UNIVERSITY OF SOUTHAMPTON

Familiarity: How Does Knowing a Face Affect Processing?

Elizabeth Lee BSc (Hons)

Thesis submitted for the degree of Doctor of Philosophy

DEPARTMENT OF PSYCHOLOGY

FACULTY OF SOCIAL SCIENCES

Submitted September 2002

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL SCIENCES
DEPARTMENT OF PSYCHOLOGY

Doctor of Philosophy

FAMILIARITY: HOW DOES KNOWING A FACE AFFECT PROCESSING?

by Elizabeth Lee

The ability to distinguish a familiar face from an unfamiliar one has obvious evolutionary advantages and remains crucially important in terms of social functioning. Given the importance of recognising faces, the focus of the present thesis is how perceiver knowledge with faces, may influence the processes being used. Specifically, this thesis addresses the issue of whether all faces are processed in the same way, or whether differences exist in the processing of familiar over unfamiliar faces

Examination of the empirical findings in this area led to the construction of three hypotheses. The first predicted that familiar face processing would be more affected by stimulus manipulation than unfamiliar face processing, because the use of expert strategies would be interfered with. The second hypothesis predicted that familiar face processing would, in contrast, be less affected by stimulus manipulation than unfamiliar face processing, because familiarity permits formation of robust representations that withstand interference. The final hypothesis predicted no difference in the way familiar and unfamiliar faces are processed across stimulus manipulation. Two sets of experiments were presented comprising face recognition tasks (studies 1, 2, 3, & 4) and face classification tasks (studies 5, 6, & 7). Although the first experiment, a speeded familiarity task, produced findings consistent with the expertise hypothesis, the more robust old/new experiments using stimulus inversion (studies 2 & 3) and negation (study 4) showed that in all three cases, the robust representation hypothesis was supported. The face classification studies (studies 5, 6, & 7) produced results wholly compatible with the robust representation hypothesis.

It is proposed that the notion of robust representations provides an appealing and parsimonious way of accounting for these results. The robustness of the representation is assumed to correlate with familiarity and so determine the speed, or automaticity, with which it is accessed. This, combined with the provision of feedback mechanisms from later processing stages back to earlier ones provides an explanation for how familiarity can influence processing at fundamental stages of the face recognition system.

LIST OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGEMENTS

Completing this thesis would have been impossible without the help and support of a number of people. First and foremost, is my supervisor, Dr. Sarah Stevenage. Thank you, Sarah, for your expertise, guidance and friendship, and for allowing me the opportunity to begin and complete this work. Thanks, also, to Dr. Nick Donnelly for the many invaluable discussions and for making me think, and to Prof. Stevan Harnad for sewing the seeds of interest in Cognitive Psychology.

It is difficult to describe accurately the level of support provided by the post-graduate students within the Psychology Department, suffice to say that the last few years would have been much less rewarding without the friendship and support of Wendy Kneller, Wendy Smith, Jane Prichard, Julie Winstone, and Vered Rafaely. I'll always be grateful for the ideas, coffee, tennis and fun we had together, as well as for all your help and advice, especially over the last few months. Special thanks to Wendy Kneller, who helped in ways too numerous to mention, and to Wendy Smith for all the stats (and other) advice. Thanks, too, to Judy Chatwin, for listening to my moaning and being tactful enough not to keep asking when this thesis would be finished.

Finally, to Jon, Tom and Fran, the people who have been, and always will be, the most important things in my life. Thank you for putting up with the disruption to domestic life, and for keeping everything in perspective.

CHAPTER 1

How does Familiarity Fit with the Processing of the Facial Image?

Background

Within a second of seeing a face, we are able to determine gender, age, and mood, as well as identity, thus allowing appropriate social interaction (Carey, 1992). With so many judgements being made on the basis of facial appearance, face perception has become recognised as being fundamental to our existence as psychosocial animals. This is borne out by the demonstration that even neonates with a median age of just 9 minutes, show a preference to attend to a face-like configuration over a scrambled or blank one (Goren, Sarty & Wu, 1975). Consequently, face perception has been at the centre of decades of research at both a psychological and clinical level, and the literature in this area is now vast. The focus, however, of the present thesis is the examination of how perceiver knowledge or expertise with faces may influence the processes being used. Specifically, this thesis addresses the issue of whether all faces are processed in the same way, or whether qualitative differences exist in the processing of familiar over unfamiliar faces.

The purpose of the present chapter is to outline the background literature in the pertinent and related areas. This chapter will cover the basic skills that perceivers display when processing faces, the theoretical models that have been advanced to account for these skills, and the possible influence that perceiver knowledge, in the form of expertise and stimulus familiarity, may have. Evidence will be presented to suggest that stimulus familiarity may actually play an important role in the low-level and high-level processing of faces. Two theoretical positions are advanced to account for this. At a process-level (that is, the way in which faces are treated by the visual system), the notion of expertise and its reliance on configural processing is explored, whereas at the representational level (that is, the way in which faces are characterised by the visual system), the notion of robust representations is examined. The predictions that follow from these two positions form the basis of empirical enquiry within this thesis.

Accounts of Face Processing

Bruce and Young's (1986) Information Processing Model

In simple terms, the task of face recognition is one of identifying the object in front of you. A large literature exists at the level of object recognition (e.g., Biederman, 1987; Marr & Nishihara, 1978), and these models are influential in describing how we distinguish between objects at a basic category level, i.e., distinguishing between a dog and a cat. However, such models fall down when what is required is a description of how we distinguish between different exemplars within the same basic category, i.e., one dog from another dog, or indeed, one face from another face. This subordinate level discrimination is exactly what is called for when accounting for face perception. To fulfil this goal, a specialised model is required, and the most thoroughly developed model of face processing is presented by Bruce and Young (1986; see Figure 1.1). Bruce and Young used an Information Processing approach based on that used by Hay and Young (1982) in order to describe the likely stages involved in face processing. Specifically, their model offers an account of familiar face recognition, and this is a point that will be returned to later.

Bruce and Young's model presents a description of how we process identity, expression, speech-related cues, and other cues directly gleaned from the face. The model is essentially a hierarchically organised, sequential stage model in which the output of one stage serves as the input for the next stage of processing. The first stage (*structural encoding*) involves the formation of a viewer-centred description of the face, in which the surface layout of the face is outlined relative to the viewer. The viewer-centred description assists the analysis of expression (for instance, is the person happy, sad, or angry?), and facial speech analysis (lip-reading, conversational turn-taking, etc.), but more critically, the viewer-centred description leads to the extraction of a description that is independent of expression. Both the viewer-centred description and the expression-independent description are required to inform *directed visual processing*. It is this mechanism that allows a perceiver to process unfamiliar faces according to particular features, or to make person-related inferences from the visual presentation of a face (i.e., Mr X must be a chef because he is wearing a chef's hat).

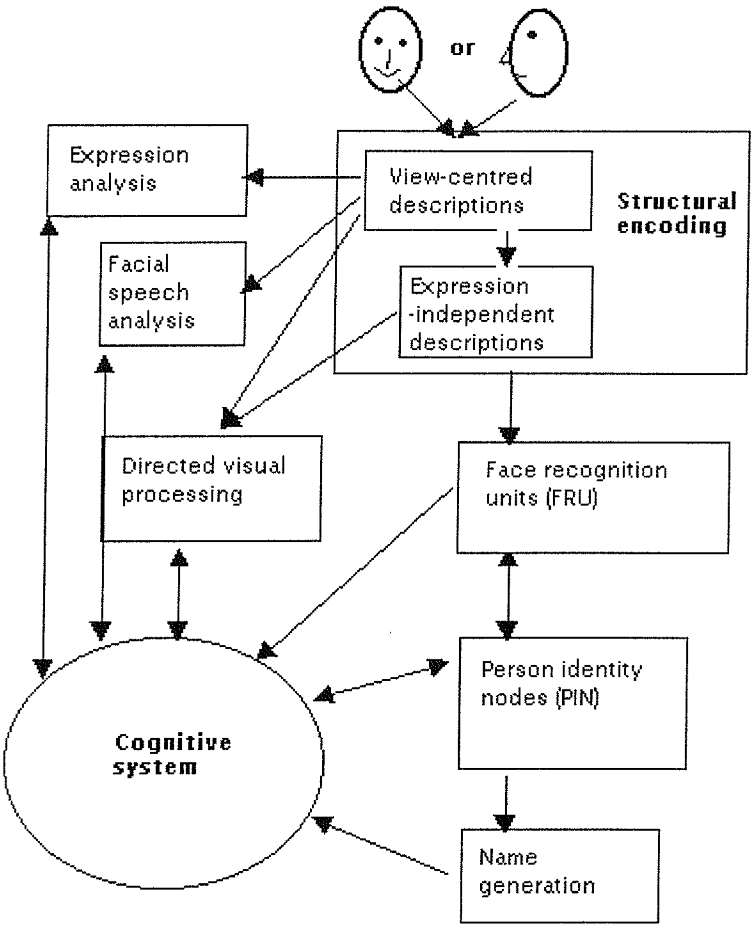

Figure 1.1: Bruce and Young's (1986) Information Processing model of familiar face recognition (adapted from source).

Disruption to processing at this structural encoding stage would be devastating and is exemplified by the clinical condition 'metamorphosia' in which a patient is unable to reconcile the perceptual input as being a complete or intact face (Hécaen & Angelergues, 1962). Furthermore, disruption at this stage would impact on all subsequent processes as the descriptions of the face act as the input for all subsequent stages. However, as is clear from Figure 1.1, the processes that arise after structural encoding appear to diverge, and both behavioural studies (Bruce, 1986; Campbell, 1989; Young, McWeeny, Hay & Ellis, 1986) and neurophysiological studies (Hasselmo, Rolls & Baylis, 1989) suggest that the processing pathways should be considered as cognitively and functionally distinct and independent of one another.

The expression-independent description also forms the input to the next stage of the recognition route – the *face recognition units* (FRUs). The purpose of the face recognition units is to determine whether the person has been seen before. As such, this is the first stage dedicated to processing identity. If a person is familiar, then a face recognition unit will have been laid down to represent that person. The expression-independent description of the face will activate the unit such that a feeling of familiarity is experienced. This activation is passed to the next stage of processing – the *person identity nodes* (PINs) which release semantic information (for example, occupation, marital status, nationality, etc.), and successful activation at this level is finally passed on to the ultimate stage of processing – *name generation.* Only if activation has successfully occurred at each of the previous layers of the model will the viewer have the possibility of activating an appropriate name.

The final component of Bruce and Young's model is the *'cognitive system'*. This is described as a 'catch-all' component, similar in function to the central executive of Baddeley and Hitch's (1974) Working Memory Model. Decision-making, selection of information, selection of processing strategy, and control of responding are all directed by the cognitive system, as is the storage of semantic memory.

Using their model, Bruce and Young (1986) were able to account for many of the misidentifications and errors of recognition in both normal and brain-damaged individuals. The types of face processing errors highlighted by Young, Hay and Ellis's (1985) diary study suggested a range of possible ways in which processing may be impaired. For example, participants may fail to recognise familiar individuals, or may mis-recognise them as someone from an associated context. Similarly, participants may inadvertently 'recognise' the face of someone they have never seen before. All of these errors would arise from insufficient or inappropriate activation at the FRU level of the model such that the presented face becomes identified as the wrong person. The diary study also revealed errors associated with misremembering, or failing to remember, the semantic information known about a familiar person. This sort of error is located at the PIN level of the model such that either inappropriate information is activated, or insufficient activation is passed on to the PINs causing failure to activate any semantic information. Finally, the diary study revealed frequent naming errors, which usually took the form of a complete naming block, or a more temporary tip-of-the-tongue (TOT) state (Yarmey, 1973). Also evident, however, were naming errors involving the activation of an incorrect but associated or phonologically similar name. These errors are located at the name generation level of the model and are explained, again, by either inappropriate, or insufficient activation reaching this level of the model.

The frequency of naming errors among the corpus of errors was highlighted by Young *et al.* (1985) and has been used to support the argument that naming is the most fragile of the processing stages because it is dependent on the successful completion of all the prior stages. It is notable, also, that no diarists reported an inability to retrieve semantic information having recognised a face as familiar and named it appropriately. This also was taken as support of the sequential stage model of Bruce and Young. Clearly it would be impossible in this model to have successful activation at the 'name generation' level if there had not been successful activation to access semantic information.

The Bruce and Young model has offered researchers an elegant and parsimonious framework for guiding further investigation, and a great deal of support, has been provided for the basic sequence of stages. For example, timing studies support the sequence by demonstrating that perceivers are faster at making a familiarity decision than they are at making a semantic level decision, and that this is faster still than a naming decision (Bruce & Young, 1986). At a clinical level also, the model has provided a framework for accounting for deficits of face processing. The condition of metamorphosia has already been mentioned and can be explained with reference to a lesion or attenuation of a neural pathway at the level of structural encoding. Similarly, lesion or attenuation at the level of the FRUs would result in a deficit when identifying known individuals, and this describes the much-researched condition of prosopagnosia (see Blanc-Garin, 1986, or De Renzi, 1986 for a definition). Finally, lesion or attenuation of the neural pathway at the level between PINs and name generation would result in an inability to name known individuals even though semantic information could be retrieved. This describes the condition of proper name anomia (Flude, Ellis, & Kay, 1989).

Modelling the Bruce and Young model using IAC

The Bruce and Young model was, however, limited by the fact that it was of a classical box-and-arrow format and, as such, failed to describe the processes that may be occurring within each of the boxes. More recent models have attempted to operationalise the processes, and this approach is exemplified by the connectionist work of Burton, Bruce and Johnston, (1990). Burton *et al.* generated an interactive activation and competition network (IAC; see Figure 1.2) based on that developed by Rumelhart and McClelland (1981). The network comprises a number of pools of units labelled to be comparable with the boxes of Bruce and Young's model. Thus there is a pool of FRUs connected to a pool of PINs, which is connected to a pool of semantic information units (SIUs) and finally, a pool of name units.

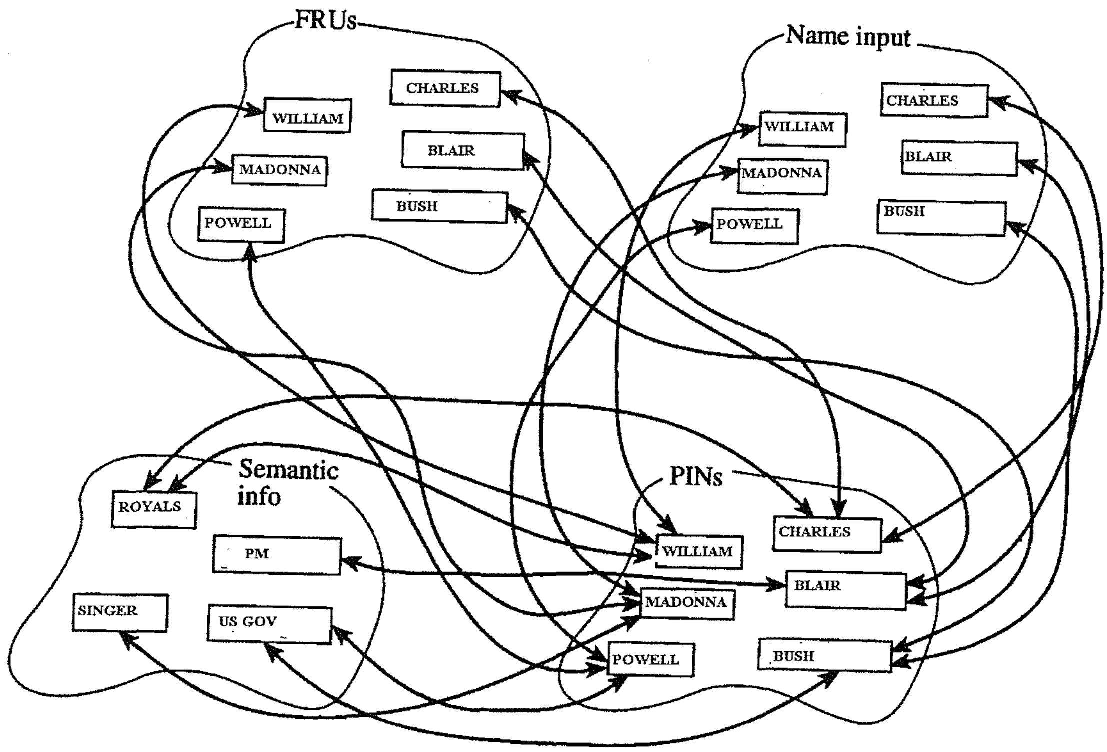

Figure 1.2: Burton, Bruce and Johnston's (1990) Interactive Activation and Competition (IAC) model of familiar face recognition (adapted from source).

The essence of the IAC model is that activation is passed between pools such that an active FRU for a known person will lead to activation of the PIN associated with that person, the semantic facts associated with that person, and finally the correct name for the person. The network is *interactive* because all units are connected via bi-directional links, which allow activation to flow backwards as well as forwards. The network is *competitive* because, within each pool of units, inhibitory links exist such that only one unit in a pool can reach activation. This reflects the fact that a face can only belong to one person.

Several changes to the Bruce and Young (1986) model have been incorporated into the IAC network, in order to improve accountability of the model. First, PINs have been re-conceptualised so that, instead of containing semantic information, they now act as a gateway to the accessing of semantic information which is held in separate SIUs. Second, the location of a sense of familiarity has also been changed. Rather than being the result of activation at the FRU level, a sense of familiarity, or knowing who you are looking at, now arises at the level of the PIN. Finally, the PIN has now been specified as receiving inputs from a variety of sources such as facial images, voices, written names, etc. These changes reflect the fact that familiarity, and so semantic information, can be retrieved following presentation of a number of different media. The proposal of a common gateway rather than separate multiple routes thus provides a parsimonious account of the situation.

With the familiarity decision now being located at the level of the PINs rather than the FRUs, the role of the FRUs now becomes one of signalling the degree of resemblance to a stored representation. If the input face provides a strong enough resemblance to a stored representation, then the relevant FRU will fire and activation will be passed on to the PIN level. Insufficient activation at the FRU level, however, results in either indecision, or failure to process the face further, and a state of familiarity is not reached.

The success of IAC in modelling familiar face recognition is that the network has made explicit the mechanisms that were alluded to in Bruce and Young's original formulation. Units are connected by links down which activation flows. When links are strong, as in the case of highly familiar faces, then activation can flow faster, with the result that response times would be quicker or activation after a given time would be higher (stronger) than if links were weak. This allows explanation of the fact that familiar faces are processed in a superior fashion (faster and more accurately) to unfamiliar ones. In addition, it is precisely this mechanism that allows IAC to explain the pattern of human performance from priming studies.

Priming refers to a facilitation of processing at 'time two' following a presentation at 'time one'. Identity (repetition) priming occurs when presentation of Mr X facilitates subsequent recognition of Mr X (i.e., Stan Laurel primes later recognition of Stan Laurel). Associative (semantic) priming occurs when presentation of Mr X facilitates subsequent recognition of the semantically associated Mr Y (i.e., Stan Laurel primes later recognition of Oliver Hardy). Identity priming can be accounted for by IAC because, when a set of connected units is activated (i.e., the FRU and the PIN for Stan Laurel) then the link between them has the potential to be strengthened. This results from application of a simple mechanism (Hebbian weight update) designed to promote association of units that are concurrently active. Subsequent presentation of the same person will thus result in faster and greater activation than if the links had not been touched (see Bruce & Valentine, 1985; Ellis, Young & Hay, 1987; Burton, Bruce, & Johnston, 1990). In a similar vein, associative priming can be explained. Presentation of Stan Laurel's face (or voice, or name) provides activation at the associated PIN and leads to the retrieval of semantically related facts. Many of these facts will be shared with Oliver Hardy and so the PIN for Oliver Hardy will become activated very slightly above its usual resting level. The consequence is that on subsequent presentation of Oliver Hardy's face (or voice, or name), PIN activation will be achieved more quickly than if Stan Laurel had not been seen. The fact that the PIN acts as a multi-modal gateway to accessing semantic information is what allows one input to prime an associated other, no matter what the medium of that input is (see Bruce, 1983; Burton, Bruce, & Johnston, 1990; Schweinberger, Burton, & Kelly, 2001; Young, Hellawell & DeHaan, 1988).

IAC has also shown the capacity to model a phenomenon that has previously eluded explanation – covert recognition. This occurs when prosopagnosic patients who are unable to exhibit overt or conscious recognition of known individual are nevertheless able to demonstrate that they knew who they were looking at. This familiarity can be exhibited through the measurement of physiological measures such as skin conductance responses (Bauer, 1984), or through patterns of priming when guessing at the familiarity of faces they cannot overtly name (Young *et al.*, 1988). A complete inability to register facial familiarity can easily be modelled by IAC by simply severing the links between FRUs and PINs. However, the case of covert recognition in patients requires that links are attenuated rather than severed, such that link strengths are minimal and the propagation of activation is restricted. The consequence is that resultant activation of a PIN following presentation of a known face is below the threshold required to indicate familiarity at a conscious level. However, sufficient activation exists to allow link strength to be temporarily increased, hence priming, or to allow subtle indicators of familiarity through physiological response.

Valentine's (1991) Multidimensional Face Space

In development of the IAC network, Burton *et al.* (1990) have given support to the basic structure provided by Bruce and Young (1986), but have also provided detail in terms of process and mechanism. A similar contribution has been made by Valentine (1991) through his conceptualisation of a multidimensional face space (MDFS). The multidimensional space framework suggests that faces are encoded within an exemplar space according to their attributes along a set of dimensions that describe the population of known faces. Valentine suggested that, across a large enough population, faces would be arranged according to a normal distribution along each dimension such that there would be a highly populated area towards the centre of each dimension and a relatively less populated area towards the extremes of the dimensions. With the centres of each dimension defining the space being co-located, this framework would predict a high density of faces at the centre of the space and a low density of faces towards the periphery (but see Burton & Vokey, 1998, for a contrasting view of the likelihood of most faces clustering at one particular area within the space).

Valentine noted that the faces that lay towards the centre of each of the dimensions would likely be typical in appearance, whereas faces towards the ends of at least one dimension were likely to be distinctive on at least that attribute. Consequently, Valentine proposed the organisation in multidimensional space as a way of accounting for the relative difficulty in processing typical over distinctive faces. In essence, Valentine suggests that the dense arrangement of typical faces at the centre of the space results in a high likelihood of confusion when trying to map an input representation onto a stored representation. The result is a slower response time for correct decisions, and/or a higher error rate when identifying typical faces than when identifying distinctive ones.

A Place for Familiarity

Both the IAC approach and the multidimensional space framework may be considered as contributing to our understanding of the processes that take place at the structural encoding stage of Bruce and Young's (1986) model, as both provide a way in which facial inputs are described. IAC achieved this by referring to facial descriptors which provided input to the FRU pool. These descriptors could either be thought of in terms of individual facial features (eyes, nose, mouth, etc.), or in terms of more global parameters such as eigenfaces (Turk & Pentland, 1991). Equally, Valentine's multidimensional space framework allows for faces to be described according to their attributes along a series of distinct dimensions that describe the population of known faces. These dimensional values could also act as inputs to the FRU pool. What is important, however, is the role that facial familiarity might play at this stage.

At a theoretical level, neither IAC, nor the MDFS makes a distinction between familiar and unfamiliar stimuli in the way a face is structurally processed. However, following structural processing, it is clear that familiar and unfamiliar faces start to be processed in different ways. This discussion implies that familiarity is the product of structural encoding in that a match at the FRU level is either found, or not. What follows from this is that facial familiarity would be unlikely to affect processing that occurs prior to the FRU level, but may affect processing that occurs at or subsequent to the FRU level. The question for this thesis is whether this bold statement of facts is correct. If it can be shown that facial familiarity does affect early visual processing of faces, then the sequence of processing stages described by Bruce and Young (1986) may have to be re-examined.

The Role of Familiarity during Structural Encoding

There are essentially two possibilities when considering the effect of familiarity on face processing. First, it is possible that familiar and unfamiliar faces are processed in identical ways but that this processing is more efficient for the familiar faces than for the unfamiliar ones. This would imply nothing more than a quantitative difference in processing across familiarity. The alternative position is that familiar and unfamiliar faces are processed in different ways and that the processing used for familiar faces is superior to that used for unfamiliar ones. This would imply a qualitative difference in processing across familiarity. If the latter position is supported, then the very interesting issue of how a face becomes familiar would need to be addressed. In other words, at what stage of familiarity might a qualitative change in processing be seen? The purpose of the present section of this review is to examine whether any evidence exists to suggest a qualitative change in processing across familiarity.

Neurological Evidence for a Qualitative Difference Across Familiarity

Single Cell Recording

The development of new techniques for brain imaging have brought about a sudden rise in the number of neurological studies in the area of face processing. At a general level, researchers have sought to identify the region of the brain responsible for face processing. Single-cell studies in animals have had some success in this pursuit. For example, Kendrick and Baldwin (1987) have identified cells in the temporal cortex of sheep that respond preferentially to facial images of other sheep. More recently, Kendrick, Da Costa, Leigh, Hinton and Pierce (2001) measured the responses from cells in the temporal and medial prefrontal cortices known to be involved on face recognition. The cells responded preferentially to sheep faces in general, but in addition, a small number of cells responded specifically to the images of familiar sheep from their own flock. Increasing familiarity further resulted in the identification of an even smaller number of cells that responded preferentially to the images of pen-mates. Although these data were obtained using a non-human species, they are suggestive of a role for familiarity in the neurological basis of face processing.

The damaging nature of single-cell recording has largely precluded its use as a means of investigation in humans. However, this procedure has been used in patients requiring, or undergoing, neurosurgery for intractable epileptic seizures (Kreiman, Koch & Fried, 2000). Although the superior temporal sulcus is the cortical area thought to be associated with face processing, Kreiman *et al.* took recordings from hundreds of single neurones in the hippocampus, entorhinal cortex and amygdala, while patients undertook a discrimination task for faces, spatial layouts, animals, cars, food, and visual patterns. Patients were required to make a face/non-face decision, and the results suggested that 20% of neurones recorded showed response to visual stimuli and 14% showed a selective response to faces. This suggests that face-selective cells are not confined to one cortical area in particular, and offers a pertinent explanation for why disorders of face recognition (such as prosopagnosia) are generally accompanied by other recognition failures (Young & Bruce, 1991).

Event-Related Potentials

More recently, measurement of event-related potentials (ERPs) has been used as a way of collecting neurophysiological evidence from a specific cortical region. This may be done in one of two ways: electrodes can be placed directly onto the cortical surface (Allison, Puce, Spencer & McCarthy, 1999; Williamson, Thadani, Darcey, Spencer, Spencer & Mattson, 1992), or more commonly, electrodes can be placed on the scalp (e.g., Bentin, Allison, Puce, Perez & McCarthy, 1996; Eimer, 2000a; 2000b; Paller, Gonsalves, Grabowecky, Bozic & Shamada, 2000). The area most widely cited as the focal point of face processing is the fusiform face area (FFA) which is located in the ventral occipito-temporal cortex. Furthermore, it appears that it is particularly the area of FFA within the right hemisphere that responds to facial stimuli. A typical ERP response can be measured approximately 150 to 200 msec following presentation of a face, with a peak at 170msecs. This 'N170' peak is thought to be associated with the structural encoding phase of face processing, with recognition of familiarity (akin to activation of an FRU) occurring a little later in the N400 and P600 peaks (Eimer, 2000a).

Functional Magnetic Resonance Imaging

A further neurological technique, and one that has recently been used to good effect when localising function within the brain, is functional Magnetic Resonance Imaging (fMRI). Using this technique, the FFA has again been isolated as being responsive to faces, but in addition, the right fusiform gyrus region appeared to be particularly responsive to the presentation of familiar faces (George, Dolan, Fink, Baylis, Russell & Driver, 1999). Furthermore, George *et al.* detailed the specific areas within the FFA where individual functions are carried out. They measured brain activity during presentation of famous and unfamiliar faces presented in photographic positive and negative. Their results led them to believe that the configuration of faces is encoded bilaterally in the posterior region of the fusiform, but activation by familiar faces is restricted to the anterior region of the right fusiform.

Positron Emission Tomography

A final technique that has been applied to the localisation of function is Positron Emission Tomography (PET). This examines regional cerebral blood flow by using a radioactive marker that can be visualised within the blood stream, while the participant performs a specific task. It is assumed that a high uptake of the marker in any one area of the brain indicates that this area is responsible for the task at hand. Using this methodology, Dubois, Rossion, Schiltz, Bodart, Michel, Bruyer, and Crommelinck (1999) examined the pattern of blood flow during a gender categorisation task for familiar and unfamiliar faces. They found a larger activation in the amygdala during presentation of unfamiliar faces which, they suggest, could indicate that the unfamiliar faces were seen as posing an emotional threat. More interestingly, however, the results suggested a decrease in the regional cerebral blood flow distribution in the earliest stages of the cortical visual system when familiar faces were presented. Thus, the authors propose that the familiarity of a face has a clear impact much earlier than recognition models (such as that of Bruce and Young) would imply as their findings question the strict independence of the visual processing and recognition routes inherent within that model.

PET scanning was also used by Wiser, Andreasen, O'Leary, Crespo-Facorro, Boles-Ponto, Watkins and Hichwa (2000) to support the idea that familiar faces activate a different area of the brain to unfamiliar faces. They used PET scanning on participants who had just learned a set of faces. Their data revealed different patterns and sizes of activation when participants viewed these newly learned faces compared to when they viewed novel distractor faces. More specifically, they suggested that the recognition of the novel faces was a frontal lobe task, whereas recognition of the familiar faces incorporated a wider area of circuitry including the inferior-medial frontal lobe, the left middle-temporal gyrus, and the left and right cuneus/lingual gyrus.

Taken together, these results highlight the importance of the FFA in the processing of facial images. More importantly, however, they point to the fact that slightly different areas of the brain are likely to be involved when processing familiar faces than when processing unfamiliar ones. This provides strong evidence in support of the possibility that a qualitative difference in the methods of processing used for familiar and unfamiliar faces may exist.

