (1978). *Visual information processing based on spatial filters constrained by biological data.* Unpublished doctoral dissertation, Cambridge University, Cambridge, UK.

Ginsburg, A. P. (1986). Spatial filtering and visual form perception. In. K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of Perception and Human Performance*. New York: Wiley.

Gobet, F. & Simon, H. A. (1996a). Recall of rapidly present random board positions is a function of skill. *Psychonomic Bulletin and Review, 3,* 159 – 163.

Gobet, F. & Simon, H. A. (1996b). Templates in chess memory: a mechanism for recalling several boards. *Cognitive Psychology, 31,* 1 – 40.

Goffaux, V., Gauthier, I., & Rossion, B. (2003). Spatial scale contribution to early visual differences between face and object processing. *Cognitive Brain Research, 16,* 416 – 424.

Goffaux, V., Hault, B., Michel, C., Vuonq, Q. C., & Rossion, B. (2005). The respective role of low and high spatial frequencies in supporting configural and featural processing of faces. *Perception, 34,* 77 – 86.

Goffaux, V. & Rossion, B. (2006). Faces are "spatial" – holistic face perception is supported by low spatial frequencies. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 1023 – 1039.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research, 39,* 3537 – 3560.

Goldstein, A. G. (1965). Learning of inverted and normally orientated faces in children and adults. *Psychonomic Science, 3,* 447 – 448.

Goldstein, A. G. (1975). Recognition of inverted photographs of faces in children and adults. *Journal of Genetic Psychology, 127,* 109 – 123.

Goldstein, A. G. & Chance, J. E. (1964). Recognition of children's faces. *Child Development, 35,* 129 – 136.

Goldstein, A. G. & Chance, J. E. (1976). Measuring psychological similarity of faces. *Bulletin of the Psychonomic Society, 7,* 407 – 408.

Goldstein, A. G. & Chance, J. E. (1980). Memory for faces and schema theory. *Journal of Psychology, 105,* 47 – 59.

Goldstone, R. L. (1998). Perceptual Learning. *Annual Review of Psychology, 49,* 585 – 612.

Goren, C. C., Sarty, M., & Wu, P. Y. K. (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Paediatrics, 56,* 544 – 549.

Graham, N. V. S. (1989). *Visual Pattern Analyzers.* New York: Oxford University Press.

Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzchak, Y., & Malach, R. (1999). Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron, 24,* 187 – 203.

Grossberg, S. (1978). A theory of visual coding, memory and development. In E. L. J. Leeuwenberg & H. R. J. M. Buffart (Eds.) *Formal theories of visual perception.* New York: Wiley.

Gustaffson, B. & Wigstrom, A. (1988). Physiological mechanisms underlying long-term potentiation. *Trends in Neurosciences, 11,* 156 – 162.

Haaf, R. A. (1974). Complexity and facial resemblance as determinants of response to facelike stimuli by 5- and 10-week-old infants. *Journal of Experimental Child Psychology, 18,* 480 – 487.

Haith, M. M., Bergman, T., & Moore, M. J. (1977). Eye contact and face scanning in early infancy. *Science, 198,* 853 – 855.

Halit, H., de Haan, M., Schyns, P. G., & Johnson, M. H. (2006). Is high-spatial frequency information used in the early stages of face detection? *Brain Research, 1117,* 154 – 161.

Hall, G. (1980). Exposure learning in animals. *Psychological Bulletin, 88,* 535 – 550.

Hall, G. (1991). *Perceptual and Associative Learning.* Oxford: Oxford University Press.

Han, S., Yund, E. W., & Woods, D. L. (2003). An ERP study of the global precedence effect: the role of spatial frequency. *Clinical Neurophysiology, 114,* 1850 – 1865.

Hanazawa, A. & Komatsu, H. (2001). Influence of the direction of elemental luminance gradients on the responses of V4 cells to textured surfaces. *Journal of Neuroscience, 21,* 4490 – 4497.

Hancock, P. J. B., Burton, A. M., & Bruce, V. (1996). Face processing: human perception and principal components analysis. *Memory and Cognition, 24,* 26 – 40.

Hanley, J. R., Smith, S. T., & Hadfield, J. (1998). I recognise you but I can't place you: an investigation of familiar-only experiences during tests of voice and face recognition. *Quarterly Journal of Experimental Psychology, 51A,* 179 – 195.

Hanley, J. R. & Turner, J. M. (2000). Why are familiar-only experiences more common for voices than faces? *Quarterly Journal of Experimental Psychology, 53A,* 1105 – 1116.

Harris, C. S. (1980). Insight or out of sight? two examples of perceptual plasticity in the human adults. In C. S. Harris (Ed.), *Visual Coding and Adaptability.* Hillsdale, NJ: Erlbaum.

Harvey, L. O., Jnr. (1986). Visual memory: what is remembered? In F. Klix & H. Hagendorf (Eds.), *Human Memory and Cognitive Capabilities.* The Hague: Elsevier.

Hasnain, N. & Husain, M. G. (1980). Is creative imagination equivalent to visual and auditory imaginations? *Asian Journal of Psychology and Education, 6,* 49 – 52.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Science, 4,* 223 - 233.

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

Hay, D. C. & Cox, R. (2000). Developmental changes in the recognition of faces and facial features. *Infant and Child Development, 9,* 199 – 212.

Hayes, A., Morrone, M. C., & Burr, D. C. (1986). Recognition of positive and negative band-pass filtered images. *Perception, 15,* 595 – 602.

Heinrich, T. S. & Bach, M. (2002). Contrast adaptation in retinal and cortical evoked potentials: No adaptation to low spatial frequencies. *Visual Neuroscience, 19,* 645 – 650.

Hershenson, M. (1964). Visual discrimination in the human newborn. *Journal of Comparative and Physiological Psychology, 58,* 270 – 276.

Hills, P. J., Lewis, M. B., & Honey, R. C. (2005). Stereotype priming in the recognition of unfamiliar faces: evidence for dynamic face-space. Paper presented at the BPS Cognitive Section Conference, Leeds, 13 – 15th June, 2005.

Hochberg, J. & Galper, R. E. (1967). Recognition of faces: I. an exploratory study. *Psychonomic Science, 9,* 619 – 620.

Honey, R. C. (1990). Stimulus generalization as a function of stimulus novelty and familiarity in rats. *Journal of Experimental Psychology: Animal Behavior Processes, 16,* 178 – 184.