It should be mentioned at this stage, that the initial purpose of locating areas of the brain responsible for face processing was to contribute to the argument that faces were special stimuli and were processed within a discrete and dedicated part of the brain. This argument seemed to find support with the identification of the FFA. However, recent evidence casts serious doubt on the 'special' argument. Again, by using fMRI, Gauthier, Skudlarski, Gore and Anderson (2000) have demonstrated that the FFA will actually respond to any class of stimuli with which the perceiver has sufficient expertise. They demonstrated this using non-facial stimuli known as 'greebles' which possess complexity, structure, homogeneity, and are differentiated according to spacing of features, as are faces. Half their participants were trained with these stimuli while the others were not. The trained participants showed activation of the FFA when presented with greebles as well as when presented with faces. The untrained participants, however, showed activation of the FFA to faces only. Consequently, Tarr and Gauthier (2000) suggest that rather than being special, faces may characterise a 'default special' recognition class such that expertise with any stimulus set invokes neural activity which is functionally and anatomically similar to that involved in face recognition.

Empirical Evidence for a Qualitative Difference Across Familiarity

At an empirical level, the evidence pertaining to differences in processing across familiarity presents a rather mixed picture. Some studies suggest a quantitative difference in processing, others suggest a qualitative difference in processing, and others still suggest no difference at all in the pattern of processing across familiarity. It is pertinent to examine performance across subsets of high and low familiarity faces as well as across individual exemplars of high and low familiarity, and the purpose of this section is to present an overview of the empirical data in this area.

Familiar vs. Unfamiliar Exemplars

There is a large amount of literature suggesting a difference in the way that familiar and unfamiliar stimuli are processed. This has been shown for alphanumeric stimuli (Schneider & Shiffrin, 1977) as well as for faces. At a simple level, Ellis, Shepherd and Davies (1979) found that familiar faces were more accurately recognised than unfamiliar faces using both recognition memory tasks and identification tasks. These results were confirmed by Klatzky and Forrest (1984) and Ellis *et al.* warn against the assumption that familiar and unfamiliar stimuli are equivalent within an experimental context. In addition to this advantage in terms of accuracy, Scapinello and Yarmey (1970) have shown that familiar face recognition is faster, as well as more accurate, than unfamiliar face recognition.

Internal vs. External Features

So far, these results merely suggest a quantitative different in processing: familiar faces are processed better. However, a literature does exist to suggest that familiar and unfamiliar faces may be processed using different mechanisms. For example, Ellis *et al.* (1979) found that whereas an unfamiliar face could be identified as easily by its outer features (hairline and jawline) as its inner features (eyes, nose and mouth), familiar faces were more readily identified by their inner features. This was supported by the results of Young, Hay, McWeeny, Flude and Ellis (1985) who used a matching task and found, again, that familiar faces were easier to match using internal features than when using outer features, whereas unfamiliar faces showed no difference in matching performance from inner and outer features. Further, and more recent, evidence of the importance of internal features in the recognition of familiar over unfamiliar faces comes from O'Donnell and Bruce (2001). They used a same-different decision task to explore the ease with which participants could detect changes in the spacing of the eyes and mouth (internal features), or swapping the hair or chin (external features) with other exemplars. O'Donnell and Bruce found that familiarity facilitated an attentional switch from external to internal features, with the eyes being of primary importance. These results may be a consequence of attending more to inner features than outer features for communicative purposes, or it may result from learning that outer features are unreliable as indicators of identity because they can be changed so easily. Whilst both familiar and unfamiliar face processing might benefit from a shift in focus to internal features, it appears that this is only achieved with faces which are familiar, perhaps because the viewer has experience of how variable the outer features can be with familiar faces (see Young, 1984; Bruce & Young, 1998).

Order of Feature Processing

Evidence also exists to suggest that the order in which features are processed appears to differ for familiar and unfamiliar faces. Hines and Braun (1990) hypothesised that the perception of a face could be performed using either serial or parallel processing, and that this would be dependent on the familiarity of the face. They used a same-different matching task, with response times as the dependent variable. Their results suggested a pattern of processing indicative of a linear top-to-bottom scanning strategy when faces were unfamiliar. This was interpreted as evidence of a serial processing strategy taking each feature in turn. In contrast, the processing of familiar faces showed no such linear top-to-bottom trend, suggesting the use of a parallel strategy more in line with perceiving the face as a whole rather than as a collection of component parts.

Viewpoint Dependency

Studies investigating object recognition have debated whether viewpoint-dependent, or viewpoint-independent processing is used when objects are familiar (see Lawson & Jolicoeur, 1998; Newell, 1998 for opposing sides of the argument). Certainly, there is some evidence that when processing everyday objects, certain viewpoints are better than others (Palmer, Rosch, & Chase, 1981) and these have been termed *canonical views*. With respect to face processing, Hill and Bruce (1996) hypothesised that the three-quarter view of a face might be the optimal, or canonical view. Investigation by means of a recognition test using familiar and unfamiliar faces suggested, however, that familiar faces were recognised well from any viewpoint, suggesting viewpoint-independent processing. In contrast, recognition of unfamiliar faces was optimal when, indeed, the three-quarter view was used, suggesting that a canonical, or viewpoint-dependent strategy was in use.

More recent work has confirmed these results (see Hill, Schyns, & Shigeru, 1997; Newell, Chiroro & Valentine, 1999; Troje & Kersten, 1999). However, Troje and Kersten suggest that the processing of one's own face, although highly familiar, might form a special case given that one usually sees one's own face from a limited set of angles that approximate a full-face view. As such, perception of one's own face may show more viewpoint-dependency than perception of other highly familiar faces. This point has been contested by the findings of Tong and Nakayama (1999) who suggest that highly overlearned faces such as one's own face have the benefit of extensive viewing from a wealth of different angles and lighting conditions, as well as across changes associated with movement, age, expression, changing hairstyles, etc. As a result, Tong and Nakayama suggested that highly familiar faces, including one's own face, have the capacity to be processed well precisely because they have the capacity to be stored in a viewpoint-independent manner. In contrast, unfamiliar faces are restricted to a manner of processing more indicative of a viewpoint-dependent style.

Categorical Perception

Categorical perception (CP) describes the perceptual phenomenon underlying our ability to respond appropriately to the mass of highly complex and confusable stimuli that exist within our environment. The perception of these stimuli is made easier if the perceiver can carve the stimuli up into categories (Harnad, 1987). For instance, the change in colour wavelength proceeds along a continuum, but we categorise the colours along the spectrum into the distinct bands familiar to any small child: red, orange, yellow, green, blue, indigo, and violet (Bornstein & Korda, 1984). CP has been applied to the study of face perception. Early studies suggested that faces differing in expression were perceived categorically such that faces with two expressions from the same category (both happy, or both sad) were more difficult to discriminate than two faces which showed expressions from different categories (happy vs. sad). This held, even though the physical differences were manipulated between the pairs of faces in each case, using morphing software, to represent equal sized steps along the expression continuum (Etcoff & Magee, 1992).

CP effects of this nature have been taken as evidence that, even though the physical difference that exist along a dimension might be fixed and equal between adjacent stimulus pairs, they are not <u>seen</u> as being fixed and equal. Instead, when a pair of stimuli straddle a category boundary they are seen as being more different to one another than when they do not straddle the boundary. In other words, a red and an orange are seen as more different to one another than two reds, and a smile and a frown are seen as more different to one another than two smiles, even though the physical differences between the pairs is held constant. This can be explained by the possibility that familiarity with the stimulus classes enables the perceiver to become more sensitive to changes in the stimuli along the dimension relevant for classification (Aha & Goldstone, 1990; Nosofsky, 1986).

Following from Etcoff and Magee's (1992) demonstration of CP effects for faces, CP effects have now been demonstrated across race (Levin & Beale, 2000), gender (Campanella, Chrysocoos & Bruyer, 2000), and more pertinently for this thesis, across familiarity (Beale & Keil, 1995). In this latter experiment, Beale and Keil generated a continuum of facial images by morphing between two familiar identities, John F Kennedy and Bill Clinton. Participants were presented with images from along the continuum and were required to identify each face as either Kennedy or Clinton. Following this, participants completed a traditional ABX discrimination task in which a pair of stimuli (A and B) were followed by the presentation of a single test stimulus (X). The task was to say whether X was the same as A or B. Performance indicated that when A and B were drawn from the same side of the Kennedy-Clinton continuum, completion of the ABX task was much harder than when A and B were drawn from different sides of the continuum. This indicated that perceptual discrimination of the stimuli was much harder when they fell within the same identity category than when they were perceived as coming from different identity categories.

Investigations using ERP data have been used to establish that the CP discrimination occurs approximately 150 ms following stimulus onset (Campanella, Hanoteau, Depy, Rossion, Bruyer, Crommelinck & Guerit, 2000). This is well within the time frame considered plausible for structural encoding to take place. What is important, however, is the effect that stimulus familiarity had on this pattern of CP. The use of continua of stimuli morphed between familiar figures of differing levels of familiarity (Kennedy-Clinton; Townsend-Stallone) suggested that the magnitude of the CP effect was indeed, affected by stimulus familiarity. Effects were larger when the identities at either end of the continuum were highly familiar, than when less familiar, and the effects were completely absent when a continuum between two unfamiliar faces was used (Beale & Keil, 1995). This points to the suggestion that, again, facial familiarity may be influential at the very earliest stages of perceptual processing for facial images.

Caricature Advantage

The caricature advantage refers to the better recognition of a stimulus after it has been caricatured than before. This was first demonstrated using computer software (Brennan, 1985) to generated line-quality caricatures of graded levels of distortion relative to a veridical or undistorted image. Using this approach, Rhodes, Brennan and Carey (1987) showed that the recognition of familiar figures could be achieved more quickly when viewing a caricatured image than when viewing an undistorted one. Subsequent studies suggested that the caricature advantage was robust when images were of a reduced level of complexity, as exemplified by line drawings (see Benson & Perrett, 1991; Stevenage, 1995a; 1995b). However, when photographic quality images were used, the advantage shown when recognising caricatures was small and at times, not evident (Benson & Perrett, 1994). A change of paradigm, however, from the use of a recognition task to the use of a perceptual task in which participants were asked to rate the best-likeness of a given individual, or to generate an optimal image by manipulating level of caricaturing, revealed very interesting results. Overall, it was evident using these perceptual tasks that the optimal image generated by participants incorporated a degree of caricaturing (average = + 4.4% distortion from veridical).

What was even more interesting was the role that facial familiarity played in the demonstration of a caricature advantage. In terms of performance in a recognition task (Rhodes & Moody, 1990) and a best-likeness task (Rhodes, Brennan & Carey, 1987) there was no evidence for any influence of caricaturing when faces were unfamiliar. Furthermore, when participants were required to generate optimal images of a range of public figures who varied in terms of their familiarity, a higher degree of caricaturing was incorporated into the optimal image when faces were highly familiar than when faces were less familiar (Benson & Perrett, 1991). Again, these data would suggest that familiar and unfamiliar stimuli may be processed in qualitatively different ways and so respond to experimental manipulation such as caricaturing to differing degrees.

Inversion Effects

By far and away the most frequently used manipulation within the face literature is stimulus inversion. This can be achieved either by rotating the face through 180 degrees, or by flipping the face in the vertical axis. The only difference across these two manipulations is in the resultant left-right positioning of features. The advantage of using stimulus inversion as a mechanism for disrupting presentation is that, notwithstanding the left-right shift in the positioning of features, the face remains structurally unaltered and so is entirely matched with upright presentations, but the perception of the face is affected quite substantially. The mechanisms governing this inversion effect are well researched and will be reviewed below. Following this, the empirical data will be examined in order to evaluate the impact that stimulus inversion has on face perception in general, and on the question of processing differences across familiarity in particular.

The first demonstration of the disruptive influence of inversion on face perception was presented by Yin (1969). He compared recognition memory for a range of visual stimuli presented upright and inverted and suggested that memory for faces was far more affected by inversion than memory for non-facial objects such as houses, aeroplanes and running stick-men. This became known as the *differential inversion effect* and the finding has since been confirmed by a series of studies (i.e., Scapinello & Yarmey, 1970; Yin, 1970).

Accounting for the Inversion Effect

Explanations for the inversion effect are varied and essentially, the task is to explain why faces are affected by inversion over and above the effect seen for other mono-oriented classes of stimuli. One suggestion is that inversion of a face leads to a loss of ability in detecting facial expression (Kohler, 1940; Yin, 1970). Whilst it is indeed difficult to determine expression from an inverted face, the fact that expression analysis and identification analysis have been shown to be independent pathways (see Bruce & Young, 1986) casts doubt on this as a way of accounting for the inversion effect.

Another suggestion is that inversion of a face leads to an inability to perceive the relationships that exist between the face and its frame of reference (Rock, 1973). In this respect, Rock suggested that the encoding of an object is executed with reference to an external (environmental) frame of reference, which allocates directions to the spatial properties of the object. Therefore, the assignations of top, bottom, and sides of the object are dependent upon the frame of reference. If a mono-oriented stimulus is inverted, encoding is disrupted because the frame of reference differs from standard (top becomes bottom and vice versa). As a result, the object comes to look different. Parks, Coss and Coss (1985) express this in terms of an interaction of the object-centred frame of reference, and the viewer-centred frame of reference based on the external environment or the position of the viewer. In the case of facial inversion, these two frames are at odds with one another and difficulty of processing arises.

Accounting for Inversion Effects with Configural Processing

Whilst the 'frames of reference' argument appears adequate in describing the impairment caused by inversion, it doesn't really provide an explanation of the effect in terms of likely mechanisms. Several lines of evidence do, however, converge to suggest that what is impaired when faces are inverted is the ability to perceive the 'whole' and as such, the ability to perceive the relationship of features to one another. This perception of the spatial layout of a face, or the relative spacings between features, has been referred to as 2^{nd} order relational processing, or configural processing (see Maurer, Le Grand & Mondloch, 2002 for a discussion of this distinction). It is contrasted with piecemeal (or featural) processing in which the facial features are considered independently of one another. A piecemeal strategy may be perfectly useful for recognising faces especially if the face is distinctive on a particular individual feature (such as Dennis Healey's eyebrows). However, when faces are more typical, reliance on piecemeal information in this way is unlikely to lead to successful processing and configural information is required. This enables the perceiver to distinguish between two faces on the basis of the fine layout of features rather than merely on the features themselves.

Configural processing can be demonstrated in a number of ways. First, the *face superiority effect* suggests that identification of a facial feature is best achieved when that feature is presented within the context of a whole face rather than when presented alone or in a scrambled face (Homa, Haver & Schwartz, 1976; Mermelstein, Banks & Prinzmetal, 1979). This gave rise to later demonstrations of what became known as a complete-over-part advantage (see Figure 1.4) for processing faces (Bruce, Doyle, Dench & Burton, 1991; Davidoff & Donnelly, 1990; Donnelly & Davidoff, 1999; Tanaka & Farah, 1993; Tanaka & Sengco, 1997).

Together these indicate that something is gained from processing the arrangement of features within the face over and above the processing of the features themselves. In other words, the capacity to engage in configural processing over and above piecemeal processing leads to improvements in subsequent performance.

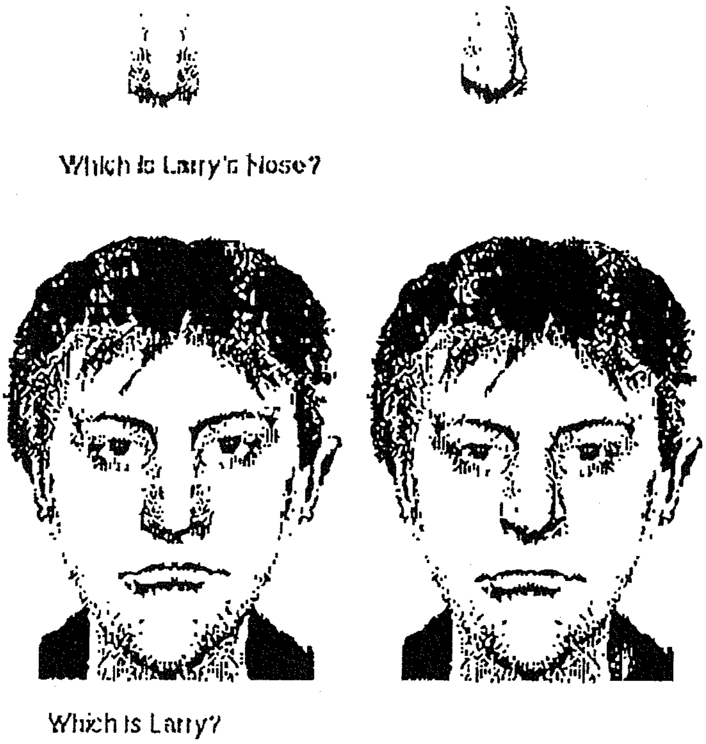

Figure 1.3: Illustration of the complete over part advantage when processing faces. Parts are recognised more efficiently when presented in the context of a whole face than when presented alone (from Tanaka & Farah, 1993).

Configural processing can also be demonstrated using what are termed *chimeric faces* (Young, Hellawell & Hay, 1987; see also Sergent, 1984). Young *et al.* created these faces using the top half of one familiar face and the bottom half of another familiar face. The two halves were either aligned so that the impression of a new composite face was created, or they were slightly off-set so as to reduce the impression of a complete new face. Young *et al.* found that identification of the original identity from the top half of the composite was incredibly difficult when the two halves had been aligned correctly. The authors felt this to be due to the fact that a strong tendency to perceive the 'whole' had been created because the configuration of the features was appropriate to that perception and dominated the percept. As such, it was difficult to pull apart the two halves of the face and consider the features separately to make the necessary identification. In contrast, when the two halves had been misaligned, the capacity to identify the original identity of the top half of the face was markedly improved, and this was felt to reflect the fact that, through misalignment, the face was seen as a collection of parts rather than as a complete whole (Figure 1.5). In this way, the respective halves were easier to disentangle and recognise than if they had been bound together perceptually by configural processing.

Figure 1.4: Illustration of Chimeric Faces (i) aligned to generated a composite, and (ii) mis-aligned to generated a non-composite (from Young, Hellawell & Hay, 1987).

Finally, configural processing has been demonstrated using a phenomenon known as the *Thatcher Illusion* (Thompson, 1980; see also Bartlett & Searcy, 1993; Searcy & Bartlett, 1996). In this demonstration, the eyes and mouth are removed, rotated and replaced in their original positions within an otherwise upright face. As such, the features of the face have remained the same, but the configuration of those features has been broken. The resultant image appears grotesque when viewed upright as a result of the obvious and extreme disruption to the configural relationship between the features forming the face. However, when viewed inverted, the grotesqueness of the image is not generally perceived (Figure 1.5).

These three demonstrations of configural processing are important when accounting for the influence of inversion. Inasmuch as face inversion is believed to disrupt the ability to relate the entire face to its frame of reference, so it is believed that inversion also disrupts the ability to relate one part of the face to another part and so perceive the essential configural information within the face (Diamond & Carey, 1977; 1986; Leder & Bruce, 2000; Rhodes, Brake, & Aktinson, 1993). Indeed, all of the above demonstrations – the face superiority effect (and complete-over-part advantage), the difficulty in identifying the top half of a chimeric face, and the perception of grotesqueness in a Thatcherised face, are all obliterated when the respective stimuli are presented upside down. This is taken as strong evidence that what is lost when a face is inverted is the ability to process the configural information or the spatial layout of features within the face

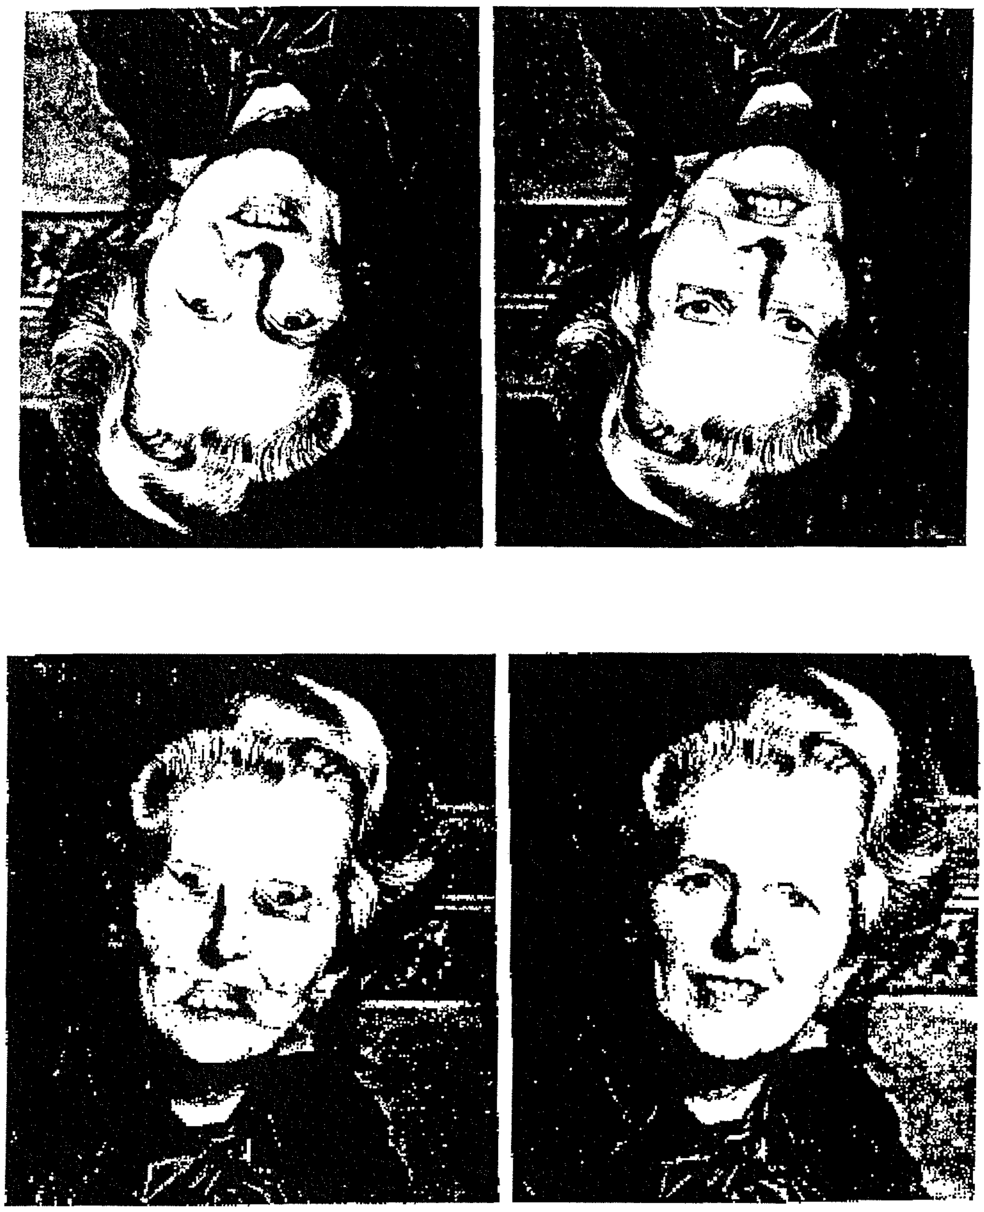

Figure 1.5: Illustration of the Thatcher Illusion. Grotesqueness is readily perceived when the face is upright but is lost when the face is inverted (from Thompson, 1980).

The Impact of Inversion across Familiarity

Having defined the role of configural processing in face perception, and demonstrated that the ability to process configural information is removed by inversion, it is now important to determine whether inversion has a consistent effect on the processing of familiar and unfamiliar stimuli. If processing of the two sets of stimuli are impaired to equivalent degrees by inversion, then it is likely that both familiar and unfamiliar stimuli are reliant on this configural method of processing. If, however, one set of stimuli is more affected by inversion than the other set, then it is possible that a different pattern of reliance on configural and piecemeal strategies is evident across familiarity.

The data in this respect are actually very surprising when set in the context of the previous evidence. Three studies have been published, and all three have cast doubt on the premise that a processing difference may exist across familiarity. Each of the studies can, however, be questioned at a methodological level. First, Scapinello and Yarmey (1970) compared the recognition of upright and inverted faces, dogs' faces, and buildings. Following Yin's (1969) proposal, they included a measure of stimulus familiarity by presenting previously novel items either once (low familiarity) or seven times (high familiarity) at study, and then noting performance at test. Their findings supported those of Yin (1969; 1970) in that faces were disproportionately affected by inversion compared to dogs and buildings. However, the data suggested no interaction of the inversion effect with stimulus familiarity. In other words, both familiar and unfamiliar faces were equally affected by stimulus inversion. This would suggest that there is no qualitative difference across familiarity in the method of processing used. Rather, all faces were processed in a configural way and processing was impaired when stimuli were inverted.

Goldstein and Chance (1981) criticised Scapinello and Yarmey's (1970) use of the familiarity variable, stating that familiarity could not adequately be manipulated in this way. Rather than being manipulated by the experimenter, Goldstein and Chance believed that familiarity was a commodity brought to the experiment by the participant. Yarmey (1971) conducted a further experiment which addressed this point. He repeated the old/new recognition task used by Scapinello and Yarmey (1970) but used faces that were either pre-experimentally familiar (through being in the public eye) or were unfamiliar. His results again demonstrated a large effect of inversion when processing faces. Again, however, this inversion effect was not modified by the familiarity of the stimulus. Ceiling level performance may, however, have limited the potential for an influence of familiarity to become evident and thought should be given to repeating this experiment under more difficult conditions to remove the ceiling performance when faces are familiar.

More recently, Collishaw and Hole (2000) have re-examined the issue of the effect of inversion across facial familiarity. This was actually embedded within a larger experiment in which the effects of inversion were compared with scrambling and blurring, separately, and in each of three pairwise combinations. The design of this experiment allowed the researchers to identify the relative contribution of configural processing (affected by inversion and scrambling) and piecemeal processing (affected by blurring) when processing familiar and unfamiliar faces. Across two experiments, Collishaw and Hole examined the influence of image manipulation on (i) a speeded familiarity task for familiar faces, and (ii) a speeded old/new recognition task for unfamiliar faces. The data suggested that both configural and piecemeal processing together provided the optimal processing conditions. However, processing could still continue, albeit more effortfully, when only one or other source of processing was available. Processing was, however, incapacitated when neither configural nor piecemeal processing was possible, as exemplified in the blurred + scrambled, and blurred + inverted conditions. More importantly for the present discussion though, the results demonstrated that inversion, along with all other stimulus manipulations, had a consistent effect on processing across facial familiarity. Again, the suggestion from these results is that familiar and unfamiliar faces are equally reliant on configural and piecemeal processing strategies and that no global process difference exists. However, it is possible to question the comparability of a speeded familiarity task (recognition of familiar faces) with an old/new recognition task (recognition of unfamiliar faces) in terms of the mental processes required by each.

Finally, it may be pertinent to examine the effect on inversion not across familiar and unfamiliar individuals' faces but across familiar and unfamiliar sets of faces. The study of inversion for own- and other-race faces provides exactly this sort of enquiry (see Meissner and Brigham, 2001, for a meta-analytic review of the studies within this area). Consideration of performance with own- and other-race faces is instructive within the present context because own- and other-race faces provide natural and comparable stimulus sets with which participants have more or less familiarity. As such, the perceptual processing of stimuli differing in levels of familiarity can be investigated without the influence of semantic knowledge that can be a factor when using 'known' faces.

Several studies have investigated the influence of inversion when processing own- and other-race faces. The results, however, have been somewhat mixed. Rhodes, Brake, Taylor and Tan (1989) suggested that the familiarity of the own-race faces may mean that perceivers are more able to extract and process configural information for these faces compared to other-race faces. Consequently, inversion, and therefore removal of configural processing, should impair the recognition of own-race faces to a greater extent. The results supported their predictions: own-race face perception was indeed more affected by inversion than other-race face perception. However, the results were in direct opposition to those previously found by Valentine and Bruce (1986*a*). Indeed, the latter study found that inversion disrupted face processing to an equivalent degree no matter what the race of the face.

Rhodes *et al.* account for this contradiction by pointing to the methodological differences between the two studies. While Rhodes *et al.* presented own-race and other-race faces for exactly the same length of time at study, Valentine and Bruce allowed a longer inspection phase for other-race faces on the grounds that participants' lack of experience with these stimuli would be a disadvantage in this condition. Valentine and Bruce's results suggested a similarly disruptive effect of inversion for both races of face, but taking into account the extra viewing time for the other-race faces, the authors proposed that recognition would have been more disrupted for the other-race set had the inspection times been equal. Notwithstanding this reasoning, the data on the effect of inversion across familiarity (and across race) are considerably mixed, and cast doubt on the earlier belief that different processing mechanisms may be at play across familiarity.

Summary of Evidence on Process Differences across Familiarity

Before moving on to consider explanations for any role of familiarity, it may be useful to summarise the data presented thus far. At a neurophysiological level, evidence was presented to suggest that faces are primarily processed within the fusiform face area of the right hemisphere. This is not a face-specific area of the brain as it is activated when experts are presented with examples of non-facial objects that have complexity, mono-orientation, homogeneity, and require differentiation on the basis of the spatial layout (or configuration) of elements. What was interesting, however, was that across a range of measurement techniques, the data suggested that familiar faces activated a somewhat different area of the brain than unfamiliar faces and that this could not merely be attributed to a differential involvement of an emotional response (within the amygdala or hippocampal area) or a mental representation (frontal lobe task).

Consideration of the empirical data again gave a strong indication that familiar and unfamiliar faces may be processed somewhat differently. However, here, the data were substantially more mixed. On the one hand, evidence existed to suggest that the familiar face was processed in a rapid, parallel, and viewpoint-independent fashion, with emphasis on internal over external features, and sensitivity to minute changes underlying CP and caricature effects. On the other hand, stimulus manipulation through inversion had a mixed effect depending on whether performance was examined for familiar face sets, or familiar individuals. Consideration of familiar face sets (own-race faces) gave some (but weak) evidence to confirm a difference in inversion effects across familiarity. In contrast, familiar individual faces suggested no difference in method of processing across familiarity.

Consequently, while it is tempting to conclude that there is a qualitative difference in processing across familiar and unfamiliar faces, the cautious position cannot be ignored, that faces are processed in the same way but with nothing more than a quantitative difference in performance. These conclusions are illustrated diagrammatically in the interim flowchart presented below. The purpose of the following section is to take these two positions as a basis for theorising about the processes at work, and the predictions for the current programme of study.

Is there any evidence for a difference in processing across familiarity?

Yes | No

Neuropsychological:	Empirical:	Empirical:
• Differential activation of FFA • Differential involvement of early visual structures	• Internal/external features • Viewpoint dependency • Order of feature processing • Categorical perception • Caricature effects • Inversion (across race)	• Inversion (across faces)

Theoretical Basis for Face Processing across Familiarity

The wealth of evidence reviewed above gives, in some sense, good grounds for accepting the position that familiar and unfamiliar faces are processed according to slightly different mechanisms. The task facing us now is to explain the basis for such a difference. With reference to the literature, it may be pertinent to consider a difference at one of two levels – a process level, or a representational level. At a process level, one candidate that emerges for consideration is that familiar faces are processed with a greater degree of reliance on configural processing than piecemeal processing compared to unfamiliar faces. At a representational level, the notion of robust representations as advanced by Tong and Nakayama (1999) may provide a fruitful basis for explanation and prediction. The conservative position evidenced by the empirical work on inversion effects is, however, that no processing difference exists across familiarity and instead, a common set of processes are used with more (familiar) or less (unfamiliar) efficiency.

The following sections will review each of these possibilities in turn, and will outline a way in which a programme of research may be developed to decide between these possibilities. In reviewing this material, however, it should be noted that although a relatively large amount of literature exists on the topic of configural processing, this should not be taken as an indication of the dominance of this strand of theorising over the remaining two strands within the current thesis.

Differences at a Process Level

The literature reviewed already has been suggestive of a difference in processing strategy across familiarity. Familiar faces appeared to be processed in a way which was more sympathetic to the configuration of internal features, or to the face as a whole whereas unfamiliar faces appeared to attract a more piecemeal level of processing. This distinction between configural and piecemeal processing across stimuli is very important and is not without precedent within the literature. Indeed, a configural/piecemeal processing shift has been implicated when accounting for expertise effects (Diamond & Carey, 1977; Sergent, 1984).

A wealth of evidence can be used to support the notion that perceivers become expert by use of configural processing. For instance, the use of configural processing is believed to underlie the ability of experts to demonstrate a preference for subordinate level processing over basic level processing of a stimulus class (Gauthier, Skudlarski, Gore & Anderson, 2000; Gauthier, Tarr, Adam, Anderson, Skudlarski & Gore, 1999; Tanaka & Gauthier, 1997; Tanaka & Taylor, 1991). This is because it is believed that the experts have the ability to process the stimuli at a configural level and hence to notice the fine-grained distinctions that allow the discrimination of one exemplar from another as is required for subordinate level decisions. In contrast, novices who lack the ability to pick up on, or process configural information may merely have sufficient perceptual skill to identify the class of stimulus rather than the exemplar within the class. This notion echoes that of categorical perception discussed earlier. It also echoes the distinction between type and token in object recognition, as discussed by Rosch, Mervis, Gray, Johnson, and Boyes-Braem, (1976).

In addition, the use of configural processing by experts within a stimulus class is believed to underlie the caricature advantage. Sensitivity to configural spacings defining the layout of features within a stimulus brings with it a sensitivity to the exaggerations of such spacings. Accordingly, it should not surprise us that experts within a stimulus class show a greater advantage when presented with caricatured stimuli than novices do. This has been demonstrated for face perception, but also for the perception of birds, as long as participants are bird experts and the stimuli are sufficiently homogeneous as to benefit from caricaturing in the first place (Rhodes & McLean, 1990).

Finally, the use of configural processing by experts is believed to underlie the inversion effect. The fact that inversion effects can be shown for non-facial stimuli such as dogs, but only when perceivers are dog experts (Diamond & Carey, 1986), or for greebles, but only when perceivers have received a minimum of 10 hours intensive training (Gauthier & Tarr, 1997; Tarr & Gauthier, 2000) suggests that again, the use of configural processing should be linked with expertise rather than merely with face processing. Indeed, Diamond and Carey propose that experts process familiar items in a qualitatively different way to novices, and that inversion of stimuli prevents this processing, reducing experts to novice levels when stimuli are inverted. In this sense, rather than being a special case, faces are just one stimulus set for which inversion will be disproportionately disruptive.

In line with this discussion, the effects of categorical perception, caricaturing, and inversion can all be seen as markers of expertise, in that all effects are shown not just when processing faces, but when processing any stimulus with which one has expertise. It is not a far cry to think of familiarity as being a further operationalisation of expertise, such that familiar faces show all the hallmarks of expertise to a stronger degree compared to unfamiliar faces. If this is correct, then one might expect the perception of familiar faces to show a greater reliance on configural over piecemeal processing than is evident for unfamiliar faces. What would follow is that any stimulus manipulation designed to upset configural processing should have a greater impact on familiar faces than unfamiliar ones. **According to the expertise reasoning then, a manipulation such as stimulus inversion should affect familiar face processing more than unfamiliar face processing.**

Differences at a Representational Level

In contrast to the location of a difference across familiarity at a process level, it is possible that differences in processing across familiar and unfamiliar faces arise because of the logical differences that exist at a representational level. The work of Tong and Nakayama (1999) is important here.

At a very basic level, one difference that exists between familiar and unfamiliar faces is that familiar faces possess a mental representation prior to the experiment while unfamiliar faces do not. In order to complete any task such as an old/new task for unfamiliar faces, a mental representation has to be created and this, necessarily, can only be informed by a single presentation. Along with that comes the fact that the representation is limited to reflecting a single orientation, viewpoint, expression, lighting and grooming style. In contrast, a highly familiar face is likely to be represented by a mental representation that captures the richness of one's visual experience with the stimulus face. Whilst it is possible that multiple instances of the individual are stored alongside one another, Smith and Medin (1981) have suggested that a composite of stored exemplars and some sort of abstracted model is more likely. More recently, the notion of prototype formation for faces has been demonstrated empirically (Bruce, Doyle, Dench, & Burton, 1991; Cabeza, Bruce, Kato, & Oda, 1999; Solso & McCarthy, 1981).

Tong and Nakayama (1999) have conceptualised of these rich representations for familiar faces as being akin to *robust representations*. They can be thought of as analogous to the FRU within Bruce and Young's (1986) Information Processing Framework, or Burton *et al.*'s (1990) IAC network. It was considered that robust representations had the opportunity to be formed and laid down when faces were highly familiar, or overlearned. Specifically, a robust representation would be expected for one's own face. On empirical trials, Tong and Nakayama (1999) found that such highly overlearned faces were processed faster whether presented in front view, three-quarter view, or in profile. The faces were not only more quickly detected when targets (compared to less familiar faces) but were also more quickly rejected when they appeared as distractors.

Tong and Nakayama proposed that these effects were reliable in that they generalised across experimental tasks. Most importantly, perhaps, they indicated that faces represented by such robust representation should be processed better than faces represented by weaker representations, no matter what the viewing conditions. In this sense, efficiency of processing is tied to the quality of the representation rather than the use of a more efficient processing strategy. Indeed, the possession of a robust representation may be seen as a way of buffering the processing against the deficits that may otherwise result from less-than-ideal viewing conditions. Of course it could be the case that the existence of a robust representation for a familiar face is exactly what permits the use of a more efficient style of processing. This more efficient processing style is, however, unlikely to be thought of in terms of configural versus piecemeal strategies, as Tong and Nakayama's data indicate the survival of superior processing for familiar faces even when stimuli are inverted. **The prediction to result from this line of reasoning is thus that a manipulation such as stimulus inversion should affect familiar face processing less than unfamiliar face processing.**

No Qualitative Difference in Processing across Familiarity

Against the two sets of predictions that have been outlined above, it is necessary to include the conservative stance and suggest the null hypothesis that no qualitative difference in processing across familiarity exists. This would be the position supported by the empirical literature examining inversion effects with familiar and unfamiliar target individuals. It would be naïve to disregard these data in preference for the data that show a consistent pattern. However, the fact that these null results have been derived exactly from the studies that have used inversion as a manipulation, is disturbing in terms of the coherency of the literature in the field. Given this, and given the concerns raised previously regarding the methodologies used in the three reported studies, the decision was taken here to use inversion as the primary tool to investigate familiarity effects. As such, this thesis permits a re-examination of the three contentious studies, using inversion across familiarity, and allows for a test of the 'process' and 'representation' predictions provided above according to a well-researched method of disrupting configural, and so expert, processing. These positions are summarised in a completed flowchart of reasoning provided below.

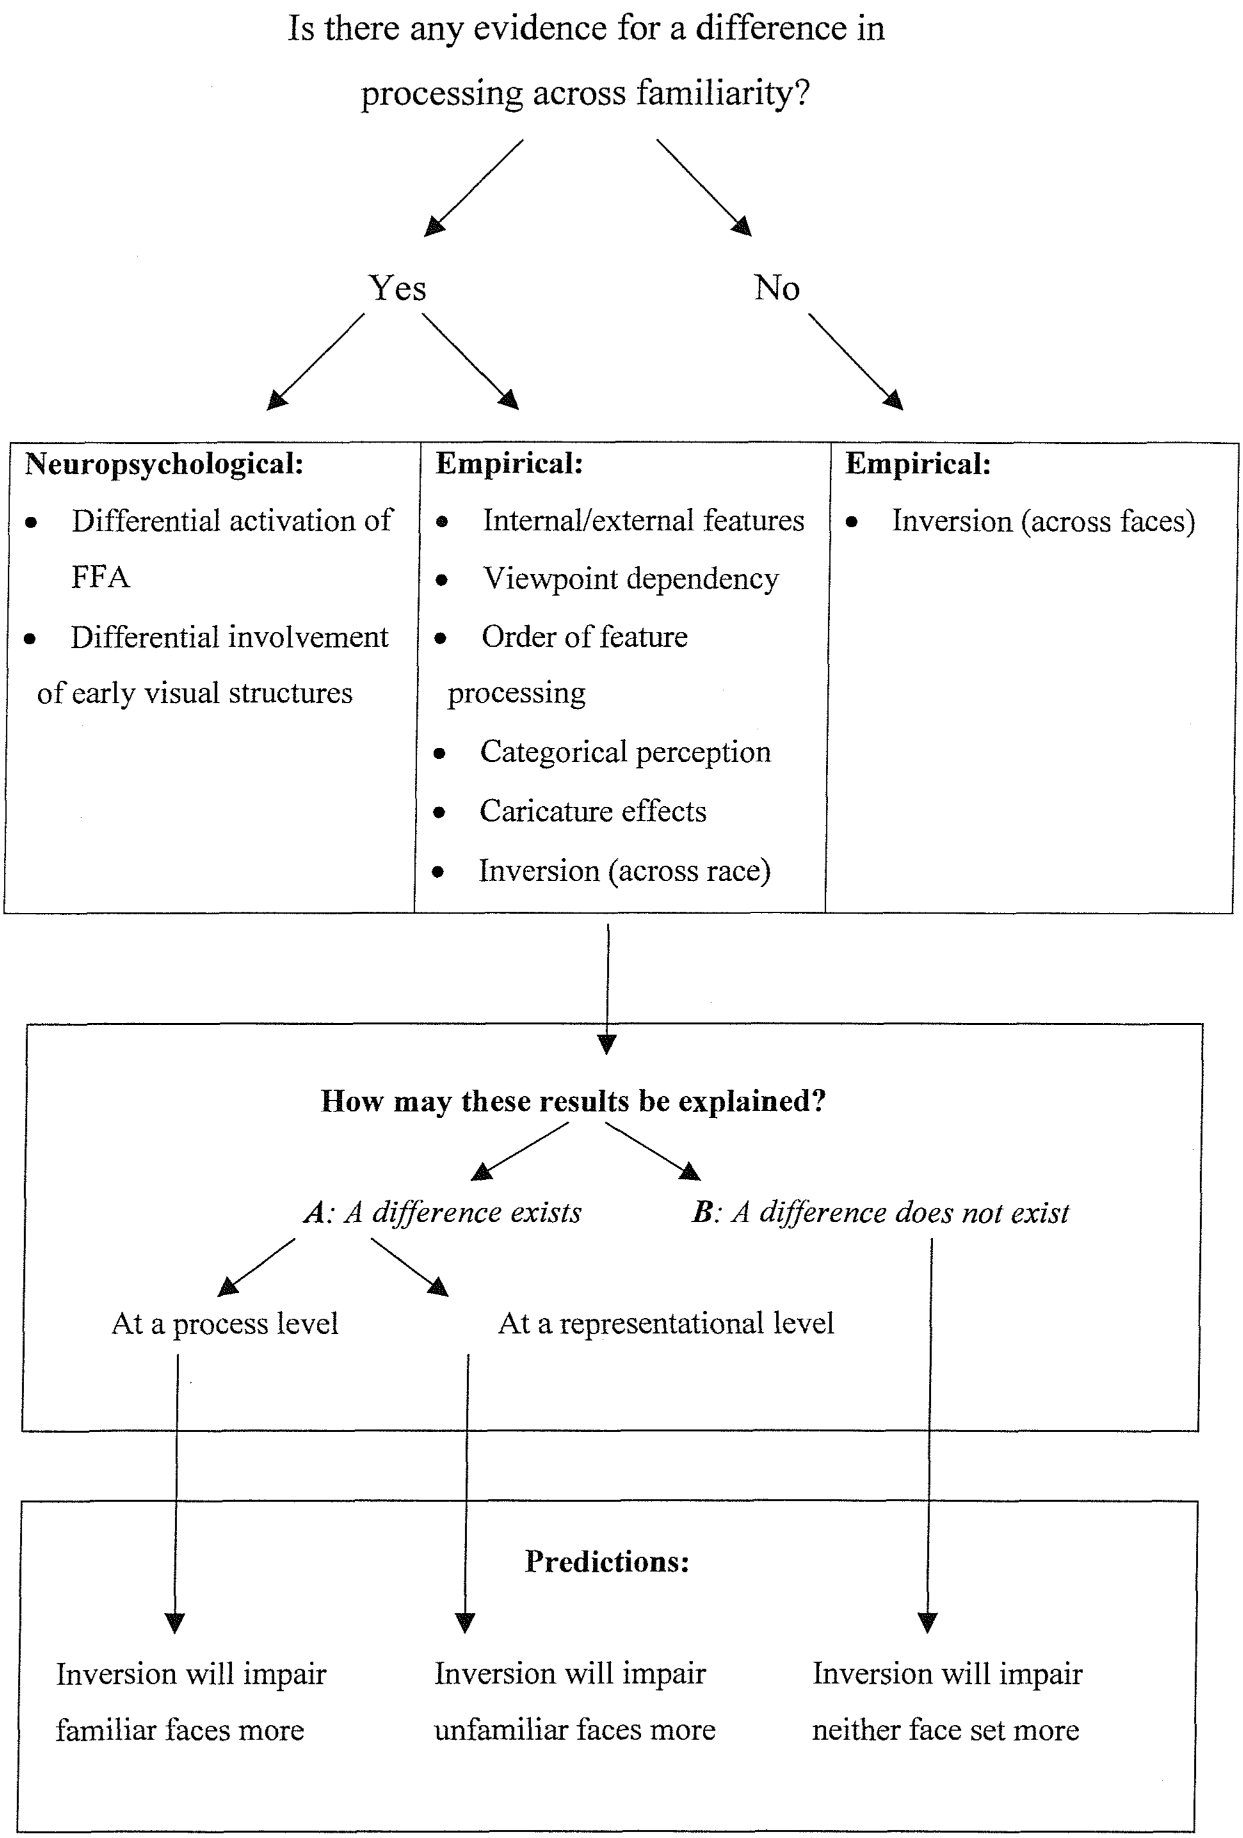
Is there any evidence for a difference in processing across familiarity?
Yes
No
Neuropsychological:
• Differential activation of FFA
• Differential involvement of early visual structures
Empirical:
• Internal/external features
• Viewpoint dependency
• Order of feature processing
• Categorical perception
• Caricature effects
• Inversion (across race)
Empirical:
• Inversion (across faces)
How may these results be explained?
A: A difference exists
B: A difference does not exist
At a process level
At a representational level
Predictions:
Inversion will impair familiar faces more
Inversion will impair unfamiliar faces more
Inversion will impair neither face set more

CHAPTER 2

Higher Level Decision Making: Does Facial Familiarity Affect The Ability to Recognise Upright and Inverted Stimuli?

Abstract

The inversion effect has been acknowledged as a consistent phenomenon within face recognition research. The aim of this study was to examine the stability of the effect of inversion across faces of different familiarity. Based on the evidence presented in the previous chapter, it was predicted that familiarity could influence recognition performance in one of two ways: either recognition of inverted familiar faces would be impaired relative to unfamiliar faces (the expertise hypothesis), or recognition would be preserved when faces were familiar (the robust representations hypothesis). A speeded familiarity task was used to investigate the influence of familiarity. Sensitivity, bias, speed, and accuracy of response were all recorded. Analysis of error rates suggested that inversion was more disruptive to recognition of familiar than unfamiliar faces, supporting the expertise hypothesis. However, evidence of a speed-accuracy trade-off and biased responding led to the results being treated with caution.

Higher Level Decision Making: Does Facial Familiarity Affect The Ability to Recognise Upright and Inverted Stimuli?

Consistent with the material presented in the Introduction, inversion is used within this experiment as a way of increasing the difficulty of the face task. The essence of the research question is whether facial familiarity lends the viewer something that enables them to overcome the difficulty of the task.

Inversion has been used extensively in the previous literature as a way of manipulating task difficulty in recognition tasks. Early studies suggested that inversion had a clear and negative effect on the recognition of faces (Goldstein, 1965; Hochberg & Galper, 1967). Yin (1969) confirmed these findings but in addition, he demonstrated that recognition of faces in an old/new task seemed to be more affected by inversion than was recognition of other classes of mono-oriented stimuli (such as running stick men, aircraft, and houses). Yin termed this the *differential inversion effect* (DIE) and concluded that the recognition of faces was more affected by inversion compared to non-facial stimuli.

More than a decade after Yin's results, Diamond and Carey (1986) reasoned that the DIE shown for faces may be the result of greater familiarity with faces as a stimulus set. This was examined by looking at the effect of inversion on the recognition of (i) faces and (ii) a non-facial stimulus set (dog profiles) with which participants were either familiar or not. As expected, their results showed a clear impairment in the recognition of faces when presented inverted as opposed to when presented upright. In addition, the recognition of dog profiles was also clearly and significantly affected by inversion but only when the participants were dog experts. Their results led Diamond and Carey to propose that, although the differential inversion effect with faces may be interpreted as evidence that faces are special, it may be more correct to conclude that impairment due to inversion will generalise to any class of homogeneous stimuli with which we have sufficient knowledge or expertise.

The most likely explanation for this effect of inversion when processing stimuli with which we are expert is that inversion disrupts an orientation-specific process that has been refined for the processing of stimuli which are usually seen upright. Whilst some believe that this process is linked with extraction of expression

(Kohler, 1940; Yin, 1969; 1970) or extraction of the relationship of the stimulus to the environment (Parks, Coss & Coss, 1985; Rock, 1973), much more evidence exists to suggest that the perception of the relationships between the features is impaired by inversion. This may be due to the difficulty in extracting shape from shading from inverted faces (Johnston, Hill & Carman, 1992). However, it may simply be more difficult to extract a sense of the configuration of features when the face is inverted (Diamond & Carey, 1977; 1986; Farah, Tanaka, & Drain, 1995; Leder & Bruce, 2000; Rhodes, Brake & Atkinson, 1993; Sergent, 1984; Thompson, 1980; Young *et al.*, 1987). If this premise is correct, then the expertise framework would suggest that stimuli with which one is more familiar with and which are processed primarily in a configural manner, should be more affected by inversion than stimuli with which one is less familiar.

The Use of Inversion Within a Stimulus Class

With this background of research using inversion as a way of disrupting recognition of facial stimuli, it is pertinent to ask whether the effect of inversion should be considered stable and consistent across different faces or whether greater familiarity with a sub-set of stimuli may afford some defence against disruption. In this light, the literature on inversion effects for own- and other-race faces becomes relevant. Based on the prediction that expert recognition will be more affected by inversion than novice recognition, Rhodes, Brake, Taylor and Tan (1989) predicted that the recognition of own-race faces would be more affected by inversion than the recognition of other-race faces. Using stimuli that were unfamiliar to the participant, and using a standard old/new recognition procedure, Rhodes *et al.*'s results broadly confirm the predictions. Their first experiment reveals a greater inversion effect for own-race than for other-race faces when considering RT data only (accuracy was at ceiling). Their second more difficult experiment shows the same effect when considering accuracy only, but the effect was only evident for a subset of the stimulus pairs. As such, these results can only be interpreted as weak support for the expertise prediction.

Some doubt is, however, cast on Rhodes *et al.*'s results by the previous work presented by Valentine and Bruce (1986). Valentine and Bruce also used an old/new recognition procedure with own- and other-race faces. However, to reduce the risk of floor effects, the exposure duration for the other-race faces was longer than that for the own-race faces. Their results showed an equivalent effect of inversion regardless of the race of the face. The authors, however, suggest that if the exposure time was equalised for both sets of faces then the effect of inversion would have been greater for the other-race set – an assumption that is clearly not in line with Rhodes *et al.*'s demonstration of a greater inversion effect for the own-race set.

Although the results are mixed, the above studies give credence to the suggestion that inversion can be used as a tool to investigate processing differences *within* a stimulus class. As such, it would appear valid to use inversion here as a way of determining whether familiarity *per se* can influence higher level recognition processes. Three published studies are directly relevant to this point. First, Scapinello and Yarmey (1970) investigated whether familiar and unfamiliar faces were equally affected by inversion. In an old/new recognition task, the results suggested that fewer errors were made on upright than inverted faces, and on familiar faces than unfamiliar faces. However, the extent of inversion effect was equivalent for the familiar and the unfamiliar stimuli.

These results are, in the present context, surprising, and go against the suggestion that expertise may modify susceptibility to task difficulties. One reason for this may lie in the way in which Scapinello and Yarmey defined familiarity. In their experiment, half the target stimuli were shown seven times (for 5 seconds each) during the inspection phase, and these were then categorised as 'familiar'. The remaining target stimuli were shown only once (for 5 seconds) during inspection, and were hence designated as 'unfamiliar'. Whilst providing an experimentally elegant design with no confound in terms of 'degree' of familiarity, or level of prior knowledge for each face, Scapinello and Yarmey's results may have been quite different had they defined familiarity according to a more usual criterion of pre-experimental knowledge.

Yarmey (1971) addressed this issue by using public figures of the day (e.g., Richard Nixon, Frank Sinatra) in an old/new task in which half the targets were presented upright at test and half were presented inverted. As with the previous study, Yarmey's (1971) results confirmed that upright faces were recognised better than inverted ones and that familiar faces were recognised better than unfamiliar ones. Critically, Yarmey's result confirmed those of the earlier study: the effect of inversion was equivalent across the familiar and the unfamiliar stimuli. Unfortunately, performance in both Scapinello and Yarmey's (1970), and Yarmey's (1971) studies appeared to be at ceiling for the recognition of upright familiar stimuli and this would have limited the potential for any interaction to emerge.

More recently, Collishaw and Hole (2000) have re-examined the issue of whether image manipulation affects stimuli to an equivalent degree regardless of their level of familiarity. They used familiar and unfamiliar images that had been either inverted, scrambled, or blurred, or both inverted and blurred, inverted and scrambled, or blurred and scrambled. Performance in each of these conditions was compared to performance with unmodified images. The rationale for the study was to examine the importance of configural and featural processing across familiarity. The use of inversion and scrambling affected configural processing while the use of blurring affected featural processing. The expectation was that impairing one or other method of processing would reduce the ability to recognise faces while impairing both methods of processing would reduce performance to chance levels. The results wholly support the expectations. What is more important is that the results were consistent across familiarity. Consequently, this study supports the previous findings in that there was no difference in the processing of familiar and unfamiliar faces when viewing was disrupted. Some consideration, however, has to be given to the task demands in Collishaw and Hole's paper. For familiar stimuli, a speeded familiarity task was used, while for the unfamiliar stimuli, an old/new recognition task was used. Whilst at pains to use an equivalent, recognition-based task for both familiar and unfamiliar stimuli, it is possible that the activation of long-term mental representations in the first case and short-term representations in the second case make the results of the task incomparable. With this in mind, the results of Collishaw and Hole's study are accepted with a degree of caution at this stage, until they can be replicated using a design which uses the same task for familiar and unfamiliar stimuli alike.

Much more recently, Tong and Nakayama (1999) have returned to the issue of whether facial familiarity can modify the processing of images presented under various unusual and difficult conditions. Tong and Nakayama's concern was with the formation of *robust representations* and they suggested that extensive visual experience or familiarity with individual exemplars, which necessarily incorporates a variety of viewpoints, lighting conditions, expressions, and poses, will lead to the formation of a robust representation. The result is quicker processing, even under difficult viewing conditions. This was borne out experimentally by the fact that highly overlearned faces (such as the participant's own face) were processed with greater speed when presented across changing viewpoints or orientations, compared to control stimuli consisting of previously unfamiliar faces that participants had been familiarised with.

The Present Study

Taking all these studies together, it is clear that inversion can be used as a way of determining higher level processing differences both across and within stimulus classes. It therefore appears to be a valid tool to investigate any processing differences across familiar and unfamiliar faces. The question is whether stimulus familiarity lends the perceiver any defence against difficult viewing conditions brought about by stimulus inversion. On this point, however, the literature is mixed. In terms of an expertise account, Rhodes *et al.* (1989) suggest that own-race (familiar) faces are *more* affected by inversion than other-race (unfamiliar) faces. However, in terms of a representational account, Tong and Nakayama (1999) suggest that familiarity allows the formation of robust representations that allow the perceiver to *withstand* the effects of difficult viewing conditions. Given this controversy, the present study is designed to address the issue further. Facial stimuli are used that are either highly familiar (personally known) or unfamiliar to the participant. A speeded familiarity task is used with stimuli presented upright or inverted. In this way, it is possible to examine the question of whether inversion has a different effect on a recognition-level decision across stimulus familiarity.

Method

Design

A within-participants 2 x 2 full factorial design was used to examine the effects of orientation and familiarity on face processing. Personally familiar and unfamiliar faces were presented sequentially in a speeded familiarity task. The independent variables were familiarity (familiar, unfamiliar) and orientation (upright, inverted). Performance was measured using accuracy (error rate), indicators of sensitivity (A') and bias (B''_D), and speed of response.

Participants

Forty-four undergraduate and postgraduate volunteers participated in this study. Twenty-two were from the Department of Psychology, and the remainder were from the Department of Archaeology (located on a different campus). Participants' ages ranged from 18 to 44 years (m = 23.5 years, sd = 8.3 years). There was a gender bias towards women, with 16 women and 6 men taking part from the Psychology Department, and 13 women and 9 men taking part from Archaeology. All had normal, or corrected-to-normal, vision.

Materials

Forty-four colour photographs of staff and postgraduate students, who were known to the undergraduate students in their respective departments, were taken using an Apple digital camera. There were an equal number of male and female faces in each group and all were Caucasian. Photographs were taken full-face, and individuals were asked to adopt a neutral facial expression "as if waiting for a bus". Twenty faces from each department were selected as stimuli (see Appendix A), having been matched for equivalence of age, gender, distinctive features (facial hair and spectacles) and colouring. Faces were scaled on the basis of inter-pupil distance using Corel PhotoPaint, and were presented against a uniform white background. The faces were then converted to greyscale bitmaps to remove any slight variation in skin coloration. Finally, both upright and inverted versions of each face were created. Examples are shown in Figure 2.1.

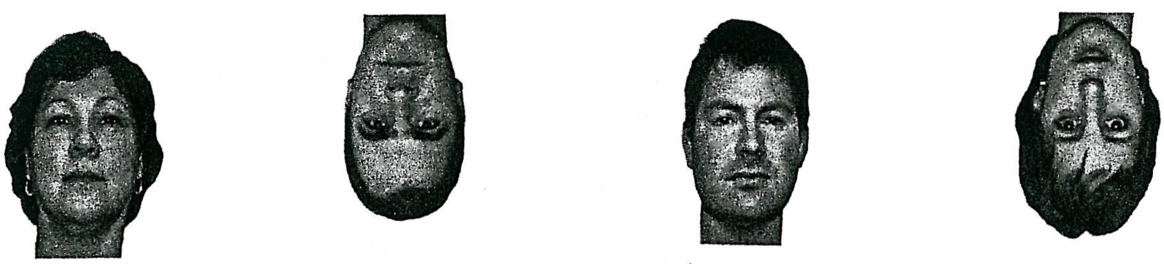

Figure 2.1: Examples of faces used from the Psychology Department (two on left) and the Archaeology Department (two on right).

The stimuli were presented on a Dell Optiplex GL 575 computer, using a 15" 16-bit colour monitor and a screen resolution of 800 x 600 pixels. Each image measured approximately 6 x 8 cm (28.34 pixels/cm) and subtended a visual angle of 5.4°. A pilot study was conducted with 4 participants from each department to confirm that the faces from the participants' own departments were recognisable. These pilot participants were drawn from a comparable sample but did not take part in the main experiment.

Procedure

Participants were seated 60 cm from the computer monitor. The stimuli were arranged so that participants saw upright and inverted familiar and unfamiliar faces in a random order. Each face was seen only once and the participant was required to provide a familiarity judgement, as quickly but as accurately as possible. The image remained in view until a response had been made, and participants gave their response by using the thumbs to press either a button marked 'familiar' or 'unfamiliar' on the two-button response pad. A 500 msec blank screen separated the response and the next trial, and allowed the participant time to prepare for the next presentation. Four practice trials were provided at the start of the test phase, using two familiar and two unfamiliar faces. This enabled participants to orient themselves to the task. The forty-four experimental trials proceeded without break after these practice trials. Speed (from presentation of image, to button-press) and accuracy of response were recorded as dependent variables.

At the end of the task, participants were electronically presented with each of the test faces. Each was seen upright and participants provided a familiarity rating using a scale from 1 (never seen this person before) to 7 (I know this person very well and can name them easily). This was done to establish the actual level of familiarity of each of the target faces for each of the participants. Responses to this task were made by using the computer mouse to click on the appropriate point of the on-screen rating scale. The total time for the task and familiarity check was approximately 10 minutes, and participants reported no signs of fatigue.