Hood, B. M., Willen, J. D., & Driver, J. (1998). Adult's eyes trigger shifts of visual attention in human infants. *Psychological Science, 9,* 131 – 134.

Hosch, H. M. & Platz, S. J. (1984). Self-monitoring and eyewitness accuracy. *Personality and Social Psychology Bulletin, 10,* 289 – 292.

Hreczko, T. A. & Farkas, L. G. (1994). Norms of the craniofacial asymmetries in North American Whites. In L G. Farkas (Ed.), *Anthropometry of the Head and Face*. New York: Raven Press

Hubel, D. H. & Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology, 158,* 267 – 293.

Hughes, H. C., Nozawa, G., & Ketterle, F. (1990). Global precedence, spatial frequency channels and the statistics of natural images. *Journal of Cognitive Neuroscience, 8,* 197 – 230.

Hulbert, A. (2001). Trading faces. *Nature Neuroscience, 4,* 3 – 5.

Ishai, A., Schmidt, C. F., & Boesiger, P. (2005). Face perception is mediated by a distributed cortical network. *Brain Research Bulletin, 67,* 87 – 93.

Issac, A. R. & Marks, D. F. (1994). Individual differences in mental imagery experience: developmental changes and specialization. *British Journal of Psychology, 85,* 479 – 500.

Itier, R. J. & Taylor, M. J. (2004). Effects of repetition and configural changes on the development of face recognition processes. *Developmental Science, 7,* 469 – 487.

Ito, M., Tamura, H., Fujita, I., & Tanaka, K. (1995). Size and position invariance of neuronal responses in monkey inferotemporal cortex. *Journal of Neurophysiology, 73,* 218 – 226.

Jiang, F., Blanz, V., & O'Toole, A. (2006). Probing the visual representation of faces with adaptation: a view from the other side of the mean. *Psychological Science, 17,* 493 – 500.

Jiang, F., Blanz, V., & O'Toole, A. (2007). The role of familiarity in three-dimensional view-transferability of face identity adaptation. *Vision Research, 47,* 525 – 531.

Johnson, M. H., Dziurawiec, S., Ellis, H. D., & Morton, J. (1991). Newborns' preferential

tracking of face-like stimuli and its subsequent decline. *Cognition, 40,* 1 – 19.

Johnson, M. H. & Vecera, S. P. (1996). Cortical differentiation and neurocognitive

development: the parcellation conjecture. *Behavioural Processes, 36,* 195 – 212.

Johnston, R. A. & Barry, C. (2006). Repetition priming of access to biographical

information from faces. *Quarterly Journal of Experimental Psychology, 59,* 326 –

339.

Johnston, R. A. & Ellis, H. D. (1995). Age effects in the processing of typical and

distinctive faces. *The Quarterly Journal of Experimental Psychology, 45a,* 447 –

465.

Joseph, J. E., Gathers, A. D., Liu, X., Corbly, C. R., Whitaker, S. K., & Bhatt, R. S.

(2006). Neural developmental changes in processing inverted faces. *Cognitive,*

Affective, and Behavioral Neuroscience, 6, 223 – 235.

Kanwisher, N. McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module

in human extrastriate cortex specialized for face perception. *Journal of*

Neuroscience, 17, 4302 – 4311.

Kanwisher, N., Stanley, D., & Harris, A. (1999). The fusiform face area is selective for

faces not animals. *Neuroreport, 10,* 183 – 187.

Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the

human fusiform face area. *Cognition, 68,* B1 – B11.

Kersten, D., Troje, N. F., & Bulthoff, H. H. (1996).. Phenomenal competition for poses of

the human head. *Perception, 25,* 367 – 368.

Kessen, W., Salapatek, P., & Haith, M. (1972). The visual response of the human

newborn to linear contour. *Journal of Experimental Child Psychology, 13,* 9 – 20.

Khazipov, R., Sirota, A., Leinekugel, X., Holmes, G. L., Ben-Arf, Y., & Buzsaki, G. (2004). Early motor activity drives spindle bursts in the developing somatosensory cortex. *Nature, 432,* 758 – 761.

Kinchla, R. A. & Wolfe, J. (1979). The order of visual processing: "top down," "bottom up," or "middle out." *Perception and Psychophysics, 25,* 225 – 231.

Kitterle, F. L. and Selig, L. M. (1991). Visual field effects in the discrimination of sine-wave gratings. *Perception and Psychophysics, 50,* 15 – 18.

Klatzky, R. L., Martin, G. L., & Kane, R. A. (1982a). Semantic interpretation effects on memory for faces. *Memory & Cognition, 10,* 195 - 206.

Klatzky, R. L., Martin, G. L., & Kane, R. A. (1982b). Influence of social-category activation on processing of visual information. *Social Cognition, 1,* 95 - 109.

Kleiner, K. A. (1987). Amplitude and phase spectra as indices of infants' pattern preferences. *Infant Behavior and Development, 10,* 49 – 59.

Kleiner, K. A. (1993). Specific vs non-specific face recognition device. In B. de Boysson-Bardies, S. de Schonen, P. W. Jusczyk, P. McNeilage, & J. Morton. (1993). *Developmental Neurocognition: Speech and Face Processing in the First Year of Life.* New York: Kluwer Academic.

Kleiner, K. A. & Banks, M. S. (1987). Stimulus energy does not account for 2-month-olds' face preferences. *Journal of Experimental Psychology: Human Perception and Performance, 13,* 594 – 600.

Komatsu, Y., Fujii, K., Maeda, J., Sakaguchi, H., & Toyama, K. (1988). Long-term potentiation of synaptic transmission in kitten visual cortex. *Journal of Neurophysiology, 59,* 124 – 141.

Konorski, J. A. & Petersik, J. T. (2003). Effects on face recognition of spatial-frequency information contained in inspection and test stimuli. *Journal of General Psychology, 130,* 229 – 244.

Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Malikovic, V., Weise, S. B., Chabris, C. F., et al. (1993). Visual mental imagery activates topographically organized visual cortex: PET investigations. *Journal of Cognitive Neuroscience, 5,* 263 – 287.

Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Chabris, C. F., Rauch, S. K., & Anderson, A. K. (1994). Identifying objects seen from different viewpoints: a PET investigation. *Brain, 117,* 1055 – 1071.

Kreiman, G., Koch, D., & Fried, I. (2000). Imagery neurons in the human brain. *Nature, 408,* 357 – 361.