Results

The design of the present study allowed several analyses to be conducted. First, error rate and speed of response for target stimuli were examined to see whether the size of any inversion effect was consistent across the familiar and the unfamiliar stimuli. Following this, signal detection theory was applied to the accuracy data in order to generate measures of sensitivity (A') and bias (B''_D). These measures were also examined to ascertain the effect of inversion on participants' decision-making processes. For each measure, there was no evidence of effects of order, gender, or participant department. Hence, all data were pooled across these variables. The analyses are taken in turn below.

Accuracy

The results from four participants from the Archaeology Department were excluded from the analysis as their post-experimental rating task indicated that they rated only 2 out of the 20 'familiar' faces as sufficiently familiar[1] for inclusion. Data from 40 participants remained (22 from Psychology, 18 from Archaeology). For the purposes of analysis, these data were collapsed across departments.

Table 2.1 *Mean proportion of errors (with standard deviation) on a speeded familiarity decision for upright and inverted familiar and unfamiliar faces.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.03	(.07)	.29	(.19)	40
Unfamiliar	.06	(.14)	.12	(.16)	40

The accuracy data, presented as proportion of errors, are summarised across experimental condition in Table 2.1 above. From this it appears that recognition of upright faces was uniformly good. However, inverted face recognition was much more error prone, especially when faces were familiar. Analysis by means of a

[1] Where faces were rated at 3 or less on the 7 - point rating scale (3 = This person is familiar but I cannot place them).

repeated measures analysis of variance (ANOVA) on the proportion of errors showed main effects of familiarity ($F(1, 39) = 5.08, p < .05$) and orientation ($F(1, 39) = 63.87, p < .001$), plus a significant interaction ($F(1, 39) = 38.27, p < .001$). *Post hoc* analysis of the simple main effects revealed that this interaction was driven by a significant difference between familiar and unfamiliar faces when inverted ($F(1, 39) = 15.06, p < .001$) but not when upright ($F(1, 39) = .86, p > .05$). Consequently, the prediction that inversion might affect familiar and unfamiliar faces differently was upheld.

Response Time for Correct Decisions

The mean response time (RT) for correct decisions was calculated for each participant within each experimental condition, having removed outliers greater than or less than the mean + 2.5 standard deviations. The resultant data are summarised in Table 2.2 below. This suggests that familiar faces were recognised faster than unfamiliar ones, and upright faces were recognised substantially faster than inverted ones. There is, however, little indication that inversion had a different effect across familiarity.

A repeated measures ANOVA confirmed the main effect of familiarity ($F(1, 39) = 16.41, p < .001$) and orientation ($F(1, 39) = 40.97, p < .001$). No significant interaction emerged to qualify these effects ($F(1, 39) = .032, p > .05$) suggesting that, although inversion had a slightly greater effect on correct responses to familiar faces than unfamiliar ones, this did not reach significance.

Table 2.2: *Mean response times (and standard deviations) for correct classifications as 'familiar' and 'unfamiliar' across upright and inverted presentation.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	1225	(291)	1544	(502)	40
Unfamiliar	1452	(362)	1732	(525)	40

The possibility of a speed-accuracy trade-off cannot, however, be discounted. The data indicated that a high error rate typically occurred in conditions where participants were responding quickly. This is especially so in the case of inverted stimuli, where significantly more errors were made to familiar stimuli than to unfamiliar ones even though participants responded with greater speed ($t(39) = 2.63$, $p < .05$).

The response time data were analysed a second way in order to examine the potential for any interaction fully. This time, response time for correct decisions was calculated for each level of familiarity as signalled by the participant on the post-experimental familiarity rating task. Consequently, instead of the stimuli being partitioned according to a-priori designations as merely familiar or unfamiliar, the stimuli were partitioned into finer categories of familiarity according to subjective and individual ratings. Data from all 44 participants were considered, but data from one participant was excluded due to their recognition of only 1 out of 20 faces from the (nominally) familiar category. The data from the remaining 43 participants were analysed. The inclusion of data from three participants who were excluded from the rest of the analysis was possible, as they had indicated at least a low level of familiarity, which could be regressed against their response times. With these data, response times were then regressed against familiarity ratings in order to examine the gradient for upright and inverted presentations. However, a paired samples t-test showed no significant difference between the two gradients ($t(42) = 1.10$, $p > .05$). This analysis again suggests that familiarity did not interact with inversion to affect the speed with which faces were classified as familiar/unfamiliar.

Sensitivity

Thus far, only the accuracy data have given any indication that familiarity modifies the effect of inversion on face perception. There are, however, logical difficulties in the interpretation of the accuracy data, which arise from consideration of the task. A speeded familiarity task requires a 'familiar' decision to known faces and an 'unfamiliar' decision to new faces. This difference in required response makes comparison of the accuracy data across familiarity difficult, as it is possible that the 'familiar' decision invokes a different process from that required in rejecting a face as 'unfamiliar'. Similarly, interpretation of the response time data is also problematic. The same process differences underlying the accuracy data may preclude clear interpretation of the response time data and, in addition, a speed-accuracy trade-off leads us to accept these data with caution.

Consequently, the data are subjected to analysis by application of signal detection theory. A measure of sensitivity will allow examination of correct and incorrect decisions within each experimental condition, and a measure of bias will record whether participants showed a tendency to respond one way or another regardless of the stimulus type.

The non-parametric sensitivity measure A' was used here. It is calculated from hit rates and false alarms and varies between 0 and 1. A value of 1 indicates perfect performance, and a value of 0.5 indicates chance level performance (see Appendix H for sensitivity and bias formulae). A' scores for upright faces (M = .97, sd = .04) were substantially greater than those for inverted faces (M = .88, sd = .07), as verified by a paired samples t-test ($t(39) = 7.41, p < .001$). Consequently, discrimination of familiar from unfamiliar stimuli is better when faces are presented upright than when inverted. It is, unfortunately, impossible to break this finding down further to examine performance for each of the four cells of the design because of the confound between manipulated and measured variables.

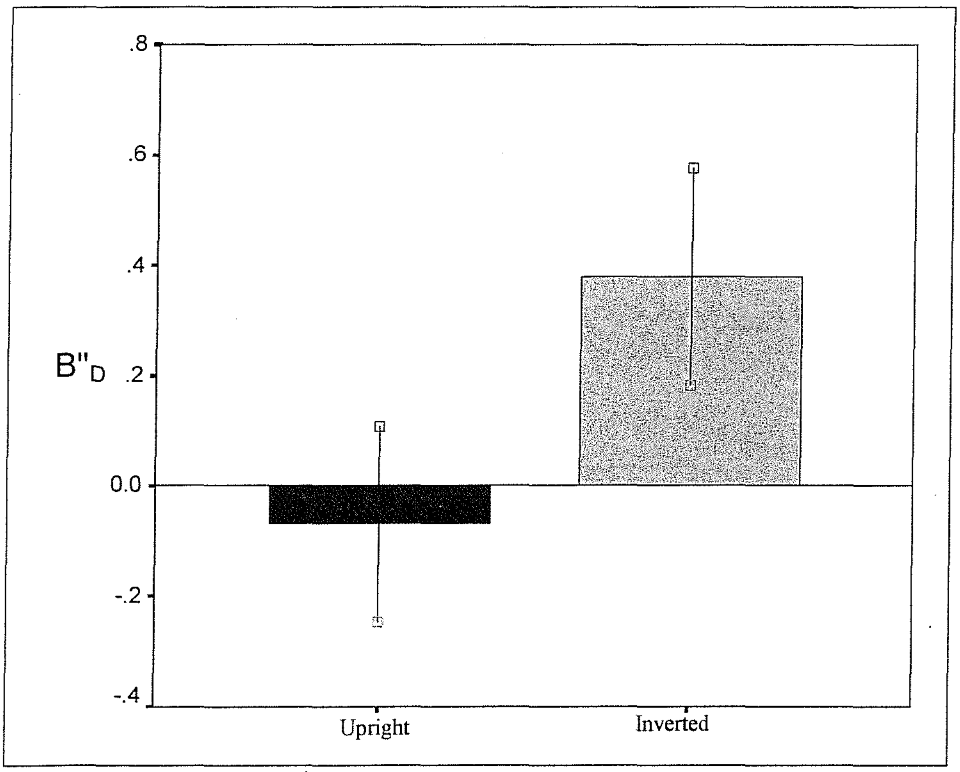

Figure 2.2: Mean B''_D values (with standard error bars) showing the effect of inversion on the bias to respond 'familiar' or 'unfamiliar'. Upright faces were subject to little bias but inverted faces were more likely to be called 'unfamiliar'.

Bias

B''_D is a measure of bias, and is calculated from misses and correct rejections. Values range between –1 and +1, with negative values indicated a bias to say 'familiar', positive values indicating a bias to say 'unfamiliar', and a value of 0 indicating no response bias. Examination of the B''_D values indicates that participants were acting cautiously for inverted faces (M = .38, sd = .62) with responses biased towards saying 'unfamiliar' (see Figure 2.2). No bias was evident for the upright faces (M = -.07, sd = .56). Analysis by means of a repeated measures t-test confirmed a significant difference in bias across orientation ($t(39) = 5.01, p < .001$).

Summary of Results

Analysis of response times showed no differences in performance between familiar and unfamiliar faces, however, the analysis of error rates suggested that inversion was more disruptive to recognition of familiar than unfamiliar faces, supporting the expertise hypothesis. The impact of this finding is diminished by the evidence of a speed-accuracy trade-off and biased responding, suggesting that the disadvantage shown in the recognition of familiar faces should be treated cautiously.

Discussion

The results of the present experiment provide no evidence to suggest that familiarity influenced the effect of inversion on face recognition. Participant accuracy was analysed first in terms of error rate. The data were suggestive of a greater inversion effect for the familiar faces than for the unfamiliar ones. However, these data are in some sense not comparable as different responses would be required for the two sets of stimuli in order for the participant to be correct. Put another way, the familiar faces beg a 'familiar' decision while the unfamiliar faces beg an 'unfamiliar' decision in order for the participant to be correct. In this light, it is highly possible that an 'unfamiliar' decision is taken only when a previously stored (familiar) mental representation cannot be located in memory. As such, the familiar and unfamiliar decisions are 'yoked' and the unfamiliar decisions will always be made after any familiar decisions. As can be seen in Table 2.2, the RT data support this view. Indeed, formal analysis confirmed that participants were quicker when faces were (i) familiar, and (ii) upright.

In examination of the hypothesis, the RT data also gave no indication that familiarity affected the degree of inversion effect. Both through the ANOVA and through the regression analysis, the data suggested an interaction of familiarity and inversion, however, this did not reach significance in either analysis. The response time data should, however, be treated with some caution as a speed-accuracy trade-off was apparent in the data. Participants appeared to respond relatively quickly but make a large number of errors particularly when presented with familiar inverted stimuli.

Given the methodological problem with the interpretation of the accuracy and RT data as they stood, signal detection theory was applied to the accuracy data in order to yield information on the sensitivity and bias in participant responding. Analysis of the sensitivity measure suggested that participants were better able to discriminate familiar from unfamiliar faces when stimuli were upright than when inverted. This is somewhat unsurprising and confirms the well-documented finding that face processing performance declines with inverted presentation.

The sensitivity data are complemented by the fact that the bias measure suggested a cautious approach when stimuli were inverted. Participants revealed a strong bias to respond 'unfamiliar' no matter what the status of the inverted face. This is a potentially interesting result as it suggests a change in the meta-cognitive strategy used when processing inverted faces. Instead of searching memory to verify existence of a match to a mental representation, as might be done with upright faces, the irregular viewing conditions have the result of making all stimuli look unfamiliar even if they belong to highly familiar individuals. Thus, when stimuli are inverted, it is possible that the starting point for decision making is altered such that individuals begin the processing task with an 'unfamiliar' expectation. Alternatively, given the unusual viewing conditions, it is possible that the level of confidence required for a participant to respond 'familiar' is higher when faces are inverted than when upright, resulting in few familiar responses being made. Both of these explanations point to something different happening when processing inverted faces and the likelihood is that both explanations would lead to a differential effect on the processing of truly familiar over truly unfamiliar inverted stimuli. It is unfortunate that the present experiment does not allow an adequate test of that prediction. The purpose of the following experiment is, however, to carry these themes forward using an alternate methodological procedure.

Despite the lack of a reliable significant interaction to indicate that facial familiarity influenced processing at the recognition level, several aspects of the data are worthy of note. For example, it is interesting to note that recognition of inverted familiar faces, although poor, is still above a chance level of performance. This suggests that a sense of familiarity is not completely lost when familiar faces are inverted. Consequently, although configural processing may be impaired by inversion, sufficient identity-specific information may be obtained through slow and effortful configural processing made possible by mental rotation, or through featural processing, to indicate the identity of the face.

This line of reasoning echoes that expressed by Collishaw and Hole (2000) who suggested that while both configural and featural processing together form the ideal processing conditions, either configural or featural processing in isolation will still be sufficient to allow recognition to be achieved. Bartlett and Searcy (1993) and Searcy and Bartlett (1996) echo the notion of dual routes to identification. On the basis of empirical work, it was suggested that one factor that may set apart the processing of familiar from unfamiliar stimuli is that familiarity provides the perceiver with recourse to a number of processing strategies. When inversion interferes with the use of one strategy, the perceiver can effortlessly switch to another in order to complete the task. Hence, when presented with inverted yet familiar faces here, the perceiver may have been able to achieve recognition through recourse to featural as well as configural information that reliably signalled identity.

Alternatively, above chance performance with inverted familiar faces may be the result, not of featural processing, but of configural processing made possible by mental rotation of the stimulus. When one is used to seeing an upright face, inversion makes it appear peculiar. Rather than trying to process the inverted face, one tries to rotate it to an upright position before the recognition process can begin. This has the effect of increasing the response times as shown here, but provides a strategy by which recognition can still be achieved for familiar inverted faces (see Rakover & Teucher, 1997).

These notions are not unlike that of *robust representations* presented by Tong and Nakayama (1999) earlier. However, Tong and Nakayama's belief was that the robust representations held for familiar stimuli would result in greater recognition of the familiar over unfamiliar stimuli when inverted. The data presented here do not support this strong prediction.

Conclusion

On the face of it, the present results appear to offer support for Rhodes *et al.*'s (1989) view that faces processed in an expert manner would be more susceptible to inversion. As such, the suggestion presented by Tong and Nakayama (1999), that facial familiarity might generate representations that can withstand difficult viewing conditions, is cast into doubt. At this stage, however, these conclusions should be accepted with a degree of caution as the design of the present study was not optimal in terms of allowing for a valid exploration of the hypotheses. Given this, the following experiment will re-address the hypothesis underlying the present study by using an old/new task in which measures and manipulations are not confounded and responses are independent of one another.

CHAPTER 3

The Influence of Familiarity on the Inversion Effect Within an Old/New Recognition Paradigm.

Abstract

The previous study examined the issue of whether familiar and unfamiliar facial stimuli were equally affected by inversion. However, the paradigm used in the previous study did not provide an ideal test of the prediction. The purpose of the present study is to re-examine the effect of familiarity on inversion by using an old/new recognition task. Personally familiar and unfamiliar (once seen) faces were used, and performance was measured using accuracy rates, measures of sensitivity and bias (A' and B''_D), and speed of response. Across measures of accuracy, response time, and sensitivity of responding, results suggested that inversion had a greater effect on the recognition of unfamiliar faces than on the recognition of familiar ones. In other words, recognition was preserved when faces were inverted but familiar. These results are discussed with reference to the notion of robust representations, and to the idea of automaticity when processing facial familiarity, even when this level of processing is not required by the task at hand.

The Influence of Familiarity on the Inversion Effect Within an Old/New Recognition Paradigm.

The rationale for the present study is the same as for the previous study. Again, the focus is on whether familiarity influences the degree to which face processing is disrupted by inversion. According to the most dominant theoretical models of face processing such as Bruce and Young's (1986) Information Processing Model, and subsequently, Burton, Bruce and Johnston's (1990) Interactive Activation and Competition Model of Face Processing, familiarity is something that is computed as a consequence of finding a match to a mental representation. It is difficult to see from these sorts of models how familiarity *could* influence the effect of a disruption to the processing of a face. This view is supported by empirical work (Collishaw & Hole, 2000; Scapinello & Yarmey, 1970, Yarmey, 1971). However, the issue has been raised again at a theoretical level by the work of Tong and Nakayama (1999) through their concept of *robust representations*, and in the body of evidence suggestive of a qualitative difference in the processing strategies used for familiar (expert) and unfamiliar faces. Given the controversy over the existing results, and the obvious implications for the theoretical literature in this area, the issue is examined again here. The focus within the current chapter, however, is on the use of a more appropriate methodology in order to enable the predictions to be addressed adequately.

The change in methodology is an important consideration in and of itself. Different tasks are likely to elicit different strategies, both in terms of processing and in terms of decision-making. Indeed, task effects, along with differences in stimuli, are a likely source of inconsistencies between reported studies (Archer, Hay & Young, 1992; Valentine, 1988). Both Payne, Bettman and Johnson (1992) and Bartlett and Searcy (1993) comment on the probability that decision processes are likely to vary across tasks, and that task complexity alone can affect decision strategy. For instance, in a task where the identification of the face is required, participants may set their criterion for response at a more conservative level and respond positively only when their confidence is high. This was apparent in the previous study and was revealed through the identification of a marked cautious response bias when faces were presented upside-down. In contrast, a different task, which can be performed without the need to identify a face, may attract the use of a less stringent response criterion. As a result, a different pattern of responding may emerge.

With this in mind, it is pertinent to re-examine the published literature and to focus on the methodologies used. Significant effects of familiarity on face processing have been revealed in several studies to date. For example, Baudouin *et al.* (2000) demonstrated an advantage for familiar faces when the task involved making judgements of expression under difficult viewing conditions. The presentation of familiar and unfamiliar faces was impaired either through masking of the mouth area, or through reducing the exposure duration from 400ms to 15ms. Baudouin *et al.*'s results suggested that expression judgements were faster for familiar faces under both impaired conditions (as compared to the respective control conditions). Consequently, it was suggested that facial familiarity might be processed automatically, even when not required by the task at hand, and that this familiarity may come to play a role when the primary task is handicapped in some way.

A further study to report an advantage for familiar over unfamiliar faces was presented by Roberts and Bruce (1988). They conducted two tasks, a recognition task and a gender decision task, to familiar and unfamiliar faces. The results of the two tasks are quite different and underline the importance of task considerations when examining familiarity effects. For the recognition task, a main effect of facial familiarity emerged with participants being faster to identify familiar faces as 'known' than to identify unfamiliar faces as 'unknown'. In itself, this result is unsurprising, but what is important is that when the task was made difficult through feature masking, processing of the familiar and unfamiliar faces was affected to an equal degree. In other words, whilst performance in the task was impaired by feature masking, the decrement in performance was consistent across the familiar and unfamiliar faces. These results mirror those of the study reported in the previous chapter and suggest that, when the task involves explicit identification, familiarity does not modify the impact of an impairment to facial processing.

Consideration of Roberts and Bruce's (1988) gender decision task, however, provides a slightly different picture. Again, participants were required to process familiar and unfamiliar faces under conditions where the whole face was visible, or under conditions of feature masking. Here, the results suggested no main effect of facial familiarity on speed of response. However, when the task was made especially difficult, by masking the nose region, then knowing the face helped in making the gender decision. As with Baudouin *et al.*'s (2000) study, these results suggest that the processing of familiarity can aid decision making in a task that was ostensibly identity-unrelated.

Finally, Young Hay, McWeeny, Flude and Ellis (1985) revealed qualitative differences in processing across familiar and unfamiliar faces using a feature-matching task. In their first experiment, participants were simultaneously shown two photographs depicting (i) a complete face, and (ii) either the internal or the external features of a face. The task was to decide whether the two photographs were of the same person or not. The results suggested faster matching of the whole face with the internal features when the faces were familiar than when unfamiliar. This pattern of results was not apparent when matching whole faces with external features, and nor was it apparent when it was possible to match the photographs on picture-related factors rather than on face-related factors (Experiment Two). However, when treated as a face, and when using facial parts from which a configuration may be extracted, the processing of familiarity again aided performance in a task where the processing of identity was not explicitly required.

Set against these results, three experiments have been reported which fail to find any qualitative difference in the processing of familiar and unfamiliar faces. First, Collishaw and Hole (2000) reported no differences across familiarity using an identification-based task. Across a range of manipulations designed to provide impairment (inversion, scrambling, blurring, and the three pair-wise combinations of these), the pattern of performance decrement, relative to a baseline condition, was consistent both when recognition of highly familiar celebrities was considered and when recognition of a once-seen face was considered. This suggests no qualitative difference in the processing of familiar and unfamiliar (once seen) faces. However, some consideration of the task demands for each stimulus set is warranted. Familiar faces were tested using a familiar/unfamiliar decision to target and distractor stimuli. In contrast, unfamiliar faces were tested using an old/new paradigm requiring 'old' responses to previously studied faces and 'new' responses to distractor faces at test. Consequently, while the authors were at pains to equalise task demands across the two sets of stimuli by using a recognition task for each set, the slightly different tasks used may have had some effect on the comparability of the results across familiarity.

Second, the results of Scapinello and Yarmey (1970) and of Yarmey (1971) both show no difference across familiarity in the extent of the inversion effect. Both studies used an old/new recognition paradigm, and so avoided the potential problem of different tasks for the different stimulus sets. However, both studies revealed potential methodological concerns that make the results questionable. For instance, Scapinello and Yarmey (1970) equated frequency of exposure at a study phase to 'familiarity'. The lack of any effect of familiarity in this case may simply reflect the fact that familiarity could be considered to be something other than, or more than, mere exposure. In Yarmey's (1971) experiment, familiarity was manipulated in a more traditional way by contrasting performance with unfamiliar faces against performance with the faces of highly familiar public figures. However, the use of an old/new paradigm with these stimuli created a fairly easy task in which a ceiling effect prevented firm acceptance of the results. In some sense, although the old/new task appears to be ideal as a way of determining familiarity effects since the task involves no explicit need to process identity, the existing literature using old/new tasks has not allowed the issue to be resolved. The current study returns to the use of an old/new recognition paradigm but takes steps to minimise the potential for a ceiling effect.

Method

Design

An old/new recognition task was used to examine the effects of inversion and familiarity on face recognition. A 2 x 2 within-subjects design was used, in which the independent variables were stimulus orientation (upright, inverted) and stimulus familiarity (familiar, unfamiliar). It should be noted here that although the unfamiliar face set are unknown to participants prior to this experiment, the old/new recognition paradigm requires an inspection phase to precede the test phase. Therefore the unfamiliar target stimuli are more correctly known as 'once-seen.'

Performance on the recognition task was measured using accuracy, speed of correct responses, a measure of sensitivity (A') and bias (B''_D).

Participants

Twenty-five volunteers (18 females, 7 males) participated in this experiment. All were undergraduate or postgraduate students from the Department of Psychology and had normal, or corrected-to-normal, vision. Participant age ranged from 19 to 40 years (mean = 26.9 years, sd = 7.8). In order to facilitate counterbalancing of the upright and inverted stimuli, the participants were randomly assigned to one of two groups, and groups were matched as far as possible for age and gender of participants. The first group saw half of the targets and distractors upright at test while the second group saw those targets and distractors inverted at test.

Materials

A total of 72 full-face photographs were used as stimuli within the current experiment. Half of these faces were personally familiar and were identical to the 'familiar' set used in Experiment One. These were staff and postgraduate students from the Department of Psychology (Appendix B) and were rated as very well-known to the psychology students taking part in the present study. The remaining 36 stimuli were unfamiliar to the participants and comprised staff and postgraduates from another department, and from the Stirling PICS database (Appendix C).

Within the familiar and unfamiliar stimulus sets, an equal number of male and female faces were used, and all were of Caucasian origin. Unfamiliar faces were selected to be matched to the familiar faces on the basis of hair colour, hair length, age, and presence of any distinguishing features (such as facial hair, spectacles, etc.). All faces were extracted from their backgrounds and mounted on a uniform white background. They were then scaled, using Corel PhotoPaint such that all images were matched according to inter-pupil distance. Converting the faces to 8-bit greyscale images ensured the removal of any slight variations in colouring.

The familiar and unfamiliar face sets were each divided into three subsets to provide 12 targets, 12 distractors, and 12 filler faces. The targets were viewed at both study and test. The distractors were seen at test only, and the filler faces were seen at study only. Finally, two versions of each of the target and distractor faces was generated to create one upright image and one inverted image. This resulted in 120 images in total (12 upright targets, 12 inverted targets, 12 upright distractors, 12 inverted distractors, and 12 fillers for each of the familiar and unfamiliar face sets). In addition to these experimental stimuli, a number of images were prepared to serve as stimuli in the practise trials prior to the test phase. These practise images comprised familiar and unfamiliar upright and inverted faces drawn from the same populations, and prepared in the same way as the main stimuli. None of these practise images were used in the main experiment.

The stimuli were presented on a Dell Optiplex GL 575 computer, using a 15" 16-bit colour monitor and a screen resolution of 800 x 600 pixels. Participants sat approximately 50 cm from the computer monitor and images measured approximately 6 x 8 cm (28.34 pixels/cm).

Procedure

The experiment consisted of two phases: a study phase, and a test phase. During the study phase, participants were shown 48 upright faces (24 familiar, 24 unfamiliar: half targets, half fillers). The faces were presented one at a time in a fixed random order and at a rate of 3 s per image. Following image offset, participants were asked to indicate the age of the person depicted. Responses were made using a 5-point scale with the points corresponded to the following age-brackets: 18-25 years; 26-35 years; 36-45 years; 46-55 years; 56-65 years. Participants used the computer mouse to click on the point of the scale which corresponded with their perception of age.

Immediately following the study phase, a test phase was presented. Participants viewed an instruction screen which allowed a self-paced break prior to beginning the test phase. The procedure at this phase allowed the provision of a number of practice trials in which faces were presented one at a time. Half of these practise images were familiar and half were unfamiliar, and faces in each familiarity level were presented either upright or upside down. Participants were warned that the orientation of the face may have changed but in every case, regardless of orientation, the task was to determine whether the face had been seen in the previous phase ('old') or had not ('new'). Participants made their response by pressing the appropriate button on a response pad. Feedback was provided during this practise phase, and the main test phase only proceeded when the participants had reached a criterion of 5 correct responses in a row.

During the main test phase, participants were shown the 24 target faces seen previously (12 familiar, 12 once seen) together with 24 distractor faces (12 familiar, 12 unfamiliar). As before, half the targets and distractors in each familiarity level were presented upright and half upside down. Again, and regardless of orientation, the task was to determine whether the face had been seen in the previous phase ('old') or had not ('new'). No feedback was provided during the main test phase.

Items were presented one at a time, in a fixed random order, and remained on screen until the participant responded with a button press. A 500 msec blank screen separated the response and the next trial, and allowed time for the participant to rest in preparation for the next stimulus. Care was taken to counterbalance the position of the old/new response buttons across participants. In addition, the orientation at test of each face was counterbalanced across the two participant groups. In this way, any results could not be attributed to a bias to press the left (or right) response button, or to a specific set of stimulus faces. Finally, care was taken to ensure that the first two and last two stimuli shown in the test phase were additional distractor stimuli, with one being shown upright and one inverted at both the start and the end of the testing sequence. These were dropped from all subsequent analyses, in recognition of possible practise and fatigue effects.

Performance was measured by recording accuracy of response, together with speed of response from image onset to button-press. The study and test phases together took a maximum of 15 minutes to complete, and participants did not report any fatigue.

Following completion of the task, participants were shown each of the familiar and distractor faces and were asked to name each one. This post-experimental check was conducted to ensure that the familiar stimuli were indeed familiar and that the unfamiliar stimuli were not known by the participants. Mean post-experimental identification of the 24 familiar faces across the 25 participants was 90.6% (sd = 10.8). Following this, participants were thanked and debriefed.

Results

Of the 25 participants, one participant failed to recognise more than 50% of the familiar faces. Data from this participant were consequently dropped from all subsequent analyses. Accuracy and speed of correct responses were recorded, and the data are analysed to examine these measures together with measures of sensitivity of response (A') and response bias (B''_D). Only target data were analysed, and response times which were greater than 2.5 sd + mean were excluded from the analysis.

Accuracy

The accuracy of old/new recognition (as recorded through proportion of errors) for upright and inverted familiar and unfamiliar (once seen) targets is shown in

Table 3.1 below. From this table it appears that familiar faces produced fewer errors than unfamiliar faces, and upright faces produced fewer errors than inverted faces. In addition, it is possible that the effect of inversion was smaller when faces were familiar than when unfamiliar. However, the effect of inversion on familiar faces led to a three-fold increase in error rate whereas inversion reduced accuracy on unfamiliar faces by two-and-a-half times. Examination of these results through a repeated measures ANOVA revealed a significant effect of familiarity ($F(1, 23) = 40.53, p < .001$), and a significant effect of orientation ($F(1, 23) = 31.72, p < .001$). Both main effects were, as predicted, modified by the emergence of a significant interaction ($F(1, 23) = 9.08, p < .01$). Examination by means of two repeated measures t-tests, with a Bonferroni correction for multiple tests, showed that there was a significant effect of inversion both when faces were familiar ($F(1, 23) = 15.79, p < .001$) and when faces were unfamiliar ($F(1, 23) = 32.76, p < .001$). A t-test on the difference between upright and inverted performance showed that the effect of inversion, whilst significant for both levels of familiarity, was far greater when faces were unfamiliar ($t(23) = 3.01, p < .01$).

Table 3.1: *Mean proportion of errors (and standard deviations) for familiar and unfamiliar faces when upright and inverted.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.04	(.09)	.13	(.15)	24
Unfamiliar	.15	(.13)	.37	(.18)	24

Response Time for Correct Decisions

The response times for correct decisions to target stimuli are summarised in Table 3.2 below. These follow the same pattern as the error rate data and, as such, there is no evidence of any speed-accuracy trade-off in the current experiment. In examination of the data, it appeared that the familiar faces were recognised as old more quickly than the once seen faces, and similarly, the upright faces were recognised as old more quickly than the inverted ones.

A repeated measured ANOVA on the data revealed a significant main effect of familiarity ($F(1, 23) = 27.81$, $p < .001$) and orientation ($F(1, 23) = 22.00$, $p < .001$). The interaction between these two variables tended to significance ($F(1, 23) = 3.29$, $p = .083$) with the suggestion that inversion had a greater effect when faces were unfamiliar than when familiar. These data thus mirror the pattern obtained with error rates.

Table 3.2. *Response times (msec) and standard deviations for familiar and unfamiliar upright and inverted faces.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	1036	(210)	1250	(273)	24
Unfamiliar	1328	(346)	1659	(527)	24

Sensitivity

Sensitivity was calculated using the non-parametric A'. This measure was chosen due to the small number of trials contributing to each experimental condition. The value of A' indicated the ability to discriminate old from new, with values varying between 0 and +1 (perfect discrimination). The mean A' values for each of the four conditions are summarised in Table 3.3 below. From this, the pattern suggested by the error data is confirmed. Discrimination is better when stimuli are familiar and when they are upright, and the data suggest a tendency for inversion to affect discrimination to a greater degree when faces are unfamiliar.

Analysis by means of a repeated measures ANOVA confirms the main effect of familiarity ($F(1, 23) = 24.53, p < .001$) and orientation ($F(1, 23) = 30.75, p < .001$). These effects are, as predicted, qualified by a significant interaction ($F(1, 23) = 4.48$, $p < .05$). Again, the interaction was examined by means of repeated measures *t*-tests, with application of a Bonferroni correction for multiple tests. These confirmed that the effect of inversion on sensitivity was significant both when faces were familiar ($F(1, 23) = 7.78, p < .01$) and when faces were unfamiliar ($F(1, 23) = 26.65, p < .001$). Analysis of the difference scores between upright and inverted performance again confirmed that while inversion affected the sensitivity of performance whatever the familiarity of the face, the effect was greater when faces were unfamiliar ($t(23) = 2.10, p < .05$).

Table 3.3: *Sensitivity of response (A') and standard deviations for familiar and unfamiliar upright and inverted faces.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.94	(.05)	.88	(.10)	24
Unfamiliar	.88	(.10)	.76	(.10)	24

Bias

Finally, the accuracy data were manipulated to generate a measure of response bias (B''_D). This measure allows any bias to respond 'old' or 'new' across each of the four conditions to become apparent. Values vary between +1 and –1, with positive values indicating a bias to respond 'new' and negative values indicating a bias to respond 'old'. The data are summarised in Figure 3.1 below, from which it is apparent that there is a bias to say 'old' to familiar faces and 'new' to unfamiliar ones. This bias appears strongest when the stimuli are inverted.

Analysis by means of a repeated measures ANOVA confirmed a main effect of familiarity ($F(1, 23) = 12.99, p < .01$). However, no main effect of orientation ($F(1, 23) = .04, p > .05$) emerged and surprisingly, no interaction between the two variables was apparent ($F(1, 23) = .21, p > .05$). These results merely suggest a bias to say that familiar faces are 'old' even if they have not been seen in the experimental context previously and a bias to say that unfamiliar faces are 'new' irrespective of prior exposure.

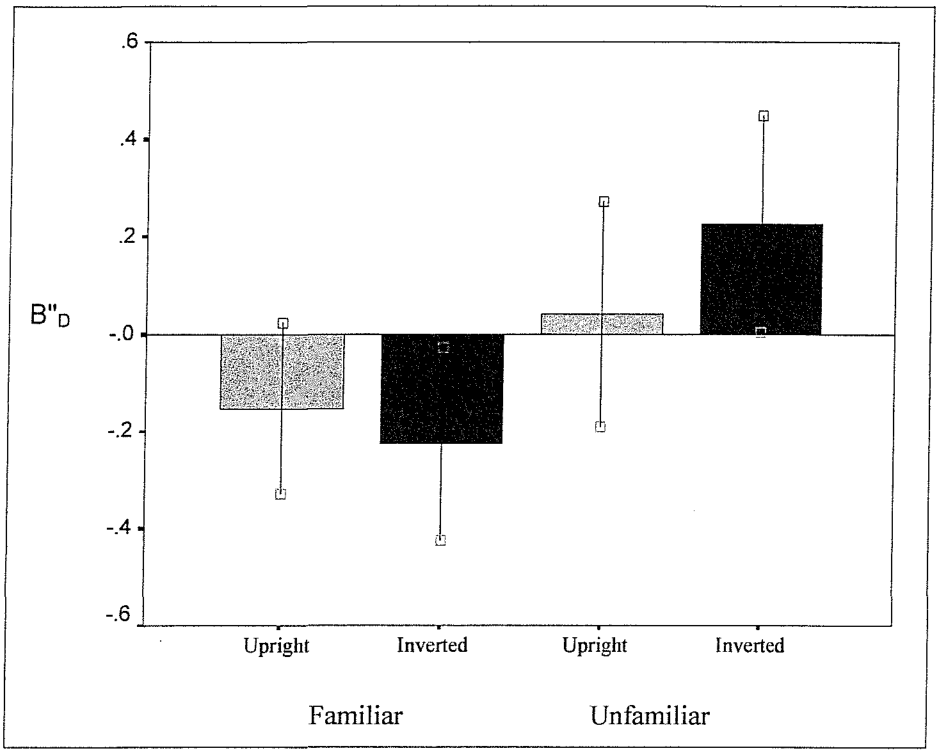

Figure 3.1: Mean B''_D values (bars represent standard error) across familiarity and orientation showing a propensity to respond 'old' to familiar faces and 'new' to unfamiliar faces.

Summary of results

The present experiment has shown that familiarity does influence the impact that inversion has on face processing. In an old/new recognition task, both accuracy and sensitivity measures showed a greater deleterious effect of inversion when faces were unfamiliar (once seen) than when familiar. This suggests that facial familiarity does provide the perceiver with some element of protection against difficult viewing conditions. This pattern of results was not echoed by the response time data and neither was it echoed by the pattern of response bias.

Discussion

The present results have echoed those of Experiment One in certain respects. First, familiar faces were processed with greater accuracy and speed than unfamiliar ones. This confirms that familiar items are processed with greater efficiency than unfamiliar items. Second, upright faces were processed with greater accuracy and speed than inverted ones. Again, this confirms the negative influence that inversion has on the processing of faces, and these results support the wealth of literature on inversion effects.

In contrast to the previous experiment, however, the results of the present experiment provide evidence that familiarity can reliably influence the magnitude of the inversion effect. Indeed, two things are notable about the results of this experiment: First, the results show consistency across accuracy, speed of correct decisions, and sensitivity of responding. Second, the pattern of results suggests that the more familiar the face, the less significant the effect of inversion. In this sense, item familiarity protects the viewer against the deficit that can arise when viewing is disrupted. As such, the present results do not support the prediction based on expertise theory, or the null hypothesis based on selected previous studies. Instead, these results wholly support the predictions arising from the robust representations perspective (Tong & Nakayama, 1999).

Having said this, examination of the present accuracy data reveals the possibility that the smaller effect of inversion when faces are familiar arises because of performance at ceiling in the 'familiar upright' condition. In other words, although performance could improve from 'familiar inverted' to 'familiar upright' conditions, performance was so good already that it may not have been able to improve to the extent that performance on the unfamiliar faces could. If this is correct, then the interaction of familiarity with inversion in the present study could be the result of a ceiling effect when processing familiar stimuli.

Several modifications to the present task would allow this possibility to be examined. First, a shorter exposure duration at study may make the old/new recognition task more difficult and hence reduce performance from its ceiling levels in these conditions. Second, the use of a larger target item set may increase the cognitive load such that, again, performance levels do not approach ceiling. Finally, the use of different stimuli of a lower familiarity level may be effective in reducing task performance. This latter modification has the advantage of being theoretically driven rather than methodologically driven. If, indeed, robust representations are responsible for the reduced impact of inversion on familiar items, then it might be expected that if the item were still familiar but the representation not so robust, then inversion may have a larger effect on performance, and may approach the influence that has been shown for the unfamiliar items. The following experiment will allow this possibility to be examined.

Examination of the present results also reveals a very interesting finding in terms of response bias. Experiment One suggested that when presented with inverted faces, participants showed a cautious bias to respond 'unfamiliar' no matter what the status of the item. Here, however, the results suggest a slightly different picture. Instead of a bias when faces were inverted, the present results suggest a bias when faces are familiar. Participants had a tendency to respond 'old' to familiar stimuli, and 'new' to unfamiliar ones no matter if they had been seen or not.

This result is intriguing, but can be explained with reference to the notion of *fluency of processing* (Jacoby, Kelley & Dywan, 1989) from the memory literature. Fluent processing refers to the ease with which items are processed in memory. An item can be processed fluently for several reasons. It may have been seen recently and hence its recognition has been primed by prior exposure. It may be highly familiar and hence it is likely to have been seen recently *and* it is likely to have a strong representation that is easily activated. From this, it is clear that highly familiar items have the potential to be processed with fluency and this may be mistaken for having been seen recently in a prior phase of the experiment. As a result, a tendency to say 'old' to familiar faces would result. Of course, for the genuinely old items, the response would be correct. However, for the new items, the source of fluency is the presence of a stored representation facilitating performance rather than a recent presentation and thus an error resulting from *false fluency* is evident (see Whittlesea & Williams, 1998).

As hinted above, the notion of fluency of processing is not unlike the perceptual notion of priming. Processing of an item can be primed by prior exposure such that recognition of an item at time 2 is facilitated by exposure to that item at time 1. This has been shown to occur for all faces, and can successfully be modelled using a neural network such as IAC (Burton, Bruce, & Johnston, 1990). However, recent empirical and computational work has suggested that the degree of priming can be influenced by the familiarity of the stimulus such that familiar items generate less priming than unfamiliar ones (Stevenage, *in prep*). The consequence of such an influence of familiarity would be that familiar items should generate less priming and so less false fluency. This clearly does not fit with the data from the present study.

Consequently it can only be concluded that the source of the false fluency observed here is the result of existence of a strong (familiar) mental representation that can be easily activated, over and above any influence of prior exposure.

Conclusions

The results of the present study provide strong evidence to suggest that familiarity can modify the effect that inversion has on face processing. In an old/new recognition task in which identification of the faces was not explicitly called for, inversion had a detrimental effect for all stimuli but this was clearly smaller when faces were familiar. As such, these results fail to support the position arising from the expertise literature that expert (familiar) items would be more affected by inversion than less expert (unfamiliar) ones. In addition, these results fail to support the results of previous workers who found that the familiarity of the faces had little bearing on the effect of inversion. Instead, the predictions arising from consideration of robust representations receive support. According to this position, familiar items have the benefit of the formation of a representation that incorporates a wealth of information from different viewpoints, expressions, poses, angles, lightings, hairstyles, etc. Such a robust representation is capable of withstanding changes in viewing conditions such that a smaller disruption to processing results. This is exactly the pattern of results that Tong and Nakayama (1999) present, and this pattern of results is echoed within the current experiment. The next chapter provides a replication of this finding with attention to the issue of ceiling performance through the use of a different set of familiar stimuli.

CHAPTER 4

The Importance of Type of Familiarity in the Demonstration of Differences in Facial Processing.

Abstract

The purpose of the present study is to determine whether the emergence of familiarity as a moderator of the inversion effect in Experiment Two generalises when another stimulus set is used. Specifically, the issue of concern here is whether these results hold, no matter what the basis of familiarity. The present study uses faces that are familiar through belonging to well-known celebrities rather than being familiar through personal contact. An old/new recognition paradigm is used, as in the previous study, in order to compare the effect of inversion on the recognition of familiar (celebrity) and unfamiliar (once seen) faces. Performance was measured using accuracy and speed of response to correct target identifications, along with measures of discrimination (A') and response bias (B''_D). As in Experiment Two, the results suggested that inversion could not remove all traces of familiarity in a known face, and that what remained was sufficient to perform an old/new task. However, the issue of whether stimulus familiarity modified the effect of inversion was unclear. The accuracy data confirmed expectations that processing of familiar faces was less affected by inversion than processing of unfamiliar faces. However, this pattern was not borne out by the sensitivity measure (A') or the response time data These data are discussed with respect to an alteration of task difficulty and a reduction in the robustness of familiar representations.

The Importance of Type of Familiarity in the Demonstration of Differences in Facial Processing.

The previous two experiments have shown mixed results on the issue of whether facial familiarity can moderate the effect of difficult viewing conditions. Experiment One, using a speeded-familiarity decision, showed an effect of inversion that was untouched by the familiarity of the stimulus. In contrast, Experiment Two, using an old/new recognition paradigm, showed a greater effect of inversion when faces were unfamiliar (once seen) than when faces were familiar. The same stimuli were used across both experiments and so the contrasting results cannot simply be attributed to a change in stimulus properties. However, the design used in Experiment Two provided a more appropriate test of the predictions and consequently, it is these data which should carry more weight. Having said this, the results of Experiment Two obviously require replication in order to verify the effects. The purpose of the current study is to provide this replication.

In a theoretical sense, the emergence of a moderating effect of familiarity on stimulus inversion should, perhaps, not surprise us. Indeed, the fact that familiar faces were less affected by inversion than unfamiliar faces fits perfectly with the concept of robust representations put forward by Tong and Nakayama (2000). One pertinent question is then how much familiarity is required to form a sufficiently robust representation that it withstands disruption from altered viewing conditions. This issue is highly relevant to current research in the area of face perception. Indeed, Collishaw and Hole (2000) discuss the issue of whether increasing familiarity allows the viewer to process a face in a qualitatively different way, or whether one merely becomes able to process the face using the same strategies but more efficiently. In this respect, it may be appropriate to investigate the processing of faces of an intermediate level of familiarity. At a practical level, such an investigation would require a set of stimuli which are recognisable but are not as well known as those used in Experiment Two. It is proposed that the faces of celebrities might provide such a stimulus set.

Many studies exploring the nature of familiar face recognition use images of actors, sportsmen and women, pop stars, television personalities, politicians, and others in the public eye (e.g., Collishaw & Hole, 2000; Ellis *et al.*, 1992; Yarmey, 1971). Celebrities are, by definition, well-known people but there are ways in which they differ to personally known faces. For example, a viewer has the opportunity to interact with personally known individuals on a regular basis, whereas interaction with a celebrity is unlikely. As such, one is not privy to the range of changes in expressions, hairstyle, and hair colour, as well as age-related changes, that one sees in those we are personally familiar with. It could be argued that some celebrities (and certainly those in the acting profession) present a range of expressions and changes in appearance congruent with role-playing, and that this allows a strong representation of that particular celebrity to be developed. However, the big difference is that, in general, celebrities are viewed from afar or via a two-dimensional screen, and this has the potential to limit the robustness of any resultant representation.

One concern that cannot be ignored when using the faces of celebrities within an experimental setting is one of *fluency of processing*. Within an old/new recognition paradigm, where the task is to identify whether a given face has been seen in a prior study phase or not, the faces of celebrities are particularly prone to attracting an 'old' response. Indeed, the results of Experiment Two confirm this very pattern. This situation arises because it can be difficult for a viewer to determine whether a familiar celebrity is familiar from the experimental context, or is familiar from 'life' (Whittlesea & Williams, 1998). A recent study by Bruce, Carson, Burton, and Kelly (1998) examined exactly this issue by investigating context effects in priming recognition of familiar faces. Bruce *et al.* compared the effects of priming obtained in two different contexts, the first, from the recruitment posters used to attract participants to the experiment, and second, from a laboratory task which preceded the test phase of the experiment. Although priming effects were seen for both contexts, the results indicated that priming was greater for the faces seen in the laboratory than on the posters. Furthermore, priming was seen to be greatest when identical pictures were seen on the posters and at test. Based on debriefing sessions of their participants, Bruce *et al.* submit that episodic memory for faces seen out of the context of the laboratory was weak. This suggests that participants have little problem in recalling whether a presented face was last seen as part of an experimental procedure, or on the television.

With this in mind, the focus of the present study is twofold. First, this study serves as a replication of Experiment Two in order to verify the impact of familiarity on inversion effects. Second, the present study examines the effect of differing levels of familiarity on the pattern of results through using faces that are familiar because they are famous rather than because they are personally known. An old/new recognition task is used again here for the purposes of comparability and the same set of unfamiliar faces as used in Experiments One and Two serves to provide a baseline level of performance. In line with Experiment Two, it is predicted that familiar faces will be recognised faster and more accurately than unfamiliar ones. In addition, it is predicted that upright faces will be recognised faster and more accurately than inverted ones. The results of Experiment Two will receive support if, in addition to the main effects, the effect of familiarity is greater when faces are unfamiliar than when faces are familiar. Finally, the present study allows a re-examination of the issue of response bias in responding 'old' for familiar faces and 'new' for unfamiliar faces.

Method

Design

The design for this experiment was identical to Experiment One. The independent variables were stimulus orientation (upright, inverted) and stimulus familiarity (familiar, unfamiliar). Performance on the recognition task was measured using accuracy, speed of correct responses, a measure of sensitivity (A') and bias (B''_D). As before, the use of an old/new paradigm suggests the unfamiliar faces be more properly referred to as 'once seen'.

Participants

Fifty volunteers (35 male, 15 female) participated in the experiment. These were drawn from the undergraduate and postgraduate population within the Psychology Department and undergraduates were given course credit for their participation. Participant age ranged from 17 to 46 years (mean = 24.8 years, sd = 7.9) and all had normal, or corrected-to-normal vision.

As before, stimulus presentation was counterbalanced across two subsets of participants such that one subset saw half of the targets and distractors upright at test while the other subset saw these stimuli inverted at test. Participants were randomly assigned to one of these two subsets and age and gender of participants within the two groups were matched as far as possible.

Materials

A total of 72 full-face photographs were used in the present experiment. Half of these were celebrities and were thus familiar through either stage, screen, music or sport (see Appendix D). The celebrity faces were obtained via the Internet and were piloted before inclusion in this study in order to ensure the familiarity of the person and the recognisability of the image.

The remaining faces were unfamiliar to the participants and were identical to the unfamiliar faces used in Experiment Two (Appendix C). An equal number of male and female faces comprised each set, and all were of Caucasian origin. Images were selected to have few distinguishing features, and were free of facial hair, scars and spectacles. Every effort was made to match a familiar (celebrity) face with an unfamiliar equivalent on the basis of age, hair colour, hair length, and attractiveness as assessed by visual inspection. All images were prepared within Corel PhotoPaint in the same way as for Experiment Two to produce a well-controlled set of stimuli (figure 4.1)

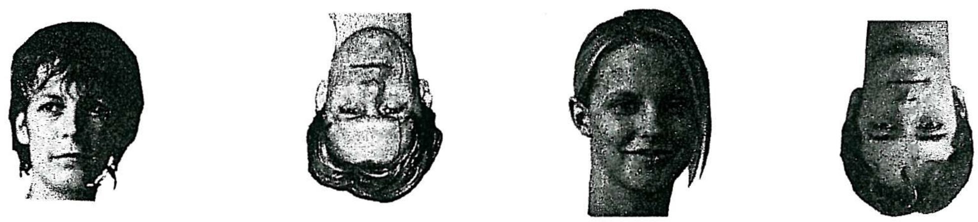

Figure 4.1: Examples of familiar (celebrity) faces (Jamie Lee Curtis and Bryan Adams on the left), and unfamiliar upright and inverted stimuli (right).

The familiar and unfamiliar face sets were each divided into three subsets to provide 12 targets, 12 distractors, and 12 filler faces. The targets were viewed at both study and test. The distractors were seen at test only, and the filler faces were seen at study only. Finally, two versions of each of the target and distractor faces was generated to create one upright image and one inverted image. This resulted in 120 images in total (12 upright targets, 12 inverted targets, 12 upright distractors, 12 inverted distractors, and 12 fillers for each of the familiar and unfamiliar face sets).

The stimuli were presented on a Dell Optiplex GL 575 computer, using a 15" 16-bit colour monitor and a screen resolution of 800 x 600 pixels. Participants sat approximately 50 cm from the computer monitor and images measured approximately 6 x 8 cm (28.34 pixels/cm).

Procedure

The procedure was identical to that used in Experiment Two. As before, the experiment consisted of two phases: a study phase, and a test phase. During both phases, presentation of stimuli was blocked according to the familiarity of the stimuli, and the order of blocks was counterbalanced across participants.

During the study phase, images were presented upright, one at a time, for a duration of 3 seconds. Instead of rating perceived age, as in Experiment Two, participants here rated perceived attractiveness. Participants used a 7-point attractiveness rating scale, ranging from 'not at all attractive' to 'very attractive indeed', and made their response by using the computer mouse to click on the point of the on-screen scale which corresponded with their perception. Immediately after the rating was made, the next image appeared and this sequence continued until all target and filler faces had been viewed. Participants were not warned that a recognition test would follow.

Following the study phase, a surprise recognition memory test was administered. This allowed for a self-paced break and a series of practice trials before the test trials-proper began. The practice trials involved the presentation of a series of upright and inverted familiar and unfamiliar (once seen) faces (these images were not used in the test phase). Participants were warned that some of the faces may be inverted but the task, regardless of orientation, was to determine whether each face had been seen before ('old') or had not ('new'). The images were presented one at a time and remained in view until the participant responded. Participants responded by using their dominant hand to press the appropriate button on a two-button response pad and care was taken to counterbalance the labelling of the response buttons across participants. Following the participant's response, a 500 msec ISI allowed the participant to reset before the next image was presented. During this practise series, feedback on the accuracy of each decision was provided and the main test phase began only when participants had reached a criterion of five correct responses in a row.

The main test phase involved the presentation of the 24 target faces (12 familiar, 12 once seen) together with 24 distractor faces (12 familiar, 12 unfamiliar).

Half of each face set were presented upright and the remainder were presented upside down and the orientation of each face was counterbalanced across the two subsets of participants. As in the practise series and regardless of orientation, the task was to determine whether the face was 'old' or 'new'. No feedback was provided. Care was taken to ensure that the first two and last two images in this test sequence were additional distractor images. Data for these images were omitted from all analyses in recognition of possible practice and fatigue effects.

The study and test phases together took approximately 15 minutes to complete and participants reported no noticeable fatigue. Following completion of the experiment, participants were shown a hard copy of the familiar faces and were asked to name each one. This was to ensure that each participant could recognise the images of each of the celebrities. Mean post-test identification of the 24 famous faces was 90.7% (sd = 8.8). Finally, participants were thanked and debriefed.

Results

All participants performed adequately on the post-experimental familiarity test, and so data from all 50 participants was included for analysis. There were no order effects, or within-group gender differences, therefore the data for male and female participants across the two groups were pooled. As with Experiment Two, accuracy data in terms of proportion of errors, together with speed of correct response to target stimuli were analysed. Response times which were greater than 2.5 sd + mean were excluded from the analysis. Measures of sensitivity and response bias were also calculated and analysed, and results are presented in this sequence.

Accuracy

It was predicted that familiar faces would produce fewer errors than unfamiliar (once seen) faces, and similarly, upright faces would produce fewer errors than inverted ones. On the basis of Experiment Two, an interaction between these two variables was also predicted. Accuracy of performance, expressed as the proportion of incorrect responses, is summarised in Table 4.1 below. From this it is clear that the present results seem to mirror those obtained in Experiment Two. A repeated measures ANOVA confirmed a main effect of familiarity ($F(1, 49) = 28.78, p < .001$) and a main effect of orientation ($F(1, 49) = 54.43, p < .001$). The interaction of these two terms tended to significance ($F(1, 49) = 3.22, p = .07$). Examination of the planned comparisons revealed that inversion affected the accuracy of recognising both familiar ($F(1, 49) = 13.42, p < .001$) and unfamiliar faces ($F(1, 49) = 18.22, p < .001$) but that the inversion effect was, indeed larger for the unfamiliar stimuli, as expected ($t(49) = 1.79, p_{1\text{-tailed}} < .05$). As such, these results support those obtained in Experiment Two, although the magnitude of the interaction effect is somewhat smaller.

Table 4.1: *Mean proportion of errors (and standard deviation) for familiar and unfamiliar faces when upright and inverted.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.05	(.11)	.18	(.16)	50
Unfamiliar	.12	(.13)	.32	(.22)	50

Response Time for Correct Decisions

Response times for correct decisions to familiar and unfamiliar target faces were analysed. As with the previous experiment, it was predicted that participants would respond faster to familiar faces and to upright faces, and that an interaction of these variables may emerge. Prior to analysis, the present data were examined to remove any outliers for each individual. Outliers were classed as being above or below the mean response time + 2 standard deviations. This resulted in the removal of 2.4 % of the data overall.

The remaining data are summarised in Table 4.2 below from which it appears that response times are, indeed, faster when stimuli are familiar and when stimuli are upright. A repeated measures ANOVA confirmed a main effect of familiarity ($F(1, 49) = 4.15, p < .05$) and a main effect of orientation ($F(1, 49) = 83.45, p < .001$). Surprisingly, however, the interaction of familiarity and orientation failed to reach significance ($F(1, 49) = .50, p > .05$).

Table 4.2. *Response times (msec) and standard deviations for familiar and unfamiliar upright and inverted faces.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	1133	(312)	1247	(324)	50
Unfamiliar	1651	(517)	1759	(761)	50

Sensitivity

Analysis of the error rates for target stimuli allows only half the data to be incorporated into the analysis. A', as a measure of sensitivity, allows responses to the distractor stimuli to be included as well as responses to the target stimuli. Again, predictions were made on the basis of the results from Experiment One to the effect that better sensitivity was expected for familiar over unfamiliar faces and for upright over inverted faces. An interaction of these two variables was also expected. Table 4.3 below summarises the A' values across the four experimental conditions. From this, the data seem to support the predictions.

Table 4.3: *Sensitivity of response (A') and standard deviations for familiar and unfamiliar upright and inverted faces.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.91	(.08)	.81	(.15)	50
Unfamiliar	.82	(.14)	.72	(.16)	50

Analysis by means of a repeated measures ANOVA confirms the existence of a main effect of familiarity ($F(1, 49) = 27.24, p < .001$) and a main effect of orientation ($F(1, 49) = 44.42, p < .001$), as expected. However, the interaction between familiarity and orientation did not reach significance ($F(1, 49) = .06, p >$

.05). As such, the sensitivity data do not entirely mirror the pattern of results produced by the accuracy data.

Bias

As before, B"$_D$ values were calculated in order to investigate response bias. These values are summarised in Figure 4.2, with positive values indicating a bias to say 'new' and negative values indicating a bias to say 'old'. The data suggested that familiar faces and upright faces showed a tendency to elicit an 'old' response.

Analysis by means of a repeated measured ANOVA confirmed the main effect of orientation ($F(1, 49) = 6.31, p < .05$). However, both the main effect of familiarity ($F(1, 49) = .84, p > .05$) and the interaction of familiarity with orientation ($F(1, 49) = .98, p > .05$) failed to reach significance. This pattern of response bias does not accord with the predictions from Experiment Two. Rather than familiar faces eliciting an 'old' response, it was the upright faces that elicited the 'old' response.

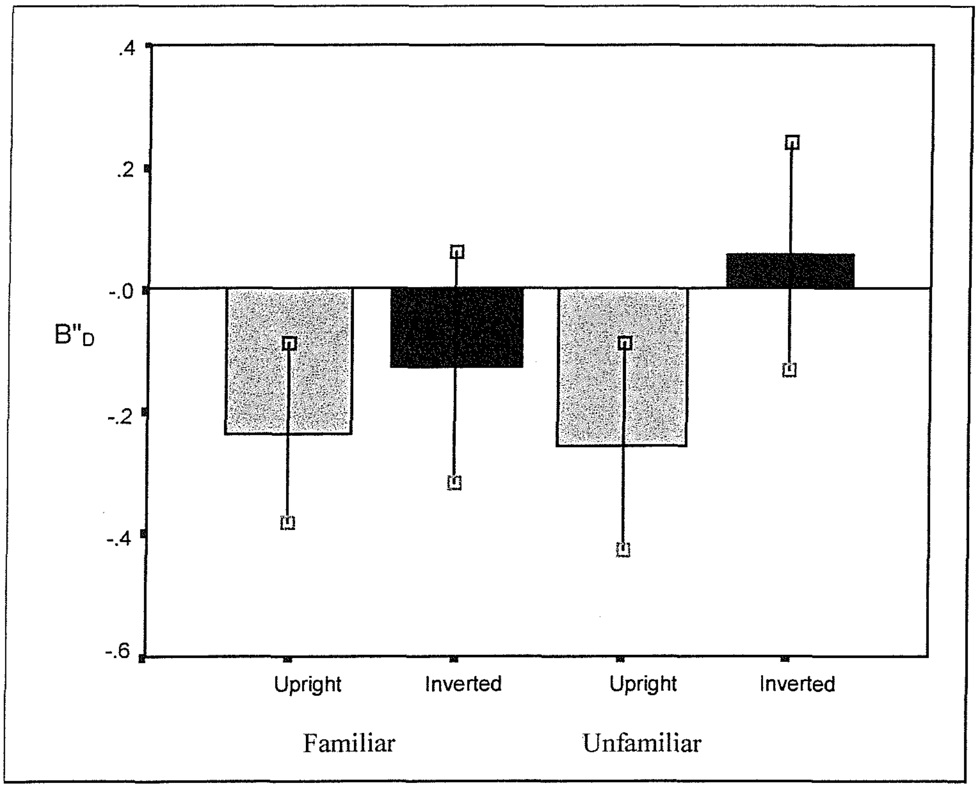

Figure 4.2: Mean B"$_D$ values (and standard deviations) for familiar and unfamiliar upright and inverted faces.

Summary of Results

The present experiment used familiar stimuli drawn from a celebrity pool, rather than a 'personal acquaintance' pool. As such, the familiar stimuli could be expected to have a different basis, and possibly a different strength of mental representation associated with them. The results of the present study still confirm the pattern of performance in Experiment Two in that the effect of inversion on recognition was affected by stimulus familiarity. This effect is, however, confined to the accuracy data, with both sensitivity and response time data showing no moderating effect of familiarity. Finally, response bias data suggests a bias to call the more usual upright orientation 'old'. No false fluency due to familiarity was apparent.

Discussion

The present results have contributed to the emerging picture regarding the influence of familiarity on the disruptive effect of inversion when processing faces. The pattern of results provided here is largely in line with predictions. Overall analyses show that familiar faces were, again, processing with greater accuracy and speed than unfamiliar faces, and that upright faces were processed with greater accuracy and speed than inverted faces. These main effects were expected, and go to support the wealth of evidence on the supremacy of processing familiar faces and the disruptive influence of inversion.