Frey, H. P., Konig, P., & Einhauser, W. (2007). The role of first- and second-order stimulus features for human overt attention. *Perception and Psychophysics, 69,* 153 – 161.

Kuhl, P. K., Williams, K. A., Lacerda, F., Stevens, K. N., et al. (1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science, 255,* 606 – 608.

Kung, C. C. (2007). Expanding the FFA debate: a re-evaluation of recent studies and evidence for the multimodal and dynamic fusiform gyrus. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 67,* 4733.

La Gasse, L. L. (1993). Effects of good form and spatial frequency on global precedence. *Perception and Psychophysics, 53,* 89 – 105.

Lamb, M. R., Robertson, L. C., & Knight, R. T. (1989). Attention and interference in the processing of hierarchical patterns: inferences from patients with right and left temporal-parietal lesions. *Neuropsychologia, 27,* 471 – 483.

Lamb, M. R., Robertson, L. C., & Knight, R. T. (1990). Component mechanisms underlying the processing of hierarchically organized patterns: inferences from patients with unilateral cortical lesions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 471 – 483.

Lamb, M. R. & Yung, E. W. (1993). The role of spatial frequency information in the processing of hierarchically organized stimuli. *Perception and Psychophysics, 54,* 773 – 784.

Lamb, M. R. & Yung, E. W. (1996a). Spatial frequency and attention: effect of level-, target-, and location-repetition on the processing of global and local forms. *Perception and Psychophysics, 58,* 363 – 373.

Lamb, M. R. & Yung, E. W. (1996b). Spatial frequency and interference between global and local level of structure. *Visual Cognition, 3,* 193 – 219.

Lavrakas, P. J., Buri, J. R., and Mayzner, M. S. (1976). A perspective on the recognition of other-race faces. *Perception and Psychophysics, 20,* 475 – 481.

Lawson, R. (2006). Local and global processing biases fail to influence face, object and word recognition. Poster presented at the Experimental Psychology Society Meeting, 10-12[th] July 2006, Plymouth, UK (Abstract).

Leder, H. & Bruce, V. (2000). When inverted faces are recognized: the role of configural information in face recognition. *Quarterly Journal of Experimental Psychology, 53,* 513 – 536.

Leder, H. & Carbon, C. (2005). When context hinders! Learn-test compatibility in face recognition. *The Quarterly Journal of Experimental Psychology, 55a,* 235 – 250.

Leinekugel, X., Khazipov, R., Cannon, R., Hirase, H., Ben-Arf, Y., & Buzsaki, G. (2002). Correlated bursts of activity in the neonatal hippocampus in vivo. *Science, 296,* 2049 – 2052.

Leek, M. R. (2001). Adaptive procedures in psychophysical research. *Perception & Psychophysics, 63,* 1279 – 1292.

Leippe, M. R. (1995). The case for expert testimony about eyewitness memory. *Psychology, Public Policy, and Law, 1,* 909 – 959.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level after effects. *Nature Neuroscience, 4,* 89 – 94.

Leopold, D. A., Rhodes, G., Müller, K., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of the Royal Society B, 272,* 897 – 904.

Levin, D.T. (2000). Race as a visual feature: using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General, 129,* 559 – 574.

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America, 49,* 467 – 477.

Lewis, M. B. (2004). Face-space-r: towards a unified account of face recognition. *Visual Cognition, 11,* 29 – 70.

Lewis, M. B. & Glenister, T. E. (2003). A sideways look at configural encoding: two different effects of face rotation. *Perception, 32,* 7 – 14.

Lewis, M. B. & Johnston, R. A. (1998) Understanding caricatures of faces. *Quarterly Journal of Experimental Psychology, 51A,* 321 – 346.

Lewis, M. B. & Johnston, R. A. (1999). A unified account of the effect of caricaturing faces. *Visual Cognition, 6,* 1 – 42.

Lewkowicz, D. J. & Ghazanfar, A. A. (2006). The decline of cross-species intersensory perception in human infants. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 6771 – 6774.

Light, L. L., Hollander, S., & Kayra-Stuart, F. (1981). Why attractive people are harder to remember. *Personality and Social Psychology Bulletin, 7,* 269 – 276.

Light, L. L., Kayra-Stuart, F., & Hollander, S. (1979). Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 212 - 228.

Lindsay, R. C. L., Smith, S. M., & Pryke, S. (1999). Measures of lineup fairness: do they postdict identification accuracy? *Applied Cognitive Psychology, 13,* S93 – S107.

Lippe, W. R. (1994). Rhythmic spontaneous activity in the developing avian auditory system. *Journal of Neuroscience, 14,* 1486 – 1495

List, J. (1986). Age and schematic differences in the reliability of eye-witness testimony. *Developmental Psychology, 22,* 50 – 57.

Loftus, E. F. & Loftus, G. (1980). On the permanence of stored information in the human brain. *American Psychologist, 35,* 409 – 420.

Liu, C. H., Collin, C. A., Rainville, S. J., & Chaudhuri, A. (2000). The effects of spatial frequency overlap on face recognition. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 956 – 979.

Luce, T. S. (1974a). The role of experience in inter-racial recognition. *Personality and Social Psychology Bulletin, 1,* 39 – 41.

Luce, T. S. (1974b). Blacks, whites, and yellows: they all look alike to me. *Psychology Today, 8,* 105 – 108.

Lunn, R. & Banks, W. P. (1986). Visual fatigue and spatial frequency adaptation to visual displays of text. *Human Factors, 28,* 457 – 464.

Lyle, K. B. & Johnson, M. K. (2004). Effects of verbalization on lineup face recognition in an interpolated inspection paradigm. *Applied Cognitive Psychology, 18,* 393 – 403.

Macchi-Cassia, V., Kuefner, D., Westerlund, A., & Nelson, C. A. (2006). A behavioural and ERP investigation of 3-month-olds' face preferences. *Neuropsychologia, 44,* 2113 – 2125.

Macmillan, N. A. and Creelman, C. D. (2005). *Detection Theory: A User's Guide.* Cambridge: Cambridge University Press.

Macrae, C. N. & Lewis, H. L. (2002). Do I know you?: Processing orientation and face recognition. *Psychological Science, 13,* 194 – 196.

Maddess, T., McCourt, M. E., Blakeslee, B., & Cunningham, R. B. (1988). Factors governing the adaptation of cells in area-17 of the cat visual cortex. *Biological Cybernetics, 59,* 229 – 236.

Magnussen, S. & Greenlee, M. W. (1985). Marathon adaptation to spatial contrast saturation in sight. *Vision Research, 28,* 1409 – 1411.

Malpass, R. S. (1974). Racial bias in eyewitness identification. *Personality and Social Psychology Bulletin, 1,* 42 – 44.

Malpass, R. S. (1981). Effective size and defendant bias in eyewitness identification lineups. *Law and Human Behavior, 5,* 299 – 309.

Malpass, R. S. & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of Personality and Social Psychology, 13,* 330 – 335.

Malpass, R. S., Lavigueur, H., & Weldon, D. E. (1973). Verbal and visual training in face recognition. *Perception and Psychophysics, 14,* 285 – 292.

Marks, D. F. (1973). Visual imagery differences in the recall of pictures. *British Journal of Psychology, 64,* 17 – 24.

Marks, D. F. & Isaac, A. R. (1995). Topographical distribution of EEG activity accompanying visual and motor imagery in vivid and non-vivid imagers. *British Journal of Psychology, 86,* 271 – 282.

Marois, R., Yi, D. J., & Chun, M. M. (2004). The neural fate of consciously perceived and missed events in the attentional blink. *Neuron, 41,* 465 – 472.

Mather, G. (1980). The movement aftereffect and a distributed-shift model for coding the direction of visual movement. *Perception, 9,* 379 – 392.

Maurer, D. (1983). The scanning of compound figures by young infants. *Journal of Experimental Child Psychology, 35,* 437 – 448.

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

Maurer, D., Lewis, T. L., Brent, H. P., & Levin, A. V. (1999). Rapid improvement in the acuity of infants after visual input. *Science, 286,* 108 – 110.

Maurer, D. & Salapatek, P. (1976). Developmental changes in the scanning of faces by young infants. *Child Development, 47,* 523 – 527.

Mayer, D. L. & Dobson, V. (1982). Visual acuity development in infants and young children, as assessed by operant preferential looking. *Vision Research, 22,* 1141 – 1151.

Mazard, A., Schiltz, C., & Rossion, B. (2006). Recovery from adaptation to facial identity is larger for upright than inverted faces in the human occipito-temporal cortex. *Neuropsychologia, 44,* 912 – 922.

McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience, 9,* 605 – 610.

McClelland, J. L. (1981). Retrieving general and specific information from stored knowledge of specifics. *Proceedings of the Third Annual Meeting of the Cognitive Science Society,* 170 – 172.

McConahay, J. B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In S. L. Gaertner, and J. F. Dovidio (Eds.), *Prejudice, discrimination, and racism* (pp. 91-126). San Diego, CA: Academic Press.

McKelvie, S. J. (1994). The Vividness of Visual Imagery Questionnaire as a predictor of facial recognition memory performance. *British Journal of Psychology, ,* 93 – 104.

McKone, E., Aitkin, A., & Edwards, M. (2005). Categorical and coordinate relations in faces, or Fechnar's law and face-space instead? *Journal of Experimental Psychology: Human Perception and Performance, 31,* 1181 – 1198.

McLaren, I. P. L. (1997). Categorisation and perceptual learning: an analogue of the face inversion effect. *The Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 50,* 257 – 273.

McLaren, I. P. L. and Mackintosh, N. J. (2000). An elemental model of associative learning: I. latent inhibition and perceptual learning. *Animal Learning and Behavior, 28,* 211 – 246.

McLaren, I. P. L. and Mackintosh, N. J. (2002). Associative learning and elemental representation: II. generalization and discrimination. *Animal Learning and Behavior, 30,* 177 – 200.

Megreya, A. M. & Burton, A. M. (2006). Unfamiliar faces are not faces: evidence from a matching task. *Memory and Cognition, 34,* 865 – 876.

Meissner, C. A. & Brigham, J. C. (2001a). A meta-analysis of the verbal overshadowing effect in face identification. *Applied Cognitive Psychology, 15,* 603 – 616.

Meissner, C. A. & Brigham, J. C. (2001b). Thirty years of investigating the own-race bias in memory for faces - A meta-analytic review. *Psychology, Public Policy, and Law, 7,* 3 – 35.

Memon, A. & Bartlett, J. (2002). The effects of verbalization on face recognition in young and older adults. *Applied Cognitive Psychology, 16,* 635 – 650.

Menees, S. M. (1998). The effect of spatial frequency adaptaion on the latency of spatial contrast detection. *Vision Research, 38,* 3933 – 3942.

Miller, G. A. (1956). The magic number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review, 63,* 81 – 93.

Minear, M. & Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research, Methods, Instruments, and Computers, 34,* 630 – 633.

Moradi, D., Koch, C., & Shimojo, S. (2005). Face adaptation depends on seeing the face. *Neuron, 45,* 169 – 175.

Moran, J. & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science, 229,* 782 – 784.

Morison, V. & Slater, A. (1985). Contrast and spatial frequency components in visual preferences of newborns. *Perception, 14,* 345 – 348.

Morton, J. & Johnson, M. H. (1991). CONSPEC and CONLEARN: a two-process theory of infant face recognition. *Psychological Review, 98,* 164 – 181.

Moutoussis, K. & Zeki, S. (2002). The relationship between cortical activation and perception investigated with invisible stimuli. *Proceedings of the National Academy of Sciences, USA, 99,* 9527 – 9532.

Movshon, J. A. & Lennie, P. (1979). Pattern selective adaptation in visual cortical neurones. *Nature, 278,* 850 – 852.

Mueller, J. H., Carlomusto, M., & Goldstein, A. G. (1978). Orienting task and study time in facial recognition. *Bulletin of the Psychonomic Society, 11,* 313 – 316.

Mulligan, N. W. & Lozito, J. P. (2006). An asymmetry between memory encoding and retrieval: revelation, generation, and transfer-appropriate processing. *Psychological Science, 17,* 7 – 11.

Mundy, M. E., Dwyer, D. M., & Honey, R. C. (2006). Inhibitory associations contribute to perceptual learning in humans. *Journal of Experimental Psychology: Animal Behavior Processes, 32,* 178 – 184.

Mundy, M. E., Honey, R. C., & Dwyer, D. M. (2007). Simultaneous presentation of similar stimuli produces perceptual learning in human picture processing. *Journal of Experimental Psychology: Animal Behavior Processes, 33,* 124 – 138.