What is more pertinent and much more important, however, is consideration of the mediating effect that facial familiarity can have on the inversion effect. Based on the literature, arguments could be made to suggest that familiar faces, if expertly processed, should be *more affected* by inversion than unfamiliar faces. Equally, empirical literature might suggest that *no qualitative difference exists* between the processing of familiar and unfamiliar faces. Finally, the idea of robust representations would predict that familiar faces, having robust representations, should *be less affected* by inversion than unfamiliar faces. The results of Experiment Two supported the latter argument. The results of the present experiment provide support for this position – again, familiar face processing appeared to be less affected by inversion than unfamiliar face processing.

The moderating effect of familiarity on inversion shown here is, however, substantially smaller than the effect noted in Experiment Two. Indeed, familiarity only influenced the magnitude of the inversion effect when accuracy (proportion of errors) was considered. Examination of the sensitivity measure (A') and the speed of correct responses showed no evidence for any difference across familiarity in the effect of inversion. This pattern of results warrants consideration in order to determine the source of the difference across Experiments Two and Three.

First, it should be remembered that the purpose of the present experiment was to replicate Experiment Two but with a more difficult stimulus set. This action was taken in order to minimise the possibility of ceiling level performance affecting the pattern of results. Consequently, the items making up the familiar face set were drawn from a pool of celebrities rather than from a pool of personal contacts. It was anticipated that the participants in the current experiment would be unlikely to be personally familiar with the celebrity stimuli. As such, the basis for celebrity familiarity could be argued to be somewhat artificial and two-dimensional. In terms of the notion of robust representations, it might be anticipated that known celebrities would have less robust representations associated with them compared to personally known contacts. The processing of these stimulus images should thus render the old/new recognition task more difficult than in Experiment Two and remove or reduce the potential for ceiling effects.

Examination of the accuracy data (expressed as proportion of errors) confirms a slight increase in difficulty from Experiment Two to the present experiment. Indeed, consideration of performance on familiar faces only shows an increase in error rate from 8.5% (personally known) to 11.5% (celebrity) overall, and an increase in errors from 4% (sd = 9%) to 5% (sd = 11%)[2] for the experimental condition at risk of ceiling performance (upright familiar presentations). Although small, this increase in task difficulty may have been sufficient to minimise the risk that the previous findings were the result of ceiling level of performance reducing the inversion effect for familiar faces. The fact that the pattern of results again shows a moderating effect of facial familiarity on the inversion effect therefore provides support for the previous results and suggests that, while possible, it is unlikely that the previous findings were the result of a ceiling effect in the familiar case.

[2] These figures are based on the error rate averaged across familiarity for Experiments Two and Three.

Theoretically, however, it is important to consider the implications of weaker celebrity representations further. On the basis of the argument that a robust representation protects the viewer from marked disruption of face processing due to impaired viewing, it would follow that a less than robust representation would offer the viewer less protection than a more robust representation. Given this line of reasoning, it should not surprise us that, as well as making the task more difficult, the moderating effect of familiarity on inversion shown here (using celebrity faces) is reduced. It should also not surprise us that the moderating effect of familiarity emerges for the accuracy data only. An increase in task difficulty overall very often means that significant effects become apparent only in the accuracy of decisions. The fact that the influence of familiarity on inversion emerges for the accuracy data but not the sensitivity data is, however, curious and unexpected, and points again to the fact that the influence of familiarity on the present inversion effect is more fragile here than in the previous experiment.

Accounting for the present pattern of results with reference to a weaker mental representation for celebrity faces over personally familiar faces gains some support from consideration of the pattern of response bias exhibited across the two experiments. In Experiment Two (personally familiar faces), participants exhibited a response bias to say 'old' to familiar faces and 'new' to unfamiliar faces, no matter whether they had been seen in the study phase or not. This was explained with reference to the notion of fluency of processing, with the belief that a high level of familiarity could result in fluent processing which was mistaken for 'oldness' rather than familiarity. In the present experiment, however, the pattern of response bias was somewhat different. Instead of showing a bias to respond 'old' to familiar faces, participants displayed a bias to say 'old' to upright faces.

This pattern of results can be explained in several ways. First, it could imply that the celebrity faces used in the present experiment were, indeed, represented with insufficient robustness to generated the feelings of false fluency shown previously. This would lend support to the explanations being advanced on the basis of representational strength.

An alternative way of accounting for the lack of 'old' bias for celebrities does, however, exist. This rests on the principles of association and specifically associative priming, which occurs when the presentation of one individual facilitates the processing of another, related or associated individual. In Experiment Two, the familiar faces were all members of the Psychology Department. They were thus well known to the participants, but also had high associative strength to the context of psychology experiments in general. In some sense then, these faces fit well within an experiment, and the decision of whether each face had been seen from the study phase, or from the department corridors might well have been a difficult decision to make. The faces also had high associated strength to one another, through their semantic grouping and the fact that they are often seen concurrently. Thus, the presentation of one professor's face might well prime the processing of another professor's face. As such, familiar target faces at test may prime the processing of familiar distractor faces at test, and the fluency that results may be misinterpreted as 'oldness', hence the response bias in Experiment Two.

The potential for this associative priming mechanism to be at work in the present experiment, is, however, much reduced. The familiar faces are celebrities who share much less in common both with one another and with the experimental context. Hence, any associative priming that occurs will be based on tenuous semantic associations based on, perhaps, a common occupation, rather than the strong semantic associations shared by the familiar stimuli in Experiment Two.

Examination of the procedures used in Experiments Two and Three suggests one final source of disparity which may account for the differences in results. At the study phase, participants in Experiment Two performed an age-judgement for personally familiar and filler faces. In contrast, participants in Experiment Three performed an attractiveness rating for celebrity and filler faces. The difference in task during this study phase was justified on the basis of the fact that it may have been awkward and somewhat inappropriate for student participants to give attractiveness ratings for their tutors and lecturers. Equally, possession of semantic knowledge about actors and celebrities may have made an age-judgement task too easy for the celebrity faces with the result that the participants may have failed to engage with the faces at study. The result, however, is that the difference in task at study may have encouraged participants to use different levels of processing across the two experiments. This explanation would account for the present pattern of results if it is assumed that processing attractiveness (Experiment Three) leads to a lower level of engagement with the target stimuli than processing age (Experiment Two), with the result that performance overall would be impaired. Although possible, this methodological explanation for the difference across experiments is a little contrived. With hindsight, however, it would clearly have been more appropriate to equalise task demands across the two experiments.

Conclusions

The results of the present study have confirmed the prediction that facial familiarity can moderate the impact of inversion when processing faces. As in the previous experiment, familiar faces were less affected by inversion than unfamiliar faces in an old/new recognition task. As such, these results support the predictions to arise from consideration of robust representations, and cast doubt on the expertise suggestion that familiar faces should be more affected, or the empirical suggestion that familiar faces should be as affected, as unfamiliar ones. It was notable, however, that the magnitude of the influence of familiarity on inversion shown here was considerably smaller than that shown in Experiment Two. This was discussed with respect to the notion of relative robustness of representations across celebrity and personally familiar faces.

CHAPTER 5

Changing the Manipulation from Inversion to Negation: Converging Evidence for a Role of Facial Familiarity.

Abstract

The previous two experiments have demonstrated the influence of familiarity in moderating the effect of inversion on recognition. This study will examine the stability of this influence, given a change in the manipulation from inversion to photographic negation. The same task as in Experiments Two and Three was used, where familiar and unfamiliar faces were presented in an old/new recognition task. The results confirmed that familiarity preserved the recognition of faces when presented in photographic negative, and this was evident for the accuracy and sensitivity data. These findings add to those demonstrated in the previous two experiments and support the hypothesis that familiarity results in the formation of robust representations for faces which aid recognition under difficult viewing conditions.

Changing the Manipulation from Inversion to Negation: Converging Evidence for a Role of Facial Familiarity.

The previous two chapters contained details of experimental work which have provided evidence for an influential role of familiarity in a difficult recognition task. The robustness of this role may be investigated by changing the manipulation responsible for disrupting viewing. One such manipulation is to use photographic negatives. Again, the purpose is to compare recognition across familiarity.

Faces are three-dimensional objects, with facial features comprising height, width, and depth. For example, my nose is 5.2 cm long (from bridge to tip), 4 cm wide (across widest part) and stands out 3.2 cms from my face. The subtle variations in measurements across these three dimensions are among the primary factors making each face unique (along with skin texture, pigmentation, etc.). A front view of a face will allow discernment of the first two of these dimensions, but not of depth. In order to extract a three-dimensional description, a number of depth cues are required, and one of these is shape from shading (Bruce, Green, & Georgeson, 1997).

Shape from shading refers to the information derived from the pattern of light intensities reflected from the facial surface. This pattern will vary depending on the level of undulation of the facial features, and the amount and direction of the light source. For instance, changing the position of a light source will alter the subtle pattern of shadows cast over the face, and so will alter the apparent size and shape of the facial features, with dramatic effect on recognition performance. This can be demonstrated both when the position of the light source is altered relative to an unmoved face (see Hill & Bruce, 1996; Johnston, Hill & Carman, 1992) and when the orientation of the face is altered relative to an unmoved light source (Enns & Shore, 1997). Under both circumstances, shadows cast by the orbital ridges, nose, and cheekbones will be altered and performance as demonstrated by a matching task (Hill & Bruce, 1996) is consequently impaired.

Facial shadows are captured by the low spatial frequencies within a photographic image, and Hayes (1988) believes that the deleterious influence that a photographic negative has can be attributed to the corruption of these low spatial frequencies. In this way, processing of photographs is affected by stimulus negation, but processing of line images (which contain primarily high spatial frequency information) is not. In this vein, there is now mounting evidence to suggest that it may be exactly this shape from shading information that is affected by photographic negation (e.g., Cavanagh & Leclerc, 1989; Galper & Hochberg, 1971; Kemp, McManus, & Pigott, 1990). Disruption of the processing of shape from shading in this way may lead to difficulties in the perception of second order relational features, or configural properties within the face, hence the consequent difficulty in recognition.

Put this way, there is an evident similarity between the effect that stimulus negation can have and the effect that stimulus inversion has. Indeed, Maurer, LeGrand, and Mondloch (2002) have recently suggested that both manipulations may disrupt the same thing – configural processing. If this is so, and the reasoning of Collishaw and Hole (2000) is adopted, then the presentation of an image that is both inverted and in photographic negative should result in performance that is not substantially different from that to an image that is either inverted or in photographic negative. This is because the combination of manipulations would still only disrupt the method of processing that each manipulation does alone. A study by Kemp, McManus & Pigott (1990), however, presents evidence to suggest that the effects of inversion and negation are additive in a perceptual task where participants were required to locate facial features. This would suggest that two different processes are being affected by the two manipulations, and that these combine to impede performance. In the same vein, Hole, George, and Dunsmore (1999) present evidence to suggest that chimeric face effects, which rely on the perception of relational information, can be demonstrated for both photographic positives and photographic negatives. In this sense, negation cannot remove the ability to process relational information.

At this point, it is important to note that negation also has the effect of changing the pigmentation, or relative brightness levels, of the face, and this alone may make recognition difficult (Bruce & Langton, 1994). Indeed, the reversal of brightness levels seen in negation means that dark hair becomes light, and light hair will appear dark. Similarly, eyes and eyebrows will have their brightness reversed. As such, the identity of the face is effectively disguised. In fact, Bruce and Langton (1994) indicate that this change in pigmentation is more important than disruption to the extraction of shape from shading in mediating the influence of photographic negatives. Notwithstanding this demonstration, the potential similarity between the effects of inversion and photographic negation make it pertinent to investigate this manipulation in the light of the findings of the previous recognition studies (Experiments Two and Three).

The present study will use the same stimulus set as in Experiment Three, except that the manipulation of the independent variable will involve reversing the contrast of a stimulus face, rather than inversion. In terms of predictions and in accordance with the remainder of this thesis, two hypotheses may be tested against the null. First, if the expertise hypothesis is correct, and familiar stimuli are processed with better, or greater reliance on configural processing, then a manipulation which threatens configural processing will have a greater effect on familiar than unfamiliar faces. Second, if the robust representations hypothesis is correct, and familiar stimuli are represented by more detailed visual codes, then familiar stimuli will be processed better than unfamiliar ones no matter what the manipulation. Based on the results of Experiments Two and Three, the latter hypothesis is favoured.

Method

Design

The design for this experiment replicated the previous study in that it was a 2x2 within-subjects design. However, instead of manipulating orientation, image contrast (positive, negative) formed the first independent variable and stimulus familiarity (familiar, unfamiliar) formed the second. The dependent variables were accuracy (proportion of errors), sensitivity (A') and bias (B''_D), and speed of response (in milliseconds). In line with Experiments Two and Three, the unfamiliar faces are more properly referred to as 'once seen'.

Participants

Twenty-two volunteers (18 females, 4 males) participated in the experiment. These undergraduate students were recruited from the Psychology Department participant pool and were given course credit for their participation. Participant age ranged from 18 to 41 years (m = 20.9 years, sd = 4.91 years). All participants had normal or corrected-to-normal vision.

Materials

The stimuli were identical to the set used in Experiment Three, except that instead of generating inverted images of the publicly familiar and unfamiliar faces, reversed contrast (photographic negative) images were used. As before, the stimuli were full-face greyscale photographs, comprising 36 celebrities, and 36 unknown faces (see Appendix E for full set of stimuli). An equal number of male and female faces were used, and all were Caucasian. Images were selected to have few distinguishing features; that is, they had no facial hair or scars, and did not wear spectacles. Both famous and unfamiliar sets of faces were split into three groups, target images, distractor images and filler images. Targets were presented in both the inspection and test phases, distracters were seen in the test phase only, and fillers were presented in the inspection phase only. Faces were scaled, mounted and presented in the same way as for the previous studies. In total, 72 facial images were presented, using identical equipment to that in the previous studies. Images were presented in 8 bit greyscale and measured approximately 6 by 8 cms.

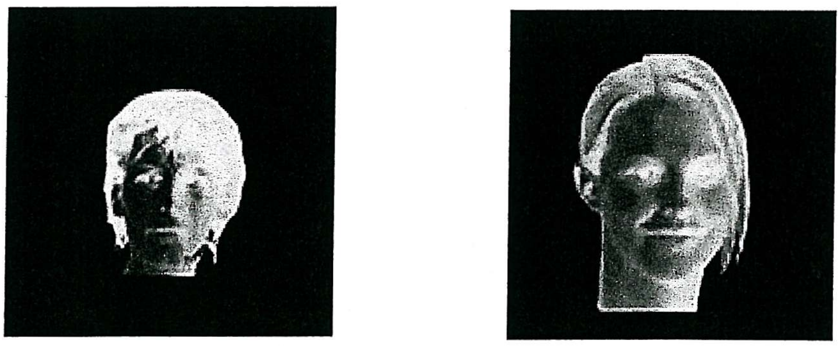

Figure 5.1. Examples of familiar (Jamie Lee Curtis on left) and unfamiliar reversed contrast stimuli.

Procedure

As before, participants were seated approximately 50 cms from the monitor and the experiment consisted of two parts; an inspection phase, and a test phase. Presentation of the stimuli in both phases was blocked by familiarity. In the inspection phase of each block, participants were shown positive items, one at a time, for 3 seconds, and were required to examine and rate each face for attractiveness, using the 7-point scale used previously. No prior warning of a recognition test was given. The test phase proceeded identically to the previous two experiments. A post-experimental identification test was used to ensure familiarity with the celebrity stimulus set, and this revealed a mean identification accuracy of 92.3% (sd = 8.7%). Finally, participants were thanked and debriefed.

Results

Results from all 22 participants were analysed. There were no order or gender effects therefore data were collapsed across groups. As previously, accuracy (proportion of errors), sensitivity (A'), bias (B''_D) and response latencies (in msec) were analysed. As before, only target data were analysed, and response times which were greater than 2.5 sd + mean were excluded from the analysis. Comparisons of results with Experiment 3 are made.

Accuracy

Based on the evidence of the previous study (Experiment Three), it was predicted that familiar faces would produce fewer errors than unfamiliar (once seen) faces, and positive faces fewer errors than negative faces and, as expected, this was found to be the case (Table 5.1). The high error rate for recognition of negative unfamiliar faces testifies to the difficulty of the task in comparison to the previous inversion tasks.

Table 5.1. *Mean proportion of errors for familiar and unfamiliar faces when in photographic positive and negative.*

	Positive		Negative		N
	Mean	(SD)	Mean	(SD)	
Familiar	.02	(.05)	.16	(.12)	22
Unfamiliar	.14	(.23)	.43	(.22)	22

A repeated-measures ANOVA revealed main effects of familiarity (F (1, 21) = 22.84, $p < .001$) and contrast (F (1, 21) = 38.02, $p < .001$), and these effects were modified by a significant interaction (F (1, 21) = 4.81, $p < .05$). *Post hoc* analyses of the simple main effects revealed familiarity was influential in normal contrast (F (1, 21) = 6.24, $p < .05$) and especially in reversed contrast (F (1, 21) = 24.35, $p < .001$). Furthermore, examination of the difference scores (familiar – unfamiliar when (a) positive, and (b) negative) showed that the effect of familiarity was significantly greater when stimuli were presented as negatives ($t(21) = 2.19$, $p < .05$). This provides further evidence for the influence of familiarity in the preservation of recognition under difficult viewing conditions. However, it should be noted that the near ceiling performance with the familiar positive stimuli could have restricted the difference between the number of errors made on positive and negative (familiar) stimuli. This factor may be responsible for the reported interaction, therefore results are accepted with a degree of caution.

Response Times

The increased response times for negated images testifies, again, to the difficulty of the task. As predicted, positive faces elicited a faster response than negative faces, and familiar faces produced a faster response than unfamiliar faces and these results are presented in Table 5.2 below. A repeated measures ANOVA on response times showed that although responses to familiar faces were faster, this finding was not significant overall (F (1, 21) = 3.51, $p = .08$). Positive images were, however, processed faster than negative images (F (1, 21) = 34.59, $p < .001$), but there was no interaction between familiarity and contrast (F (1, 21) = 2.40, $p > .05$).

Table 5.2: *Mean response times (with standard deviations) for familiar and unfamiliar face stimuli across all participants.*

	Positive		Negative		N
	Mean	(SD)	Mean	(SD)	
Familiar	1271	(321)	1756	(353)	22
Unfamiliar	1317	(376)	2012	(649)	22

Sensitivity

Analysis of A' (the ability to discriminate between old and new faces) followed the same pattern as the error data and can be seen in Table 5.3 below. Performance on this measure echoed the results of Experiment Three where discrimination was near perfect for familiar positive images, with discrimination of familiar negative images being almost as good as that for unfamiliar positive images. There was a greater disparity between familiar and unfamiliar images when shown as negatives than when shown as positives.

To test for a significant difference between these variables, a repeated-measures ANOVA was performed. This revealed main effects of familiarity (F (1, 21) = 12.22, $p < .01$) and contrast (F (1, 21) = 38.97, $p < .001$), and a significant interaction between these variables (F (1, 21) = 4.69, $p < .05$). This interaction was investigated using *post hoc* contrasts of the simple main effects, from which a significant effect was found for both positive (F (1, 21) = 11.39, $p < .01$) and negative polarities (F (1, 21) = 9.38, $p < .01$). Examination of the difference scores (familiar – unfamiliar when (a) positive and (b) negative) showed, however, that the effect of familiarity was significantly greater when stimuli were presented as photographic negatives (t(21) = 2.17, $p < .05$). As such, these data mirrored the pattern produced by the accuracy data.

Table 5.3: *A' values for familiar and unfamiliar face stimuli in positive and negative polarities.*

	Positive		Negative		N
	Mean	(SD)	Mean	(SD)	
Familiar	.97	(.04)	.87	(.10)	22
Unfamiliar	.90	(.08)	.69	(.23)	22

Bias

Finally, B''_D values were calculated to investigate the existence of any response bias. These are presented in Figure 5.2, and show a tendency to see all positive, and all familiar faces as 'old', although this was less marked for familiar negative faces. In contrast, unfamiliar negated faces were more likely to be called 'new'.

A repeated measures ANOVA with familiarity and contrast as the within-subjects variables revealed a main effect of contrast only ($F(1, 21) = 6.40, p < .05$), with no effect of familiarity ($F(1, 21) = .97, p > .05$) and no interaction between these variables ($F(1, 21) = 2.68, p > .05$). These results are broadly in line with those in the previous experiment, and as such do not show the *false fluency* effect seen in Experiment Two.

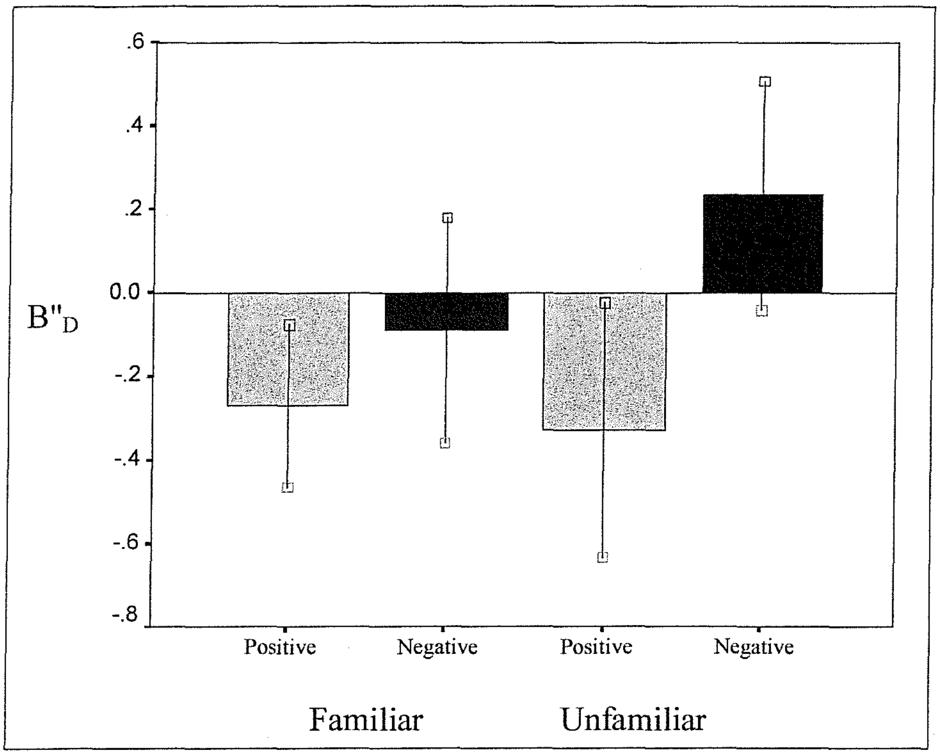

Figure 5.2. Mean B''_D values for positive and negative faces across familiarity.

Summary of results

This experiment used the same facial sets of stimuli as Experiment Three and so any contrast in findings is likely to be attributable to the change in manipulation (reversing contrast as opposed to orientation). Indeed, the findings of this experiment closely match those of the previous chapter and suggest that the manipulation of reversing the contrast from photographic positive to photographic negative elicited the same pattern of response as found in inversion. Furthermore, the results of the present study add further confirmation to the pattern of performance seen in the previously presented recognition tasks, that the impairment caused by disrupting viewing conditions was affected by stimulus familiarity. However, contrary to the previous experiment, this effect is seen not only in the accuracy data, but also in the sensitivity data. As before, response time data showed no moderating effect of familiarity. There was a tendency to see positive images as 'old' but this was not limited to familiar faces and hence there was no *false fluency* effect of familiarity.

Discussion

This study looked at the effect of a different way of manipulating viewing difficulty, and how this manipulation would affect the influence of familiarity in an old/new recognition task. Here, the manipulation used was image negation. It was predicted that recognition performance would be reduced when faces were presented in photographic negative, and this was the case for every measure investigated. Furthermore, it was predicted that familiarity would influence the effect of negation, and indeed, familiarity was found to have a preserving effect on recognition in terms of accuracy (proportion of errors) and sensitivity (A'). Negation was found to disrupt the accuracy and sensitivity of recognition performance especially for the unfamiliar faces. The fact that the influence of familiarity did not carry forward to affect response times as well as accuracy and sensitivity is somewhat surprising given that performance was good. However, the evident variance within the accuracy data make these the primary measures. Finally, the present study has revealed a consistent bias in the pattern of responding, such that participants revealed a tendency to say 'old' when stimuli were presented in the more usual photographic positive condition than when presented as negatives. This mirrors the tendency to say 'old' in the previous experiment when stimuli were presented in the more usual upright orientation than when inverted.

The pattern of results displayed in the present experiment conforms entirely with predictions, and with the pattern of results displayed in both Experiments Two and Three. In all three studies now, familiarity has been shown to influence a recognition-level task in such a way as to make the task easier when viewing conditions are difficult. Consequently, these data provide converging evidence, using a different stimulus manipulation to suggest that familiarity influences the course of face processing. The fact that familiar faces are affected less than unfamiliar ones when stimuli are both inverted and negated is important in terms of the hypotheses under investigation. According to the expertise hypothesis, familiar faces are processed with greater reliance on configural processing. As such they should be affected more by a manipulation that impairs configural processing. The data do not support this position. Neither do they support the null position that there is no difference across familiarity in the influence of a stimulus manipulation. In contrast, according to the robust representations hypothesis, familiar faces possess a more robust representation due to the greater number of instances seen, and the greater number of variations in appearance that have been encountered. As such, familiar faces should be less affected by any manipulation to the viewing conditions. The data do support this hypothesis. Indeed, Tong and Nakayama (1999) have proposed that a wide visual experience of a face over different viewing and lighting conditions will likely strengthen the visual code derived for an individual face. This in turn will lead to a robust representation for that face, such that *any* disruption to viewing is attenuated. Hence, across three studies now, the notion of robust representations emerges as a parsimonious way of explaining all sets of data.

Several other aspects of the present data are also worthy of consideration. First, comparison of the overall level of performance across Experiments Three and Four (in which the same stimuli were used) enables examination of the relative impact of inversion (Experiment Three) and negation (Experiment Four). Comparison of error rates first shows that negation was by far the more difficult manipulation. Performance in the unmanipulated condition (averaged across familiarity) was .085 in Experiment Three (inverted) and .08 in Experiment Four (negated), indicating no real difference in baseline levels of performance. However, the impact of image manipulation (manipulated – unmanipulated) shows a greater influence of negation over inversion not so much for familiar faces, which are protected by the robust representation (increase in error rate of .13 when inverted, and .14 when negated) but clearly when faces are unfamiliar (increase in error rate of .20 when inverted, and .29 when negated). This pattern of a greater influence of negation on task performance is echoed when the RT data are examined. Here, performance in the baseline conditions is again broadly comparable, however, the effect of image manipulation (manipulated RT – unmanipulated RT) again shows a far greater effect of negation than inversion for both familiar faces (increase in RT = 114 msecs when inverted; 485 msecs when negated) and unfamiliar ones (increase in RT = 108 msecs when inverted; 695 msecs when negated). These results thus confirm those of Kemp McManus and Pigott (1990) and Bruce and Langton (1994) in that image negation appears to disrupt processing more than image inversion. Given that, the effects of facial familiarity shown when negation is used might well be expected to be clearer than when inversion is used. The fact that the influence of familiarity is revealed both in accuracy and sensitivity data here but not in Experiment Three (inversion) might attest to this.

The relative impact of negation and inversion, however, raises a difficult theoretical consideration which should be explored. If it is the case that both negation and inversion disrupt the same thing – extraction of configural information – then why should one manipulation have a greater effect on performance than the other? Two possibilities arise. First, it may be the case that negation affects two things – configural processing *and* relative pigmentation within the image, while inversion affects only one thing – configural processing. In this sense, negation might quite reasonably be expected to have a greater effect than inversion, and the additive results of Kemp, McManus & Pigott (1990) are to be expected. Second, it is possible that both manipulations do indeed affect the same thing – extraction of configural information – but that we have a recovery strategy to deal with one manipulation but not the other. Put more specifically, whilst inverted faces (and objects) are not common, they do arise, and we have the capacity to (effortfully) mentally rotate a stimulus to correct for this. In contrast, negated images do not often appear in the real world. As a result we do not have a recovery strategy to hand to enable us to correct the reversal of contrast. Indeed, the figure-ground literature confirms that once perception is biased to see an ambiguous figure a certain way, it is very difficult indeed to reverse the boundaries within the image in order to see the alternative percept. Consequently, the greater effect of negation over inversion need not imply the disruption of different (or additional) processes. They might simply reflect that a recovery strategy is less available, if it is available at all.

Conclusions

Within the present study, the effect of negation has exerted a more disruptive influence than that of inversion (Experiment Three). This pattern is in line with previous results (Bruce & Langton, 1994; Kemp, McManus & Pigott, 1990). The present study has also confirmed the influence of familiarity on the effect of negation such that familiar faces were less disrupted by negation than unfamiliar ones. As such, the influence of familiarity appears to be stable across different manipulations, and this supports the idea that familiarity with a face will modify the effect of *any* disruption caused by difficult viewing. These results have been explained within the context of the notion of robust representations which buffer the viewer against image transformation, regardless of the nature of the transformation. The purpose of the next few chapters is to determine whether this protection extends regardless of the nature of the task.

CHAPTER 6

The Effect of Familiarity in a Face Classification Task using Thatcherised Faces

Abstract

Throughout the previous experiments, familiarity has been shown to affect the visual processing of faces during recognition- and identification-level tasks. The purpose of the present experiment is to determine whether familiarity might also affect the very early visual processing involved in a face classification task. Participants performed a series of 2AFC trials where they were shown normal and Thatcherised versions of a set of faces. On each trial, the task was to identify the 'odd' face. Celebrity and unfamiliar faces were presented both upright and inverted and both accuracy and speed of correct decisions indicated near-ceiling performance when stimuli were upright, but a clear impairment when stimuli were inverted. Critically, the magnitude of the inversion effect depended on the familiarity of the stimulus, with familiar stimuli being protected somewhat against disruption. These results support the pattern of results from previous experiments within this thesis, and are discussed again with reference to the notion of robust representations.

The Effect of Familiarity in a Face Classification Task using Thatcherised Faces

Previous chapters of this thesis have been concerned with high-level processing of the facial image. Questions have been asked surrounding whether facial familiarity can affect the way in which the face is processed to achieve a recognition or an identification as 'old' or 'new'. The evidence presented thus far suggests that, contrary to expectation, facial familiarity can have an influence on processing at this recognition stage.

The purpose of the next few studies is to examine whether facial familiarity can be shown to influence processing at an early visual processing level. Specifically, the question of whether familiarity influences decision making at a category level is addressed. Now, the focus is not on how the image is retrieved but rather on how the image is perceived. This is examined through a face classification task using the Thatcher Illusion as a way of disrupting face-ness.

Thompson (1980) first demonstrated the Thatcher Illusion when he inverted the eyes and mouth of Margaret Thatcher and replaced them within an otherwise upright facial frame. When the image was viewed the right way up, a grotesque expression was apparent. However, when the image was viewed upside-down, the grotesqueness was lost (see Figure 6.1). More recently, Stürzel and Spillman (2000) have revealed that the illusion of grotesqueness is apparent between 0° and approximately 94° but thereafter, the grotesqueness is lost. Furthermore, Lewis and Johnston (1997) showed that the illusion does not depend on a smiling mouth that becomes a 'biting' mouth when rotated. They asked participants simply to say whether the eyes and mouths had been inverted in a set of faces, but they used faces with neutral expressions rather than the smiling expressions used in previous studies. Their results confirmed that when inverted, the decision of orientation of the eyes and mouth relative to the other features became very difficult.

Thompson (1980) proposed that the illusion occurred because the eyes and mouth were important as conveyors of facial information such as mood or expression. This *expression hypothesis* echoed the views of Kohler (1940) and Yin (1969) who believed that inversion impaired the ability to read expression from the face and hence impaired the ability to perceive the grotesqueness.

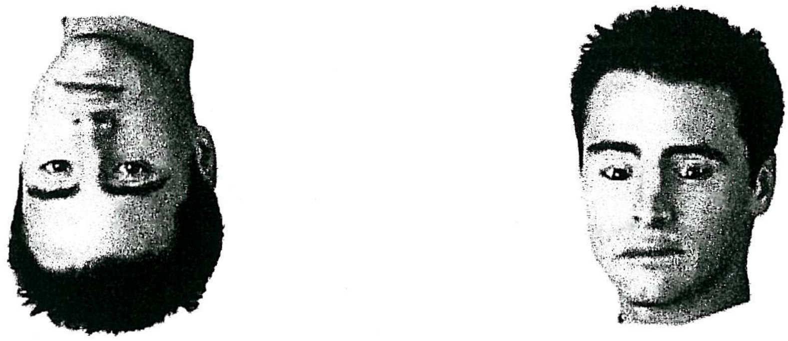

Figure 6.1: An inverted and Thatcherised version of Matt le Blanc's face (left). The grotesque expression only becomes apparent when the image is viewed upright (right).

Bartlett and Searcy (1993), however, rejected the expression hypothesis on the basis of empirical findings that revealed a lesser effect of inversion on an expression-judgement task than on a familiarity-judgement task. Disruption to the perception of expression cannot therefore be the factor that impairs identification of inverted faces.

An alternative account of the Thatcher Illusion is presented by Parks *et al.* (1985). They suggested that what was important was the positioning of the eyes and mouth relative to one another. They used simple schematic features (eyes and mouths) and asked participants to rate the 'pleasantness' of pairs of eyes, and mouths, when presented alone, and when presented above or below one another. Eyes and mouths were presented both upright and inverted. The results suggested, not surprisingly, that upright (smiling) mouths were preferred over inverted (frowning) mouths. Equally, they found that an inverted mouth appeared more pleasant when the eyes were beneath it, as this arrangement gave the overall appearance of an inverted but intact face. This contrasted with the finding that an upright mouth was rated as *less* pleasant when the eyes were placed below it than when the eyes were above. All of these results can be explained by the suggestion that the assignment of 'top' of the stimulus is dependent on the positioning of the eyes.

Parks (1983) suggested that it is the assignment process that produces the Thatcher Illusion, and that the effect should not be limited to faces. He proposed that any familiar frame shown in its normal orientation would assign 'up' to the features within it. This allows any oddities to be seen such as the grotesque expression within the Thatcherised face. However, if the image is presented in an unusual orientation (inverted), the assignment of 'up' is less reliable and so any oddities within the image become difficult to perceive. Parks demonstrated his point using words (see example in Figure 6.2).

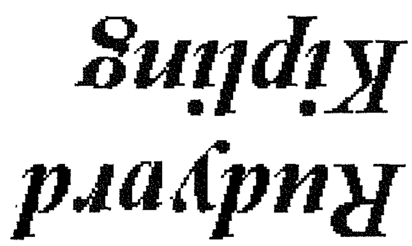

Figure 6.2: At first glance, the name appears upside-down but otherwise intact. Inverting the page reveals the abnormal 'a'.

Bartlett and Searcy (1993) have examined the robustness of Parks' (1983) observations. They reject such a simple 'frames of reference' account on the basis of the fact that it may be interpreted as nothing more than a demonstration of configural processing. The perception of the positioning of one set of features *relative* to another is, in essence, what configural processing is all about. As such, Bartlett and Searcy (1993) contest that it was the poorer processing of configural information in Parks *et al.*'s (1983) inverted eyes-mouth images that was responsible for the pattern of pleasantness ratings.

The most common explanation for the Thatcher Illusion is that, when viewed upright, the face can be processed configurally. That is, the relative position and spacing of features can be determined and so the mismatch between the orientation of features can be perceived. In contrast, when inverted, the ability to process the configural information within the face is compromised. Without the capacity to perceive the mismatch in orientation of one feature relative to another, the grotesque expression goes unnoticed (Bartlett & Searcy, 1993).

Bartlett and Searcy tested their ideas using photographic images of faces in two conditions. The first was an expression condition, where faces were presented with (i) a Thatcherised distortion, (ii) a grotesque expression, and (iii) a neutral expression. These were contrasted with a spatially distorted condition where (i) the eyes were moved down and further apart, and (ii) the mouth was moved upwards, or (iii) the face was normal. Participants were asked to rate faces for grotesqueness. The authors reasoned that if disruption of configural information caused the effect, then inversion would reduce the grotesqueness of the spatially distorted faces as well as the Thatcherised faces, but would not distort the perception of the grotesque expression face. Their findings supported these ideas, although the effect was greatest for the Thatcherised faces. Their findings also highlighted the importance of experimental task. A second experiment using similarity judgements supported the findings of the grotesqueness rating task – inversion affected the judgements for Thatcherised faces and distorted faces but not the judgements for grotesque expression faces. However, a third experiment using a same-different judgement task showed no effect of inversion. This led the authors to suppose that same-different judgements may be performed using a component-based strategy, which could be completed in both upright and inverted orientations. As such, Bartlett and Searcy suggested that the task will determine what type of information is encoded and this will be reflected in the influence of inversion on performance outcome.

More recently, Searcy and Bartlett (1996) compared the effects of inversion on faces with spatial-relational changes (eyes moved up, mouth moved down) and faces with component changes (teeth made to look like 'fangs', eyes reddened). They found that inversion reduced the grotesqueness ratings of the spatially altered faces, but made no difference to the faces with distorted components. This confirmed findings from their earlier study, and again showed that inversion altered the processing of configural information but not component (featural) information.

What is pertinent for the present experiment is the effect that facial familiarity might have on the Thatcher Illusion. In this respect, it is important to note that the illusion has been demonstrated with unfamiliar faces as well as with familiar ones, and has even been demonstrated using schematic faces (Rakover, 1999). What is not known, however, is whether familiarity affects the extent of the Thatcher Illusion. If, as the expertise literature suggests (Rhodes *et al.*, 1989) expert processing is dependent on the use of configural strategies, and familiar faces are processed more expertly than unfamiliar ones, then it might be expected that a greater Thatcher Illusion is evident when faces are familiar than when unfamiliar. In contrast, if the robust representation view of Tong and Nakayama (1999) is considered, then it might be expected that a smaller Thatcher Illusion is evident when faces are familiar than when unfamiliar due to the robustness of the familiar representation acting as a buffer to disrupted viewing. Finally, it might be considered that in a face classification task such as this, there is no requirement to identify the individual and so familiarity might have no effect whatsoever. The present experiment will examine these three opposing viewpoints.

A two-alternate forced choice (2AFC) paradigm will be used, where normal and Thatcherised versions of celebrity and unknown faces will be viewed simultaneously. The task is to indicate which, of the two faces, appears 'odd'. Faces will be presented both upright and inverted. It is anticipated that upright faces, in which configural information can be processed readily, will be responded to with accuracy and speed compared to inverted faces. It is also anticipated that the familiarity of the face, whilst unimportant to the 'normal/odd' decision, may speed processing. Of main importance, however, is the question of whether the Thatcher Illusion remains evident for both familiar and unfamiliar faces when inverted. This will be addressed through examination of the interaction effect for measures of accuracy and speed of response.

Method

Design

A 2 x 2 x 2 within-subjects design was used in which the independent variables were familiarity of the face (celebrity, unfamiliar), orientation (upright, inverted), and distortion (normal, Thatcherised). A two-alternate forced choice (2AFC) task was used where participants were presented with two versions of the same face, one normal, and one distorted. Participants were required to indicate which of the two images looked 'odd'. Accuracy and speed of response were the dependent variables.

Participants

Twenty-eight undergraduate students from the Psychology Department participant pool (6 males, 22 females) acted as volunteers in the present experiment and were given course credit for their participation. Their ages ranged from 18 to 39 years (M = 20.4, sd = 4.9 years). All had normal, or corrected-to-normal vision.

Participants were randomly assigned to one of two groups. The only difference in the groups was in the identity of the inverted stimuli. This merely ensured adequate counterbalancing of the stimuli across orientation and across the participants.

Materials

Sixteen celebrity faces were selected, these came from the set used in Experiment Three, with additional images which had also been assessed and piloted for familiarity. All were full-face greyscale images, and were selected on the basis of their lack of distinguishing features such as facial hair, spectacles, etc. Sixteen unfamiliar faces were selected from those used in Experiment Three also. These were selected to be matched with the celebrities as far as possible on factors such as approximate age, hair colour and hair length. All images had already been prepared within Corel PhotoPaint to be matched on inter-pupil distance and to be mounted on a uniform white background. Images measured 6 x 8 cms (28.34 pixels /cm)

Two versions of each of these 32 faces (16 celebrity, 16 unfamiliar) were prepared. One was unaltered and appeared exactly as in Experiment Three. The other version was Thatcherised. Using Corel PhotoPaint, the eyes and the mouth of each face were extracted, rotated through 180 degrees, and then reinserted into the upright face. The edges of the pasted areas were 'airbrushed' to remove any harsh edges so that the Thatcherisation was not apparent merely by the appearance of grid-like lines on the face. Finally, both normal and Thatcherised faces were treated with a one-pixel Gaussian blur. This was to ensure that the distorted faces were not noticeably different in quality compared to the normal faces due to the cutting and pasting process. The Gaussian blur produces a slight hazy effect by averaging the pixel information outward using bell-shaped curves. The greater the number of pixels used in the averaging process, the greater the blurring effect. The use of a one-pixel blur produces a minimal effect, with no loss of featural information. Instead, the image appears slightly 'softened' and the edges blended.

Image displays were generated such that each face was paired with its Thatcherised counterpart. The two images were presented side by side, equidistant from the midline of the canvas (see Figure 6.3 for examples). Two displays were generated for each stimulus face, one with the distorted face on the left, and one with the distorted face on the right. Finally, inverted versions of each display were generated. This resulted in a total of 128 displays (16 celebrities, 16 unfamiliar, each represented by: distorted left, distorted right, as upright and inverted displays).

Trials were blocked by orientation, so that all participants saw the upright trials followed by the inverted trials. Within each block, trials were randomised and could contain familiar or unfamiliar faces. Images were presented in 8-bit greyscale, and were viewed via a 17" colour monitor with a screen resolution of 800 x 600 pixels.

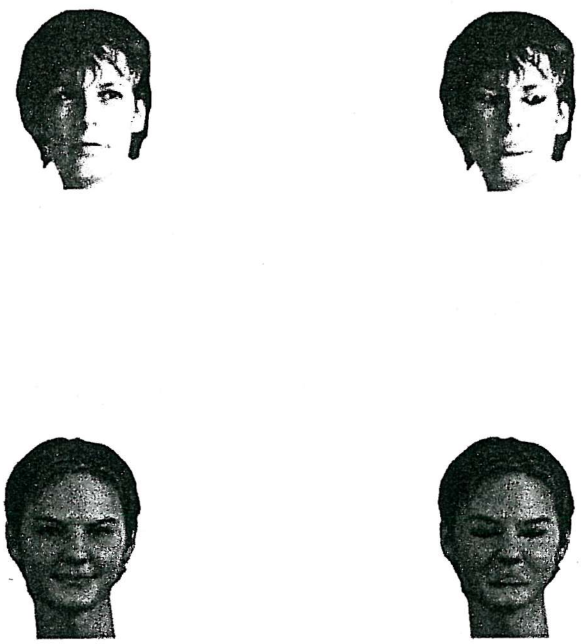

Figure 6.3: Example of Thatcherised (right) and normal (left) versions of Jamie Lee-Curtis (top) and an unfamiliar female face (bottom).

Procedure

Participants were seated approximately 50 cms from a computer monitor. They were presented first with a series of 10 practice trials. For each trial, the participant was presented with a fixation cross for 1000 msec, followed by a blank screen acting as an inter-stimulus interval (ISI) for 500 msec. Finally, a pair of faces was displayed side by side on the screen. One was always Thatcherised, and the task was to indicate, as quickly but as accurately as possible, which image looked 'odd'.

Participants responded by pressing either the right or the left hand button on a 2-button Cedrus response box. The images remained in view until a button press, and feedback was provided during this practise phase. Half of the trials during the practise phase consisted of upright images, and the remainder consisted of inverted images. As such, the participant had full preparation for the main test. Data for these practise trials were, however, omitted from all subsequent analyses and the stimuli were not repeated in the main test.

After completing the practise phase, the main test phase began. Participants saw a series of 64 upright trials, followed by 64 inverted trials. The format of each trial was identical to the practise phase, and the task was always to indicate which of the two images looked 'odd'. No feedback was provided during the main test.

Following the main test, participants were shown each of the celebrity faces, upright and intact, and asked to name them. In addition, participants saw each of the unfamiliar faces and were asked to confirm their unfamiliarity. This post-experimental check was performed to ensure that participants knew all the celebrity faces, and did not know the unfamiliar ones. The entire procedure took a total of 20 minutes to complete, and participants were then thanked and debriefed.

Results

All participants were able to name the celebrity faces and were unfamiliar with the non-celebrity faces. Therefore, data from all 28 participants was analysed. Accuracy and response times were recorded and are presented below.

Accuracy

The mean accuracy levels across participants and within each of the experimental conditions is summarised in Table 6.1 below. From this it is clear that the detection of distortion in an upright face was almost error-free, with participants performing at ceiling for both familiar and unfamiliar faces. However, when inverted, the detection of deception was substantially more difficult but the degree of difficulty appeared to be affected by the familiarity of the face.

Analysis by a repeated measures ANOVA on the proportion of accurate responses revealed a main effect of familiarity ($F(1, 27) = 12.9, p < .01$), where performance on familiar faces was more accurate than performance on unfamiliar faces. Unsurprisingly, there was also a main effect of stimulus orientation ($F(1, 27) = 30.2, p < .001$). Critically, the interaction between familiarity and orientation approached significance ($F(1, 27) = 3.96, p = .06$). The near-ceiling performance on the upright faces (both familiar and unfamiliar) compromises the validity of the interaction, and for this reason, speed of response becomes the primary dependent variable.

Table 6.1: *Mean proportion of errors (with standard deviations) during detection of distorted faces for upright and inverted familiar and unfamiliar presentations.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.02	(.03)	.11	(.08)	28
Unfamiliar	.04	(.04)	.16	(.12)	28

Response Time for Correct Decisions

The response time for correct decision was examined, having removed outliers of the mean plus or minus two standard deviations for each individual. The resultant data are summarised in Table 6.2 below. From this, it appears that the speed of response mirrors the accuracy data and there is no indication of any speed-accuracy trade-off. Performances overall were much faster for upright presentations, and were marginally faster for familiar faces. There was also some indication that inversion had a different effect depending on the familiarity of the face.

Analysis by means of a repeated measures ANOVA confirmed a large main effect of orientation ($F(1, 27) = 56.76, p < .001$) and a significant main effect of stimulus familiarity ($F(1, 27) = 7.60, p < .05$). Both main effects were qualified by the emergence of a significant interaction between familiarity and orientation ($F(1, 27) = 6.25, p < .05$). Post-hoc analysis of the interaction confirmed significant effects of familiarity when stimuli were inverted ($F(1, 27) = 7.00, p < .05$) and upright ($F(1, 27)$

= 6.98, $p < .05$), but examination of the difference scores (familiar – unfamiliar when (a) upright, and (b) inverted) showed that the effect of familiarity was significantly greater when stimuli were presented inverted ($t(27) = 2.49$, $p < .05$).

Table 6.2: *Mean response times in milliseconds, with standard deviations, showing time taken to detect distorted of faces for upright and inverted, familiar and unfamiliar presentations.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	720.28	(147.70)	2015.79	(1006.96)	28
Unfamiliar	745.76	(170.73)	2246.65	(1278.29)	28

Summary of Results

As expected, the detection of distortion was fast and accurate for upright faces, with participants performing close to ceiling. When inverted, distortion was recognised faster and with greater accuracy when faces were familiar. The distortion was still, however, detectable in unfamiliar inverted faces, as shown by the relatively accurate responses. However, performance was significantly slower in the inverted case when faces were unfamiliar. These results accord with those of the previous experiments in demonstrating a role for familiarity in the early visual perception of faces.

Discussion

The present results add weight to the suggestion that facial familiarity does play a role in the early visual processing of faces. The task used here was essentially a face classification task in which participants had to decide which of two images looked 'odd', or was a non-face (Thatcherised). The results suggested, as expected that this task was easy when the faces were presented upright, but was significantly harder, both in terms of accuracy and speed of performance, when faces were inverted. The results also suggested that facial familiarity played a role: the task was harder when faces were unfamiliar to the perceiver. Critically, however, it appeared that the role of familiarity was evident not when faces were upright, but when the task was made difficult by inversion. Under these conditions, the detection of 'oddness' was significantly better and significantly quicker if the face was familiar.

These results, combined with the results of the previous experiments within this thesis, confirm that facial familiarity has a surprising and significant role to play in both the early and the later visual processing of faces. At a face recognition level, the fact that familiarity influences processing should not, perhaps, surprise us. After all, recognition depends on the accessing of a mental representation, and familiarity rests exactly on the existence of and ease of processing such a representation. A difference between familiar and unfamiliar faces in recognition-based tasks might therefore reflect nothing more than the presence or absence of a pre-existing mental representation. In direct contrast, the fact that an effect of facial familiarity is evident in the present task – a face classification task – requires a substantially less superficial explanation. In a face classification task, where the task is essentially to report on whether a stimulus is intact or not, it is not explicitly necessary to access a mental representation. As such, no effect of facial familiarity might be expected. The fact that one emerges indicates that one cannot help but compute familiarity and that with this completed, processing at all levels can be facilitated. The question that then arises is how this facilitation might occur.

Three possibilities exist to account for the better processing of familiar inverted stimuli in the face classification task. First, it is possible that the familiarity of the stimulus allows the perceiver to switch to an alternate processing strategy with greater ease than if the stimulus were less familiar. This may result because the mental representation may contain a richer level of both featural and configural information than if the face was unfamiliar. If this was the case then it might be possible for the perceiver to switch relatively effortlessly to a featural processing strategy. As such, they could complete the 'oddness' task for inverted stimuli by processing each feature in series and then comparing each feature for orientation. Similarly, the perceiver may choose to focus on a particular component of a feature (for example, the bow of the lips). A mismatch in orientation would indicate oddness and so enable an accurate response.

Whilst this explanation is possible, it may be considered unlikely as the detection of oddness within a Thatcherised face really does depend on the comparison of features within the facial frame – that is, it depends on a relative processing strategy as exemplified by configural processing. To explain this further, the perceiver may detect the distortion within the face by focusing on the bow of the lips, and this is a *featural* strategy. However, to successfully determine the orientation of the lips in isolation from other features, or devoid of their facial context, is unlikely. Therefore the task is essentially configural in nature. Hence, it is difficult to rationalise how the Thatcher task can be performed with such levels of accuracy (87%, averaged across responses to familiar and unfamiliar inverted faces) when configural processing is impaired by inversion. An alternate process explanation is needed. In this respect, a second possibility is that, whilst inversion minimises the potential to process faces at a configural level, it might not remove this capacity completely. In this respect, one could argue that if familiar faces are processed with greater expertise, then this may be reflected by a greater reliance on configural processing. The result may be that even though inversion minimises the use of configural processing, it does not remove it and what is left means that familiar faces can still be processed with greater efficiency than unfamiliar faces. In the present task, the greater efficiency when processing familiar faces is not evident when stimuli are upright because performance is quick and accurate for all stimuli. However, the greater efficiency when processing familiar faces becomes evident when stimuli are inverted because the system is taxed and so a disparity becomes apparent.

This latter explanation is possible. However, it may also be considered unlikely given the weight of evidence to suggest that inversion removes the ability to process stimuli at a configural level. It might also be questioned as to why configural processing is relatively preserved when stimuli are familiar but not when stimuli are unfamiliar.

A third possibility, though, is located as the level of the mental representation rather than at the level of the processing strategy used. It is possible that the very familiarity of the stimuli results in a robust representation along the lines discussed by Tong and Nakayama (1999). As such, the mental representation is based on a rich input of information covering a range of views, lighting conditions, expressions, and changes in appearance across age and general grooming (i.e., hairstyle changes). The resultant mental representation might be considered more robust because the mental representation created might be more representative, or more prototypical (Cabeza, Bruce, Kato, & Oda, 1999) of the individual than if it were based on only a few perhaps idiosyncratic presentations. According to Tong and Nakayama (1999), the impact that a robust representation has on processing is that the perceiver remains able to process the stimulus despite changes in viewing conditions that usually disrupt performance. In other words, a robust representation affords the perceiver some defence against difficult viewing conditions.

This latter explanation is in accordance with the explanation offered for the results of the previous experiments. As such, the concept of robust representations is gaining support from converging lines of evidence within the thesis and offers a parsimonious way of explaining all the results hitherto. Having said this, there are two reasons for accepting the present results with a degree of caution. The first rests on the possibility of an artefact within the data. The second arises as a consequence of the task used in the current experiment.

Examination of the data revealed near-ceiling levels of performance, both in terms of accuracy and speed of correct responses, when stimuli were upright. That is, the task of detecting the 'odd' or Thatcherised face was very easy when stimuli were presented within their normal upright orientation. The fact that performance was so good when upright may, consequently, have negated the capacity to demonstrate any difference in performance across familiarity. Quite simply, performance was already so good with upright unfamiliar faces that it could not get significantly better when viewing upright familiar ones. If this is the case, then the apparent interaction of familiarity and inversion found here may merely be the result of the near-ceiling performance when stimuli were presented upright. It is proposed that increasing task difficulty (on upright as well as inverted stimuli) would provide a more sensitive tool with which to investigate familiarity effects. Therefore, a replication of the present study using a methodology designed to lower overall performance levels would be required to test this possibility.

The second potential problem within the current experiment lies with the use of a two-alternate forced choice paradigm in which participants are presented with one normal and one Thatcherised face. It is possible, therefore, to complete the 'oddness' task in three ways. First, one could make an absolute decision as to which of the two faces was 'odd'. Second, one could make a relative decision as to which of the two faces was 'more odd'. Finally, one could make an absolute decision as to which of the faces was 'normal' and then respond to the other one by default as 'odd'. As such, it is unclear from the present design whether the participants were actually engaging in a configural-level task, or whether the task could be completed using a more comparison-based strategy. In this sense, whilst not a same/different task, the current task *could* be completed using a featural strategy which loses little when stimuli are inverted. The fact that the task could still be performed with what are quite high levels of accuracy across both familiar and unfamiliar stimuli might be taken in support of this possibility.

The use of a design in which a 'normal/odd' response is made to single images will certainly address the second point of relative responding, and may also address the first point of ceiling levels of performance. This is the purpose of the following experiment.

CHAPTER 7

Confirming the Effect of Familiarity using Classification of Thatcherised and Normal Faces

Abstract

The predictions of the previous chapter were re-examined in the present experiment by using an odd/normal face classification task rather than a 2AFC (which is 'odd'?) task. Thatcherised faces were used as a basis of disrupting face-ness, and stimuli were presented both upright and inverted. It was predicted that the detection of distortion would be easy when faces were upright but would be made difficult by inversion. On the basis of the previous results, it was anticipated that under conditions of inversion, facial familiarity would help in completion of the task. Both accuracy and speed of response were recorded and the data supported all predictions. The results are discussed with reference to the notion of robust representations, and with the notion of automaticity of processing facial familiarity.

Confirming the Effect of Familiarity using Classification of Thatcherised and Normal Faces

The rationale for the present experiment is identical to that for the previous experiment. Essentially, the purpose is to determine whether the familiarity of the face can affect very early visual processing, as required for a face-classification task. The Thatcher Illusion again provides a way of disrupting facial stimuli in order to create a face/non-face distinction.

The previous study saw this effect used to compare the processing of familiar (celebrity) and unfamiliar faces. Participants viewed pairs of stimuli – one normal, one Thatcherised – and were required to indicate which of the pair looked 'odd'. The results suggested that when upright, the task was easy and could be performed quickly and accurate regardless of the familiarity of the face. However, when inverted, the task became more difficult, and a discrepancy emerged between the processing of the familiar and unfamiliar stimuli. It appeared that knowing a face helped participants perform the classification task when that task was made difficult by inversion.

These results are intriguing because they offer evidence to suggest that familiarity can help participants do a task even when familiarity is not explicitly required. Moreover, with the task of face classification assumed to precede the determination of familiarity, it is surprising to see the later decision affect the earlier one. However, the use of a 2AFC task may have had the effect of simplifying the decision process so that distortion was assessed relative to the normal face within the pair. Two consequences may have arisen. First, the possibility of completing the task by using a comparison strategy may have made the task much easier than intended, hence creating ceiling level of performance when stimuli were upright. Second, the possibility of using a comparison strategy raises the question of whether distortion *could* be detected when a comparison face is absent, and whether this detection might also show sensitivity to facial familiarity.

To address the issue of whether a 2AFC task in the previous experiment made the detection of distortion too easy, the task will be replicated here. Participants are however presented with single images and are required to indicate whether each is normal or distorted. This task will ensure that participants are responding to an individual exemplar, rather than selecting a distorted face through a comparison strategy. It is hoped that this change in task demands will serve to increase the difficulty of the task so that an explanation for the previous results based on both ceiling performance and processing strategy can be controlled for. As in the previous experiment, it is predicted that detection of distortion will be better when stimuli are upright than when inverted, and that general processing of faces will be aided by familiarity. The issue of whether familiarity affects the detection of distortion under difficult conditions of inversion will be assessed through examination of the interaction effect for both accuracy and speed of response.

Method

Design

A 2 x 2 x 2 within-subjects design was used, as in the previous experiment. The independent variables were facial familiarity (celebrity, unfamiliar), orientation (upright, inverted) and 'appearance' (Thatcherised, normal). A yes/no decision task was adopted for this study, where participants were required to decide whether the face was normal (yes) or not (no). Accuracy and speed of response were recorded, and measures of sensitivity (A') and bias (B''_D) were calculated.

Participants

Fourteen undergraduate students (2 males, 12 females) from the Department of Psychology volunteered to take part in the present study and were given course credit for their participation. The participants' ages ranged between 18 and 31 years (M = 21 years, sd = 3.6 years) and all had normal, or corrected-to-normal vision. None of the participants had taken part in any of the previous experiments.

Materials

The stimuli used in the present experiment were identical to those used in Experiment Five (see Appendix F). These comprised 16 celebrity faces and 16 unfamiliar faces selected to be matched to the celebrities on approximate age, hair colour, and hair length. As in the previous experiment, the 32 stimuli were prepared as normal and Thatcherised versions and were presented both upright and inverted. The only difference between the present experiment and the previous experiment was that stimuli were presented one at a time rather than in normal-distorted pairs. As such, there were 128 experimental trials consisting of 32 faces presented as upright and inverted, normal and Thatcherised images.

The stimuli were presented within a SuperLab environment. All images measured 8 cm x 6 cm and were presented centrally on the 17" screen. All other aspects of the stimulus preparation and presentation were identical to Experiment Five.

Procedure

The procedure was identical to that used in Experiment Five. Participants completed a practise phase of 10 trials, during which a fixation cross was shown for 1000 msec, followed by a blank screen as an inter-stimulus interval, for 500 msec. Finally a single image was shown and participants were required to indicate whether the image was normal or 'odd' (Thatcherised). Participants responded by pressing one of two labelled buttons on a 2-button Cedrus response pad, and the image remained on screen until the participant responded. Accuracy and speed of response were recorded and feedback was provided on performance during the practise phase. All data and stimuli from the practise phase were, however, not included in the main experiment.

Following completion of the practise trials, the main test phase began. Stimulus presentation was blocked by orientation such that all participants saw the upright images first. Within each block, the familiar and unfamiliar stimuli were presented in a random order. Participants responded as above, but no feedback was provided.

Finally, participants completed a post-experimental check to ensure that they knew each of the celebrity faces, and did not know the unfamiliar faces. Participants were included in the analyses if they could name all celebrities and were unfamiliar with the non-celebrities.

Results

All 14 participants performed adequately on the post-experimental check, however, data from one participant were dropped due to incorrect performance for all bar one unfamiliar inverted 'odd' face. Data from the remaining 13 participants were analysed. With the focus of the present experiment on the detection of distortion when familiar and unfamiliar faces are upright and inverted, the accuracy data will be of primary concern. However, the design of the present task provides the potential for participants to respond with a bias (i.e., to say 'normal' or 'odd' regardless of the stimulus appearance). With this in mind, sensitivity of discrimination, together with a measure of response bias, will be examined. Finally, the speed of correct responses provides a further measure to amplify the picture provided through examination of accuracy. Each of these analyses will be taken in turn.