Nagayama, R., Yoshida, H., & Toshima, T. (1995). Interrelationship between the facial expression and familiarity: analysis using spatial filtering and inverted presentation. *Japanese Journal of Psychology, 66,* 327 – 335.

Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research, 39,* 3824 – 3833.

Navon, D. (1977). Forest before the trees: The precedence of global features in visual perception. *Cognitive Psychology, 9,* 353 – 383.

Nelson, C. A. (2001). The development and neural bases in face recognition. *Infant and Child Development, 10,* 3 – 18.

Nelson, K. (1981). Social cognition in a script framework. In J. H. Flavell & L. Ross (Eds.), *Social Cognitive Development.* Cambridge: Cambridge University Press.

Nelson, K. (1996). *Language in Cognitive Development.* New York: Cambridge University Press.

Noguchi, Y., Tanabe, H. C., Sadato, N., Hoshiyama, M., & Kakigi, R. (2007). Voluntary attention changes the speed of perceptual neural processing. *European Journal of Neuroscience, 25,* 3163 – 3172.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorisation relationship. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 14,* 700 – 708.

O'Craven, K. M. & Kanwisher, N. (2000). Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience, 12,* 1013 – 1023.

O'Donovan, M. J. (1999). The origin of spontaneous activity in developing networks of the vertebrate nervous system. *Current Opinion in Neurobiology, 9,* 94 – 104.

Ogawa, H., Takeda, Y., & Kumada, T. (2007). Probing attentional modulation of contextual cueing. *Visual Cognition, 15,* 276 – 289.

Over, R. (1971). Comparison of normalization theory and neural enhancement explanation of negative after effects. *Psychological Bulletin, 75,* 225 – 243.

Parker, D. M. & Costen, N. P. (1999). One extreme or the other or perhaps the golden mean? Issues of spatial resolution in face processing. *Current Psychology: Developmental, Learning, Personality, 18,* 118 – 127.

Pascalis, O., Coleman, M., & Campbell, R. Primates first? How expertise affects the ability to identify individuals within and across species. *Journal of Cognitive Neuroscience, 61C.*

Pascalis, O., de Haan, M., & Nelson, C. A. (2002). Is face processing species-specific during the first year of life? *Science, 296,* 1321 – 1323.

Pasupathy, A. & Connor, C. E. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology, 82,* 2490 – 2502.

Pasupathy, A. & Connor, C. E. (2001). Shape representation in area V4: position-specific tuning for boundary conformation. *Journal of Neurophysiology, 86,* 2505 – 2519.

Pavlov, I. P. (1927). *Conditioned Reflexes.* London: Oxford University Press.

Pearson, D. A. & Lane, D. M. (1991). Visual attention movements: a developmental study. *Child Development, 61,* 1779 – 1795.

Pedelty, L., Levine, S. C., & Shevell, S. K. (1985). Developmental changes in face processing: results from multidimensional scaling. *Journal of Experimental Child Psychology, 39,* 421 – 436.

Pellicano, E. & Rhodes, G. (2003). Holistic processing of faces in preschool children and adults. *Psychological Science, 14,* 618 – 622.

Pellicano, E., Rhodes, G., & Peters, M. (2006). Are preschoolers sensitive to configural information in faces? *Developmental Science, 9,* 270 – 277.

Penfield, W. (1969). Consciousness, memory and man's conditioned reflexes/ In K. Pribram (Ed.), *On the Biology of Learning.* New York: Harcourt, Brace Jovanovich.

Perfect, T. J. (2003). Local processing bias impairs line-up performance. *Psychological Report, 93,* 393 – 394.

Perrett, D. I., Mistlin, A. J., & Chitty, A. J. (1987). Visual neurons responsive to faces. *Trends in Neurosciences, 10,* 358 – 364.

Perrett, D. I., Rolls, E. T., & Caan, W. (1979). Temporal lobe cells of the monkey with visual responses selective for faces. *Neuroscience Letters, 31,* S358.

Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Experimental Brain Research, 47,* 329 – 342.

Peyrin, C., Chokron, S., Guyader, N., Gout, O., Moret, J., & Marendaz, C. (2006). Neural correlates of spatial frequency processing: a neuropsychological approach. *Brain Research, 1073-1074,* 1 – 10.

Pezdek, K., Blandon-Gitlin, I., & Moore, C. (2003). Children's face recognition memory: more evidence for the cross-face effect. *Journal of Applied Psychology, 88,* 760 – 763.

Phillips, R. J. (1972). Why are faces hard to recognize in photographic negative. *Perception & Psychophysics, 12,* 425 – 426.

Piaget, J. (1952). *The Origins of Intelligence in Children.* New York: International Universities Press.

Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., & Vullieumier, P. (2005). Portraits or people? Distinct representations of face identity in the human visual cortex. *Journal of Cognitive Neuroscience, 17,* 1043 – 1057.

Regan, D. & Hamstra, S. J. (1992). Shape discrimination and the judgement of perfect symmetry: dissociation of shape from size. *Vision Research, 32,* 1845 – 1864.

Rhodes, G. (1988). Looking at faces: first-order and second-order features as determinants of facial appearance. *Perception, 17,* 43 – 63.

Rhodes, G. (1993). Configural coding, expertise, and the right-hemisphere advantage for face recognition. *Brain and Cognition. 22.* 19 – 41.

Rhodes, G., Brake, S., & Atkinson, A. P. (1993). What's lost in inverted faces? *Cognition, 47,* 25 – 57.

Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: implications for mental representations of faces. *Cognitive Psychology, 19,* 473 – 497.

Rhodes, G., Hickford, C., & Jeffery, L. (2000). Sex-typicality and attractiveness: are supermale and superfemale faces super-attractive. *British Journal of Psychology, 91,* 125 – 140.

Rhodes, G. & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision Research, 46,* 2977 – 2987.

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

Rhodes, G., Jeffery, L., Watson, T. L., Jacquet, E., Winkler, C., & Clifford, C. W. G. (2004). Orientation-contingent face after effects and implications for face-coding mechanisms. *Current Biology, 14,* 2119 – 2123.

Ridley, M. (1993). *The Red Queen.* London: Penguin.

Richardson, A. (2000). Individual differences in visual imagination imagery. In R. G. Kunzendorf & B. Wallace (Eds.), *Individual Difference in Conscious Experience.* Amsterdam: John Benjamins Publishing Company.

Richardson, A. & McAndrew, F. (1990). The effects of photic stimulation and private self-consciousness on the complexity of visual imagination imagery. *British Journal of Psychology, 81,* 381 – 394.

Robbins, R., McKone, E., & Edwards, M. (2007). After effects for face attributes with different natural variability: adapter position effects and neural models. *Journal of Experimental Psychology: Human Perception and Performance, 33,* 570 – 592.

Robertson, L. C. (1996). Attentional persistence for features of hierachical patterns. *Journal of Experimental Psychology: General, 125,* 227 – 249.

Robertson, L. C., Lamb, M. R., & Knight, R. T. (1988). Effects of lesions of temporal-parietal junction on perceptual and attentional processing in humans. *Journal of Neuroscience, 8,* 3757 – 3769.

Robertson, L. C., Lamb, M. R., & Zaidel, R. (1993). Callosal transfer and hemisphere laterality in response to hierarchical patterns: evidence from normal and commisurotomized subjects. *Neurophsychology, 7,* 325 – 342.

Roediger, H. L. (1990). Implicit memory: Retention without remembering. *American Psychologist, 45,* 1043 – 1056.

Rogers, B. & Graham, M. (1985). Motion parallax and the perception of three-dimensional surfaces. In D. J. Ingle, M. Jeannerod, & D. N. Lee (Eds.), *Brain Mechanisms and Spatial Vision.* Dordrecht, Netherlands: Martinus Nijhoff.

Rolls, E. T. & Baylis, G. C. (1986). Size and contrast have only small effects on the responses to faces of neurons in the cortex of the superior temporal sulcus of the monkey. *Experimental Brain Research, 65,* 38 – 48.

Rondan, C. & Deruelle, C. (2004). Face processing in high functioning autistic adults: a look into spatial frequencies and the inversion effect. *Journal of Cognitive and Behavioral Psychotherapies, 4,* 149 – 163.

Rossion, B. & Gauthier, I. (2002). How does the brain process upright and inverted faces? *Behavioral and Cognitive Neuroscience Reviews, 1,* 63 – 75.

Rotshtein, P., Henson, R. N. A., Treves, A., Driver, J., & Dolan, R. J. (2004). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nature Neuroscience, 8,* 107 – 113.

Ryan, R. S., Aulenbach, R., Becker, M., Johnson, T., Christman, J., Murray, K., & Bergmaier, J. (2006). Global processing improves self-assessment of face recognition ability. Poster presented at the 18[th] Annual Convention for the Association for Psychological Science, New York, May 25-28[th], 2006.

Ryu, J. & Chaudhuri, A. (2006). Representations of familiar and unfamiliar faces as revealed by viewpoint-after effects. *Vision Research, 46,* 4059 – 4063.

Saariluoma, P. (1994). Location coding in chess. *Quarterly Journal of Experimental Psychology, 47A,* 607 – 630.

Saarinen, J., & Levi, D. M. (1995). Perceptual learning in vernier acuity: What is learned? *Vision Research, 35,* 519 – 527.

Salapatek, P. (1968). Visual scanning of geometric figures by the human newborn. *Journal of Comparative and Physiological Psychology, 66,* 247 – 258.

Salapatek, P. & Kessen, W. (1966). Visual scanning of triangles by the human newborn. *Journal of Experimental Child Psychology, 3,* 155 – 167.

Saltz, E. & Sigel, I. E. (1967). Concept overdiscrimination in children. *Journal of Experimental Psychology, 73,* 1 – 8.

Scapinello, K. F. & Yarmey, A. D. (1970). The role of familiarity and orientation in immediate and delayed recognition of pictorial stimuli. *Psychonomic Science, 21,* 329 – 330.

Schooler, J. W. (2002). Verbalization produces a transfer inappropriate processing shift. *Applied Cognitive Psychology, 16,* 989 – 997.

Schooler, J. W. & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology, 22,* 36 – 71.

Schooler, J. W., Fiore, S. M., & Brandimonte, M. A. (1997). At a loss for words: Verbal overshadowing of perceptual memories. In D. L. Medin (Ed.) *The psychology of learning and motivation* (pp. 292-334). San Diego, CA: Academic Press.

Schwarzer, G. (2000). Development of face processing: the effect of face inversion. *Child Development, 71,* 391 – 401.

Schyns, P. G., Bonnar, L., & Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. *Psychological Science, 13,* 402 – 409.

Schyns, P. G. & Gosselin, F. (2003). Diagnostic use of scale information for componential and holistic recognition. In M. A. Peterson & G, Rhodes (Eds.),

Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes. New York: Oxford University Press.

Schyns, P. G. & Oliva, A. (1999). Dr. Angry and Mr. Smile: when categorization flexibility modifies the perception of faces in rapid visual presentations. *Cognition, 69*, 243 – 265.

Scott, L. S., Shannon, R. W., & Nelson, C. A. (2005). Behavioral and electrophysiological evidence of species-specific face processing. *Cognitive, Affective, and Behavioral Neuroscience, 5*, 405 – 416.

Scott, L. S., Shannon, R. W., & Nelson, C. A. (2006). Neural correlates of human and monkey face processing in 9-month-old infants. *Infancy, 10*, 171 – 186.

Searcy, J. H. & Bartlett, J. C. (1996). Inversion and processing of component and spatial-relational information in faces. *Journal of Experimental Psychology: Human Perception and Performance, 22*, 904 – 915.

Searcy, J. H., Bartlett, J. C., & Memon, A. (2000). Influence of post-event narratives, line-up conditions and individual differences on false identification by young and older eyewitnesses. *Legal and Criminological Psychology, 5*, 219 – 235.

Searcy, J. H., Bartlett, J. C., Memon, A., & Swanson, K. (2001). Aging and lineup performance at long retention intervals: effects of metamemory and context reinstatement. *Journal of Applied Psychology, 86*, 207 – 214.

Seeleman, V. (1940). The influence of attitude upon the remembering of pictorial material. *Archives of Psychology, 258*, 69.

Shatz, C. J. (1990). Impulse activity and the patterning of connections during CNS development. *Neuron, 5*, 745 – 756.

Shatz, C. J. (1993). Spontaneous action-potential activity and the patterning of connections during CNS development. *Journal of Cellular Biochemistry, 17D*, 184.

Shatz, C. J. (1996a). Emergence of order in visual system development. *Journal of Physiology, 90,* 141 – 150.

Shatz, C. J. (1996b). Emergence of order in visual system development. *Proceedings of the National Academy of Sciences of the United States of America, 93,* 602 – 608.

Sheffield, F. D. (1965). Relation between classical conditioning and instrumental learning. In W. F. Prokasy (Ed.), *Classical Conditioning: A Symposium.* New York: Appleton-Centur-Crofts.

Shepherd, J. W. & Deregowski, J. B. (1981). Races and faces – a comparison of the responses of Aficans and Whites to faces of the same and different races. *British Journal of Social Psychology, 20,* 125 – 133.

Shepherd, J. W., Deregowski, J. B. & Ellis, H. D. (1974). A cross-cultural study of recognition memory for faces. *International Journal of Psychology, 9,* 205 – 212.

Shepherd, J. W., Ellis, H. D., & Graham, G. M. (1977). *Perceiving and Remembering Faces.* Chichester: Wiley.

Shepherd, J. W., Ellis, H. D., McMurran, M., & Davies, G. M. (1978). Effect of character attribution on Photofit construction of a face. *European Journal of Social Psychology, 8,* 263 – 268.

Shepherd, J. W., Gibling, F. & Ellis, H. D. (1991). The effects of distinctiveness, presentation time and delay on face recognition. *European Journal of Cognitive Psychology, 3,* 137 – 145.

Shulman, G. L., Sullivan, M. A., Gish, K., & Sakoda, W. J. (1986). The role of spatial-frequency channels in the perception of local and global structure. *Perception, 15,* 259 – 273.

Shulman, G. L. & Wilson, J. (1987). Spatial frequency and selective attention to spatial location. *Perception, 16,* 103 – 111.

Sierra-Vazquez, V., Serrano-Pedraza, I., & Luna, D. (2006). The effect of spatial-frequency filtering on the visual processing of global structure. *Perception, 35,* 1583 – 1609.

Simon, H. A. & Barenfeld, M. (1969). Information processing analysis of perceptual processes in problem solving. *Psychological Review, 76,* 473 – 483.

Simon, H. A. & Gilmartin, K. (1973). A simulation of memory for chess positions. *Cognitive Psychology, 5,* 29 – 46.

Sinha, P. & Poggio, T. (1996). I think I know that face… *Nature, 384,* 404.

Slone, A. E., Brigham, J. C., & Meissner, C. A. (2000). Social and cognitive factors affecting the own-race bias in Whites. *Basic and Applied Social Psychology, 22,* 71 – 84.

Snowdon, R. J. & Hammett, S. T. (1996). Spatial frequency adaptation: Threshold elevation and perceived contrast. *Vision Research, 36,* 1797 – 1809.

Spitzer, H., Desimone, R., & Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science, 240,* 338 – 340.

Sporer, S. L. (2001). Recognizing faces of other ethnic groups – an integration of theories. *Psychology, Public Policy, and Law, 7,* 36 – 97.

Stacey, P. C., Walker, S. & Underwood, J. D. M. (2005). Face processing and familiarity: evidence from eye-movement data. *British Journal of Psychology, 96,* 407 – 422.

Stark, L. & Ellis, S. (1981). Scanpaths revisited: cognitive models direct active looking. In D. Fisher, R. Monty, & J. Senders (Eds.), *Eye Movements: Cognition and Visual Perception.* Hillsdale, NJ: Erlbaum.

Stechler, G. (1964). Newborn attention as affected by medication during labor. *Science, 144,* 315 – 317.

Stillman, J. & Kemp, T. (1993). Visual versus auditory imagination: image qualities, perceptual qualities, and memory. *Journal of Mental Imagery, 17,* 181 – 194.

Suzuki, S. (2001). Attention-dependent brief adaptation to contour orientation: a high-level aftereffect for convexity? *Vision Research, 41,* 3883 – 3902.

Suzuki, S. (2003). Attentional selection of overlapped shapes: a study using brief shape after effects. *Vision Research, 43,* 549 – 561.

Suzuki, S. (2005.) High-level pattern coding revealed by brief shape after effects. In C. W. G. Clifford and G. Rhodes (Eds.) *Fitting the Mind to the World: Adaptation and After-effects in High-Level Vision.* Oxford: Oxford University Press.

Swets, J. A. (1966). *Signal detection and recognition by human observers.* New York: John Wiley & Sons Inc.

Tanaka, J. W. & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology, 46A,* 225 – 246.

Tanaka, J. W., Kay, J. B., Grinnell, E., Stansfield, B., & Szechter, L. (1998). Face recognition in young children: when the whole is greater than the sum of its parts. *Visual Cognition, 5,* 479 – 496.

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

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience, 19,* 109 – 139.

Teller, D. Y., McDonald, M. A., Preston, K., & Sebris, S. L., et al. (1986). Assessment of visual acuity in infants and children: the acuity card procedure. *Developmental Medicine and Child Neurology, 28,* 779 – 789.

Thomas, H. (1965). Visual-fixation responses of infants to stimuli of varying complexity. *Child Development, 36,* 629 – 638.

Thompson, P. (1980). Margaret Thatcher: A new illusion. *Perception, 9,* 483 – 484.

Thurstone, L. L. (1927). A law of comparative judgement. *Psychological Review, 34,* 273 – 386.

Tieger, T. & Ganz, L. (1979). Recognition of faces in the presence of two-dimensional sinusoidal masks. *Perception and Psychophysics, 26,* 163 – 167.

Troje, N. F. & Bulthoff, H. H. (1996). Face recognition under varying poses: the role of texture and shape. *Vision Research, 36,* 1761 – 1771.

Turk, M. & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience, 3,* 71 – 86.

University of Stirling Psychology Department. (2000). Psychological Image Collection at Stirling (PICS). http://www.pics.stir.ac.uk. (Date accessed 20/07/2002).

Valentin, D., Abdi, H., Edelman, B., & O'Toole, A. J. (1997). Principal component and neural network analysis of face images: what can be generalized in gender classification? *Journal of Mathematical Psychology, 41,* 398 – 413.

Valentine, T. (1986). Encoding processes in face recognition. Unpublished PhD thesis, University of Nottingham.

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. (1991). A unified account of the effects of distinctiveness, inversion and race on face recognition. *Quarterly Journal of Experimental Psychology, 43A,* 161 – 204.

Valentine, T. & Bruce, V. (1986a). The effects of distinctiveness in recognising and classifying faces. *Perception, 15,* 525 – 535.

Valentine, T. & Bruce, V. (1986b). The effect of race, inversion and encoding activity upon face recognition. *Acta Psychologica, 61,* 259 – 273.

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

Valentine, T., Chiroro, P., & Dixon, R. (1995). An account of the own-race bias and the contact hypothesis based on a 'face space' model of face recognition. In T. Valentine (ed.), *Cognitive and Computational Aspects of Face Recognition.* London: Routeledge.

Valentine, T. & Endo, M. (1992). Towards an exemplar model of face processing: the effects of race and distinctiveness. *Quarterly Journal of Experimental Psychology, 44A,* 671 – 703.

Verde, M. F. & Rotello, C. M. (2003). Does familiarity change in the revelation effect? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 739 – 746.

Verde, M. F. & Rotello, C. M. (2004). ROC curves show that the revelation effect is not a single phenomenon. *Psychonomic Bulletin and Review, 11,* 560 – 566.

Vokey, J. R. & Read, J. D. (1992). Familiarity, memorability, and the effect of typicality on the recognition of faces. *Memory and Cognition, 20,* 291 – 302.

Vurpillot, E. (1968). The development of scanning strategies and their relation to visual differentiation. *Journal of Experimental Child Psychology, 6,* 632 – 650.

Wagner, D. A. & Scurrah, M. J. (1971). Some characteristics of human problem solving in chess. *Cognitive Psychology, 2,* 451 – 478.

Walker, P. M. and Hewstone, M. (2006). A perceptual discrimination investigation of the own-race effect and intergroup experience. *Applied Cognitive Psychology, 20,* 461 – 475.

Wang, G., Tanaka, K., & Tanifugi, M. (1996). Optical imaging of functional organization in the monkey inferotemporal cortex. *Science, 272,* 1665 – 1668.

Ward, L. M. (1982). Determinants of attention to local and global features in of visual form. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 562 – 581.

Watson, T. L. & Clifford, C. W. G. (2003). Pulling faces: an investigation of the face-distortion aftereffect. *Perception, 32,* 1109 – 1116.

Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature, 428,* 557 – 561.

Webster, M. A. & MacLin, O. H. (1999). Figural after effects in the perception of faces. *Psychonomic Bulletin & Review, 6,* 647 – 653.

Webster, M. A. & Mollon, J. D. (1995). Colour constancy influenced by contrast adaptation. *Nature, 373,* 694 – 698.

Webster, M. A. & Mollon, J. D. (1999). Selective adaptation to correlated variations in luminance and chromaticity. *Perception, 18,* 352.

Webster, M. A. & Miyahara, E. (1997). Contrast adaptation and the spatial structure of natural images. *Journal of the Optic Society of America A, 14,* 2355 – 2366.

Weisstein, N. A. (1980). Tutorial: the joy of Fourier analysis. In C. S. Harris (Ed.), *Visual Coding and Adaptability.* Hillsdale, NJ: Erlbaum.

Weston, N. J. & Perfect, T. J. (2005). Effects of processing bias on the recognition of composite face halves. *Psychonomic Bulletin and Review, 12,* 1038 – 1042.

Wiggs, C. L. & Martin, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion in Neurobiology, 8,* 227 – 233.

Wilson, H. R., Ferrara, V. P., & Yo, C. (1992). A psychophysically motivated model for two-dimensional motion perception. *Visual Neuroscience, 9,* 79 – 97.

Wilson, H. R. & Humanski, R. (1993). Spatial frequency adaptation and contrast gain control. *Vision Research, 33,* 1133 – 1149.

Winograd, E. (1981). Recognition memory for faces following nine different judgements. *Bulletin of the Psychonomic Society, 8,* 419 – 421.

Wong, R. O.L. (1999a). Retinal waves: stirring up a storm. *Neuron, 24,* 493 – 495.

Wong, R. O. L. (1999b). Retinal waves and visual system development. *Annual Review of Neuroscience, 22,* 29 – 47.

Wong, R. O. L., Meister, M., & Shatz, C. J. (1993). Transient period of correlated bursting activity during development of the mammalian retina. *Neuron, 11,* 923 – 938.

Wright, D. B. & Stroud, J. N. (2002). Age differences in lineup identification accuracy: people are better with their own age. *Law and Human Behavior, 26,* 641 – 654.

Yamashita, J. A., Hardy, J. L., De Valois, K. K., & Webster, M. A. (2005). Stimulus selectivity of figural after effects for faces. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 420 – 437.

Yarmey, A. D. (1971). Recognition memory for familiar 'public' faces: effects of orientation and delay. *Psychonomic Science, 24,* 286 – 288.

Yarmey, D. A. (1995). Earwitness speaker identification. *Psychology, Public Policy, and Law, 1,* 792 – 816.

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

Yin, R. K. (1970). *Face Recognition: A Special Process?* Unpublished PhD Thesis, Massachusetts Institute of Technology.

Young, A. W. & Bion, P. J. (1981). Identification and storage of line drawings presented to the left and right cerebral hemispheres of adults and children. *Cortex, 17,* 459 – 463.

Young, A. W. & Bion, P. J. (1982). The nature of the sex difference in right hemisphere superiority for face recognition. *Cortex, 19,* 215 – 225.

Young, A. W., Flude, B. M., Hellawell, D. J., Ellis, A. W. (1994). The nature of semantic priming effects in the recognition of familiar people. *British Journal of Psychology, 85,* 393 – 411.

Young, A. W., Hay, D. C., & Ellis, A. W. (1985). The faces that launched a thousand slips: everyday difficulties and errors in recognizing people. *British Journal of Psychology, 76,* 495 – 523.

Young, A. W., Hellawell, D., & de Haan, E. H. F. (1988). Cross-domain priming in normal subjects and a prosopagnoisc patient. *Quarterly Journal of Experimental Psychology, 40A,* 561 – 580.

Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configural information in face perception. *Perception, 16,* 747 – 759.

Young, M. P. & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science, 256,* 1327 – 1331.

Yovel, G. & Kanwisher, N. (2004). Face perception: domain specific, not process specific. *Neuron, 44,* 889 – 898.

Zeki, S. (1973). Colour coding in rhesus monkey prestriate cortex. *Brain Research, 53,* 422 – 427.

Zeki, S. (1977). Colour coding in the superior temporal sulcus of rhesus monkey cortex. *Proceedings of the Royal Society of London, B, 197,* 195 – 223.

Zeki, S. (1993). *A Vision of the Brain.* Oxtord: Blackwell.

Zhao, L. & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research, 41,* 2979-2994.