Accuracy

The mean accuracy levels across participants and within each experimental condition are summarised in Table 7.1 below. From this, it is clear that the task of face classification as normal or distorted was highly accurate when faces were upright. Performance was, however, impaired when faces were inverted, as expected. Under these conditions, it became difficult to determine the relative orientation of the features within the facial frame. It was also apparent that a discrepancy emerged between performance on familiar faces and performance on unfamiliar faces when those faces were inverted.

Table 7.1: *Mean proportion of errors (with standard deviations) for detection of distorted faces for upright and inverted familiar and unfamiliar presentations.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.03	(.05)	.19	(.09)	13
Unfamiliar	.04	(.03)	.33	(.11)	13

Analysis by means of a repeated measures ANOVA revealed a main effect of orientation ($F(1, 12) = 57.06, p < .001$) and a main effect of stimulus familiarity ($F(1, 12) = 40.18, p < .001$). Both effects were, however, qualified by the expected interaction ($F(1, 12) = 19.20, p < .01$). Examination by means of post-hoc contrasts confirmed that there was no effect of familiarity on detection of distortion when stimuli were upright ($F(1, 12) = .49, p > .05$). However, a clear and significant effect of familiarity emerged when the task was made difficult by stimulus inversion ($F(1, 12) = 31.74, p < .001$).

Sensitivity of Responding

Analysis of accuracy allows examination of the success of saying 'distorted' to distorted images and 'normal' to normal images within each of the experimental conditions. It does not, however, separate out responses to the distorted and normal stimuli in the same way that a measure of discrimination does. Calculation of A' allows examination of the balance of responding across stimuli in this way. A' values were calculated for each individual within each of the experimental conditions. These are summarised in Table 7.2 below and the pattern echoes that provided by the accuracy data.

Table 7.2: *Mean A' values (with standard deviations) for upright and inverted, familiar and unfamiliar presentations.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
Familiar	.96	(.07)	.88	(.07)	13
Unfamiliar	.97	(.02)	.76	(.14)	13

Analysis by means of a repeated measures ANOVA also confirms the pattern provided by the accuracy data. A main effect of orientation was evident ($F(1, 12) = 22.29, p < .001$), with discrimination being better when stimuli were upright. A main effect of familiarity was also evident ($F(1, 12) = 6.44, p < .05$) with discrimination being better when faces were familiar. Finally, a significant interaction emerged ($F(1, familiarity ($F(1, 12) = .46, p > .05$). However, a significant interaction emerged ($F(1, 12) = 5.65, p < .05$). Post-hoc analyses identified that inversion affected the direction of the bias significantly when faces were unfamiliar (($F(1, 12) = 6.2, p < .05$) but did not affect the extent of bias when faces were familiar ($F(1, 12) = 4.4, p > .05$).

Response Speed for Correct Decisions

Finally, the speed of correct responses was examined in order to investigate the influence of familiarity on detection of distortion further. Responses for 'odd' and 'normal' decisions were separated and are summarised across all participants within the experimental conditions in Table 7.3 below. From the table, it appeared that when upright, decision times were faster for odd faces over normal ones, and for unfamiliar over familiar ones. When inverted, however, decision times showed the opposite pattern. Decisions were made faster when faces were normal rather than odd, and when faces were familiar rather than unfamiliar.

Table 7.3: *Mean response times (with standard deviations) for detecting distortion in Thatcherised faces when familiar and unfamiliar and when presented upright and inverted.*

	Upright		Inverted		N
	Mean	(SD)	Mean	(SD)	
NORMAL TRIALS					
Familiar	881	(282)	1283	(557)	13
Unfamiliar	860	(268)	1393	(670)	13
ODD TRIALS					
Familiar	755	(187)	1464	(542)	13
Unfamiliar	715	(183)	1660	(564)	13

Analysis by means of a 3-way repeated measures ANOVA, with familiarity, orientation, and appearance, as factors, confirmed a main effect of orientation ($F(1, 12) = 28.79$ $p < .001$), but not of familiarity ($F(1, 12) = 1.89, p > .05$) or appearance ($F(1, 12) = .60, p > .05$). In addition, a significant interaction between familiarity and orientation emerged ($F(1, 12) = 5.06, p < .05$), and between orientation and

12) = 15.13, $p < .01$). As anticipated, post-hoc contrasts confirm this interaction to be due to a significant difference across familiarity when stimuli were inverted ($F(1, 12) = 14.27, p < .01$) but not when upright ($F(1, 12) = .59, p > .05$).

Response Bias

Within the present design, it is possible for participants to reveal a biased pattern of responding. This would be evident if a participant utilised a cautious approach, for example, and responded 'normal' more often than not, irrespective of the actual appearance of the stimuli. In order to investigate this possibility, a measure of response bias (B''_D) was calculated. This is summarised across participants and within each experimental condition in Figure 7.1 below, from which it appeared that there was a bias to say 'normal' when faces were inverted. When faces were upright, there was much less bias evident. However, there was a tendency to say 'distorted' when the upright stimuli were unfamiliar.

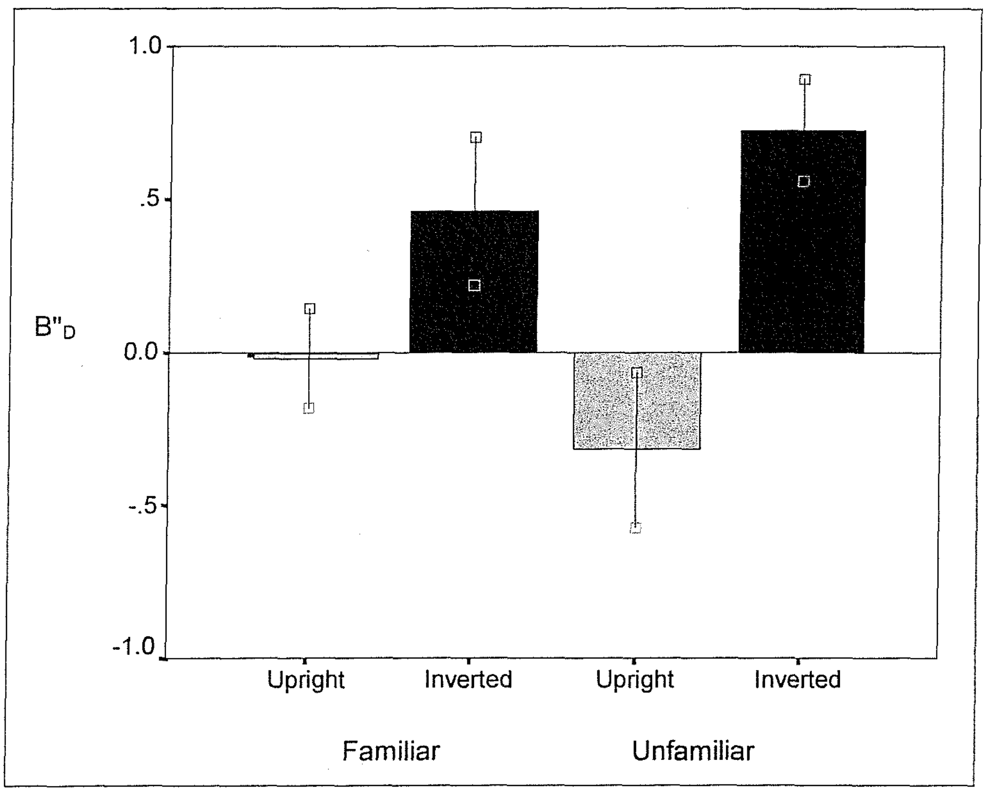

Figure 7.1: Mean B''_D (with standard error bars) showing the bias to see inverted faces as 'normal'.

These data were examined by means of a repeated measures ANOVA. The results confirmed a main effect of orientation ($F(1, 12) = 160.16, p < .001$), with a tendency to say 'normal' for inverted faces. There was no main effect of facial appearance ($F(1, 12) = 8.29, p > .05$) but no other 2-way, or 3-way interactions reached significance. Post-hoc examination of the interaction of familiarity and orientation suggested that although familiarity had some influence on detection of distortion when faces were upright ($F(2, 11) = 6.19, p < .05$), this was more marked when faces were inverted ($F(2, 11) = 20.62, p < .001$) such that detection of distortion was aided by familiarity when processing was made difficult by inversion. These results echoed the pattern revealed by the accuracy and sensitivity data.

Summary of Results

The data from all analyses were consistent and suggested that when upright, performance was accurate, quick, and uniform across familiarity. However, when inverted, performance differences emerged across familiarity to suggest that the detection of distortion was less impaired by inversion when faces were familiar than when unfamiliar. Analysis of the bias measure revealed a bias to see all inverted faces as normal, thus showing the relative difficulty in detecting distortion in the inverted face. This tendency was especially marked for unfamiliar faces. In contrast, when upright, familiar faces yielded no bias in responding at all but there was a tendency to call upright unfamiliar faces 'odd' even if they had not been transformed. These results are discussed with respect to the notion that facial familiarity can facilitate performance on a task that does not explicitly demand an identity-level decision.

Discussion

The results presented in this experiment are consistent with one another, and are consistent with those presented in the previous chapter. Whether tested using a 2AFC task, or a distorted/normal task, performance was influenced by the familiarity of the stimuli when task difficulty was increased by inversion. More specifically, under conditions in which the detection of distortion was impaired, performance was nevertheless facilitated if the face was familiar. This pattern was evident when examining both accuracy and speed of correct responses across tasks, and when examining sensitivity of responding in the present experiment. As such, the influence of familiarity in a face-classification task using Thatcherised faces can be considered a robust finding.

The influence of familiarity in these face classification tasks is somewhat surprising if one takes the view that face classification precedes a judgement of facial familiarity. According to this logic, it is plausible that familiarity might affect judgements that are made after the familiarity decision is reached, however, it is difficult to see how familiarity can affect judgements made before this point. The facts of the matter go against this logic as it is clear that familiarity does indeed have an influence on the early visual processing of facial images. The task of the present discussion is to explain how this situation can come about.

One explanation may be that the previous assumption that classification precedes determination of familiarity is faulty. It might, perhaps, be the case that both classification and familiarity judgements arise as a consequence of structural encoding. This is an appealing notion and would certainly allow explanation of the current data, as decisions about both familiarity and face-ness would be taken at the same point in the processing sequence. This notion is, however, countered by the clinical dissociation of the classification and familiarity decisions: The condition of metamorphosia affects the ability to perceive a stimulus *as a face* while the condition of prosopagnosia affects the ability to determine the *familiarity* of a stimulus whilst classification as a face is unimpaired.

The previous chapter did, however, discuss several possibilities to account for the involvement of familiarity in a difficult face classification task. First, the possibility that familiarity allowed the perceiver to switch processing strategies with ease was considered. Such a processing switch may allow the perceiver to process an inverted face with greater ease through the use of a featural strategy and this strategy may be sufficient to perform the face classification task. Whilst possible, this option was considered unlikely given the need to process the features of the face in relation to one another in order to determine whether the image had been Thatcherised.

The second possibility was that familiar faces could be processed with greater expertise, as exemplified by greater reliance on configural processing, compared to unfamiliar faces, and that inversion minimised but did not remove the capacity to process the face configurally. There are two suppositions embedded within this position: first that familiar exemplars are processed with greater reliance on configural processing, and second, that inversion does not obliterate the potential to process faces in this way. In terms of the first proposal, this notion is not without precedent and certainly there is evidence to support the suggestion that subsets of faces with which a perceiver has more familiarity can show a processing advantage in a task that relies on configural processing (see Rhodes *et al.*, 1989). However, it is somewhat difficult to consider the notion of *degrees of configural processing.* One might, perhaps consider that configural processing can be used with more or less efficiency, as exemplified perhaps through speed of processing. However, it is unclear whether configural processing can be implemented in degrees as opposed to being an all-or-nothing approach. In terms of the second proposal, the notion that some capacity to process faces configurally might be retained even when stimuli are inverted is not without precedent. Again, however, by the same logic as above, it might be considered that what is affected is the speed of extracting and utilising the configural information rather than the degree of configural information that is used. In both cases, if configural information is used by degrees, either because of an influence of familiarity or because of an influence of inversion, then an effect on performance might be expected when examining response times rather than accuracy. The fact that the influence of familiarity is evident across both measures negates this argument.

The third possibility to be examined in the previous chapter was that the processing of familiar faces was protected from the negative impact of inversion because the perceiver possessed robust representations for these faces. Based on a wealth of rich visual experience with the face, this would allow the perceiver to remain able to process a familiar stimulus despite a disruption such as inversion. As such, the perceiver remained able to process the inverted familiar faces, albeit not quite so accurately or quickly, and so remained able to detect distortion within the familiar stimuli.

A fourth possibility is, however, apparent and arises through consideration of the present results. The purpose of this experiment over the last experiment was to increase the task difficulty and avoid ceiling level performance. This was addressed through the use of a normal/odd decision here, rather than a 2AFC (which face is 'odd') decision previously. The additional benefit of such a change in task was to eradicate the possibility that responding in the previous experiment could have been based on a comparative strategy rather than an absolute strategy. Comparison of the results across studies suggests that while accuracy remained uniformly good and near-ceiling when faces were upright, the accuracy of responding was successfully lowered in the present task when faces were inverted. This was more evident for the unfamiliar faces (66% vs. 83%) than for the familiar faces (80% vs. 89%). Moreover, examination of speed of response for correct decisions again verifies that the task appeared more difficult in the present experiment than in the previous 2AFC task.

One implication of this increased difficulty is that, although the task of detecting distortion can still be done, it is done more slowly now than before. It is possible, now, that the automatic processing of facial familiarity might have reached completion before the face classification task had been resolved. As such, the processing of familiar faces would benefit from the increased amount of information available, whilst the processing of unfamiliar faces would remain untouched.

This explanation would account for the effect of familiarity when detecting distortion in inverted faces (when the classification task is difficult and slow) and the absence of such an effect when detecting distortion in upright faces (when the classification task is easy and completed with speed). This explanation is also in line with the findings of Bruce (1986) and Baudouin *et al.* (2000). Bruce showed that familiarity could help gender decisions under ambiguity – when gender decisions were difficult, knowing a face helped resolve the decision. Similarly, Baudouin *et al.* found that when viewing conditions were made more difficult, for example by decreasing presentation time or obscuring the mouth region, expression decisions were faster for familiar than for unfamiliar faces. In these cases, as in the present study, the facilitating effect of familiarity is revealed when task difficulty is increased and so the task is slowed. In all cases then, it is possible that the familiarity decision has time to be resolved before the primary task and that this familiarity decision then has a window of opportunity to facilitate performance on this primary task.

It should be noted at this stage that the concept of robust representations for familiar faces, and the notion of automatic processing of familiarity which feeds back to early visual processing stages, are not wholly incompatible with one another. It is perfectly plausible, for example, that it is the very robustness of the representation for familiar faces that allows familiarity to be determined so quickly and within the time frame taken for the difficult primary task of face classification. The only factor left to consider is whether it is plausible to think of facial familiarity as being an unstoppable and automatic goal of face processing, even though it is not explicitly required by the primary task. Intuitively, it would seem sensible that the goal of a face processing system is not simply to report on whether a face is intact or not, but is to report on *who* you are looking at. Empirically as well, the present demonstration of an influence of familiarity on an early visual task such as classification implies that one cannot help but process familiarity and perhaps, indeed, identity, when presented with a facial stimulus. Given this, the potential for familiarity to influence a whole host of face-related decisions would thus appear to exist.

Conclusions

The present study has confirmed the results of the previous study in demonstrating an influence of familiarity on a face classification task. More specifically, it has shown that familiarity can aid processing of inverted faces such that distortion is more rapidly detected when a face is known than when a face is unknown. One explanation for this influence is that familiarity permits the formation of a robust representation that can withstand disruption to processed caused by, for example, stimulus inversion. An alternative explanation, and one that is not wholly incompatible with the notion of robust representations, is that the difficulty of the face classification task when stimuli are inverted means that the primary task of classification cannot be completed before the (here) secondary and automatic processing of familiarity is achieved. Thus, facial familiarity has the capacity to facilitate processing in an otherwise difficult early visual task.

CHAPTER 8

Converging Evidence for a role of Familiarity within a Face Classification Task using Scrambled and Intact Stimuli

Abstract

The previous two experiments have used the Thatcher Illusion as a way of distorting the facial image. In both a 2AFC (which is odd) task, and an odd/normal task, facial familiarity was shown to facilitate performance when task difficulty was increased by inversion. Thatcherisation is, however, a very subtle form of distortion, generating an image that is still face-like but possesses an 'odd' appearance. The purpose of the present experiment is to see whether the role of familiarity generalises to a task where the face is distorted to a greater degree through feature scrambling. Normal, moderately and highly scrambled stimuli of familiar and unfamiliar faces, were presented in a face classification (face/non-face) task. Both accuracy measures (proportion of errors, sensitivity) and speed of response indicated that familiarity did influence performance when stimuli were inverted, and when scrambling was moderate. In this case, familiarity hindered performance. These results are discussed with respect to the automaticity of processing facial identity, and the influence that such an identity decision can have on a classification task as opposed to a normal/odd task.

Converging Evidence for a role of Familiarity within a Face Classification Task using Scrambled and Intact Stimuli

In two Thatcher Illusion tasks, one a 2AFC task (Experiment Five) and the other, a typical/odd decision task (Experiment Six), an advantage was shown for familiar faces, where classification of faces as 'typical' or 'odd' (Thatcherised) was more efficient if participants knew the face. The Thatcher Illusion was used to show that distortion due to a (relatively) subtle change in configuration was detected more easily in familiar faces. This difference in processing has not been shown previously, and suggests that the processing of familiarity may be automatic and may mediate the negative influence of stimulus inversion. The following task seeks to provide converging evidence for this novel finding through using a face/non-face decision rather than a typical/odd decision. The face stimuli are either normal or contain gross changes in configuration brought about by feature scrambling. It is anticipated that the results will again show an influence of facial familiarity in completing the face classification task despite the change in task decision and the alteration in stimulus preparation.

Previous studies have used scrambled faces to investigate face classification; which involves the basic level categorisation of stimuli as a 'face' or a 'non-face' (Sergent, 1984). This level of decision is assumed, within Bruce and Young's (1986) hierarchical framework, to occur as part of the structural encoding process, and its existence is supported by the clinical dissociation of the conditions *metamorphosis* and *prosopagnosia*. However, Bruce and Young did not specify how the face classification decision was accomplished because they suggested that to do so would infer the need for a face-specific analysis rather than a structural analysis that could fit any object. Whilst it is likely that there is indeed a face-specific analysis in operation, this has not been demonstrated categorically as yet, and the accumulation of data suggesting similarities in the way that faces and objects are processed casts doubt on there being such a simple distinction. Moreover, Bruce and Young consider it unclear as to whether a face 'switch' exists and is activated when an input is classified as a face, or whether selective analysers exist which tune themselves to pick up the input and implicitly allow classification as a face. Finally, Bruce and Young remain unconvinced of the utility of proposing such a face classification stage as, at present, it is unclear whether such a stage would precede other processing stages or whether it would be computed in parallel with them.

Therefore, although there is a general acknowledgement that face classification is probably accomplished at an early stage of processing, there is no clear definition of how, or when, this occurs. Notwithstanding this shortcoming, scrambled stimuli have been used to great effect to investigate differences in early facial processing. For example, Davidoff and Donnelly (1990) and Tanaka and Farah (1993), using Mac-a-Mug images, have shown what is known as a complete-part advantage (CPA). This is seen when a feature (probe) is recognised more efficiently when presented in the context of a complete face than when presented alone. This effect can be demonstrated when faces are presented intact, but not when they are presented as scrambled stimuli. This was taken to suggest that intact and scrambled stimuli are processed in different ways and put different emphasis on the 'whole'. More specifically, it was argued that upright intact faces were processed configurally, through which attention to the relationship between parts was informative, whilst scrambled faces were processed featurally, through which the features other than the probe gave no help in processing the probe.

Subsequent work by Donnelly *et al.* (1994) used inversion to confirm that intact and scrambled faces were likely to be processed in different ways. The classification of intact faces was significantly affected by stimulus inversion, as might be expected if processing relied on configural and parallel strategies. In contrast, the classification of scrambled faces was affected only minimally by stimulus inversion, indicating the use of a strategy which is not orientation specific, such as a serial and self-terminating featural strategy.

The question for this thesis is whether any evidence exists on the effect of scrambling across different levels of facial familiarity. To this end, only one published study has been found that has a bearing on this issue. Embedded within Collishaw and Hole's (2000) paper is the use of a scrambling manipulation across two experiments; one with familiar faces and one with unfamiliar faces. The data suggested that both accuracy (in terms of d'), and speed of correct responses, showed no difference in the magnitude of the scrambling effect whether faces were familiar (celebrities – Experiment 1), or unfamiliar (seen only once – Experiment 2). Indeed, in consideration of the accuracy data, the authors comment:

> … there does not appear to be any difference in the relative impact of blurring, scrambling, and inversion [across familiarity] when applied alone (p902).

and in consideration of the speed of response, the authors note that:

> … all manipulations have a greater effect on RTs for familiar face recognition, but the patterns of RTs across the different manipulations do not differ for the two tasks [across familiarity] (p903).

In evaluating these data, two things are worth noting. First, the task being performed was not a face classification task in which participants made a face/non-face decision to an image. Instead, participants made a recognition-level decision. It is possible that, given a change in task demands from recognition to classification, quite a different pattern of results may have been obtained. Second, the recognition-level task being performed differed somewhat across the familiar and unfamiliar stimuli. Experiment 1 required a speeded familiarity judgement to celebrity faces, and required the involvement of a long-term mental representation. In contrast, experiment 2 required an old/new judgement to previously unfamiliar faces, and thus required the involvement of a short-term mental representation. As suggested elsewhere in this thesis, the slight difference in task demands across the two experiments makes direct comparison of the data somewhat problematic.

Given this discussion, the issue of whether facial familiarity affects a face classification decision remains unresolved. The purpose of the present study is to address exactly this issue.

The Present Study

A face/non-face classification task will be used, where the stimuli consist of intact and scrambled, familiar and unfamiliar faces. Unlike previous studies using scrambled stimuli, the present study uses photographic images in which the features were electronically cut out and pasted back into the face (in a changed order) and the edges of the pasted sections were airbrushed to remove any lines or obvious differences in tone. This gave the images a natural appearance and had the advantage of not forcing the viewer into any particular processing strategy. This contrasts with the scrambling process used in the majority of previous studies in which photographic images had been cut into horizontal strips which were then reassembled in a rearranged order. This may have had the effect of forcing a strip-by-strip analysis of the face, hence a serial feature-by-feature search.

The use of a face/non-face classification task here follows the principles of the studies presented in the previous two chapters, such that identification of the face is not the focus of the task. As such, familiarity might plausibly be expected not to influence the task at all. However, if the results of the previous Thatcher tasks are reliable, then an influence of facial familiarity at this early visual processing stage, is to be expected. Accuracy and speed of response will be examined in order to test the robustness of these predictions.

Method

Design

A 2 x 2 x 3 within-subjects design was used to investigate the effects of familiarity, orientation and scrambling on face processing. Familiar (celebrity) and unfamiliar faces were presented sequentially either as intact or scrambled images in both upright and inverted orientations. The task was a face/non-face decision task. Performance was measured by recording accuracy and speed of response. In addition, response sensitivity and response bias were calculated and examined.

Participants

Thirty undergraduate and postgraduate volunteers (8 males, 22 females) participated in the current experiment. Ages ranged between 19 and 25 years (m = 23.9 years, sd = 5.8 years), and none of the participants had taken part in any of the previous studies. All participants had normal, or corrected-to-normal vision.

Materials

A total of 16 stimulus faces were used in the current experiment. Half of these were familiar (celebrity) faces, and these images were piloted before inclusion in the test to ensure that they were recognisable. The remaining faces were unfamiliar, and were drawn from the Stirling PICS database. An equal number of male and female faces was used in each set, and all were Caucasian. Faces were selected to have few distinguishing features; that is, they had no facial hair, or scars, and did not wear spectacles.

Faces were cut out and mounted on a uniform white background using Corel PhotoPaint. They were then scaled and standardised on the basis of inter-pupil distance. Converting the images to greyscale ensured the removal of any slight variations in skin colouring.

Six versions of each face were generated. One was undistorted. The remaining five versions displayed varying arrangements of features (eyes, nose, mouth) to create either a moderately scrambled face (where one feature was in the correct place, n = 3) or a highly scrambled face (where all features were out of place, n = 2). Examples of the stimuli are shown in Figure 8.1 below.

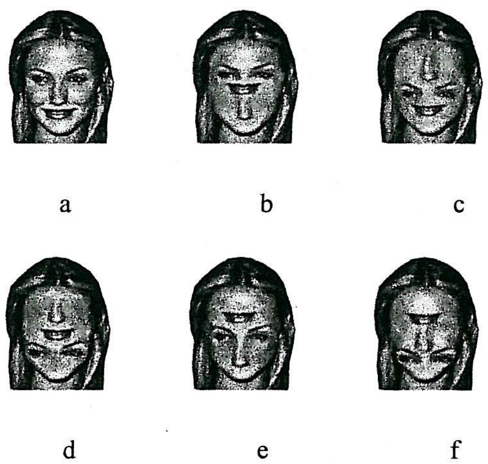

Figure 8.1: Examples of non-scrambled (normal) and scrambled versions of a celebrity face (Cameron Diaz). Image a is normal. Images b, c and f are moderately scrambled. Images d and e are highly scrambled.

Scrambling was achieved using Corel PhotoPaint. The eyes, nose and mouth of each face were extracted from the facial outline and re-ordered to generate the five scrambled versions for each face. Feature edges were disguised by using an airbrush technique to blur any variation in pixel colouring at the junction of each feature section. In addition, a one-pixel Gaussian filter was placed over all images (scrambled and normal), to reduce the effects of differences in shading of the repositioning features, and to ensure that scrambled and normal versions of the face possessed equivalent picture quality.

Presentation of the stimuli was organised to ensure an equal number of 'intact' and 'scrambled' trials. Thus, each face was seen 24 times (intact x 12, moderately scrambled x 6 (3 versions twice each) and highly scrambled x 6 (2 versions three times each). This yielded 384 trials. In addition, trials were shown upright and then inverted and were blocked by orientation to give 768 trials in total.

Stimuli were presented and responses were recorded within a SuperLab testing environment. Stimuli were shown on a 17" colour monitor, with a resolution of 1024 x 768 pixels. Images measured approximately 6 x 8 cms (28.34 pixels/cm) and subtended a visual angle of 6°. Responses were made using a six-button Cedrus Response Box, and care was taken to ensure that allocation of response keys was counterbalanced across participants.

Procedure

Participants were seated approximately 50 cms from the computer screen. They were presented with a series of stimuli, shown one at a time, and their task was to decide whether each stimulus was 'normal' or 'scrambled'. Decisions as to the amount of scrambling (moderate, high) were not required.

Each trial consisted of the presentation of a fixation cross for 500 msec, and then the stimulus face which appeared in the centre of the visual field and remained in view until a response had been made. A 500 msec blank screen separated the response from the next trial. Participants were required to make their response by using the outer buttons of a 6-button Cedrus response box. These were marked 'N' (normal) and 'S' (scrambled) and button labelling was counterbalanced across participants.

For all participants, the presentation of stimuli was blocked by orientation, and all participants saw the upright trials first. Finally, a post-experimental check was performed by all participants to ensure that they were familiar with all celebrity faces. This check consisted of the presentation of each face in turn. Participants were required to try to name the celebrities, and to confirm that the unfamiliar faces were indeed unknown.

Results

All participants were able to name the celebrity faces, and so data from all 30 were analysed. Data were analysed to examine the influence of familiarity on the face classification task for intact, moderately scrambled, and highly scrambled faces. Accuracy and speed of correct decisions are analysed, as are measures of sensitivity (A') and response bias (B''_D). These analyses are presented in turn below.

Accuracy

Mean accuracy across familiarity, orientation, and degree of scrambling, is summarised in Table 8.1 below. From this it is clear that performance when stimuli are upright is very good indeed, with little variation across conditions. When stimuli are inverted, the face classification task is more difficult, especially when faces are moderately scrambled, and when faces are familiar.

Analysis of the proportion of accurate responses by means of a 2 x 2 x 3 repeated measures ANOVA, with familiarity, orientation, and scrambling as factors, showed significant main effects of all factors. These were moderated by significant two way interactions (familiarity x orientation, $F(1, 29) = 11.31, p < .01$; familiarity x scrambling, $F(1, 29) = 22.10, p < .001$; orientation x scrambling, $F(1, 29) = 8.88, p < .01$). Finally, a significant three way interaction moderated all effects ($F(1, 29) = 8.88, p < .01$).

Table 8.1: *Mean proportion of errors (with standard deviations) showing overall accurate recognition of scrambled faces (n = 30).*

	Upright			Inverted		
Scrambling	Intact M (SD)	Moderate M (SD)	High M (SD)	Intact M (SD)	Moderate M (SD)	High M (SD)
Familiar	.02 (.03)	.04 (.05)	.03 (.03)	.02 (.03)	.11 (.10)	.05 (.05)
Unfamiliar	.02 (.03)	.03 (.05)	.04 (.05)	.04 (.04)	.05 (.08)	.04 (.05)

Further examination of the three-way interaction was performed, using two separate two-way ANOVAs to examine the effects of familiarity and scrambling first for upright faces and then for inverted faces. When stimuli were upright, the results showed no effect of familiarity ($F(1, 29) < 1, p > .05$), or scrambling ($F(1, 29) = 3.35, p > .05$) and no interaction ($F(1, 29) = 2.96, p > .05$). However, when stimuli were inverted, the results were quite different. A main effect was noted both for familiarity ($F(1, 29) = 12.76, p < .01$) and scrambling ($F(1, 29) = 11.98, p < .01$) and these were moderated by the emergence of a significant interaction ($F(1, 29) = 20.96, p < .001$). Examination of the interaction through tests of the simple main effects showed no effect of familiarity when faces were highly scrambled ($F(1, 29) = 2.33, p > .05$). However, familiar faces were classified with slightly but significantly greater accuracy when intact ($F(1, 29) = 6.42, p < .05$), but with markedly less accuracy when moderately scrambled ($F(1, 29) = 22.15, p < .001$).

In sum, familiarity has little or no effect when the classification task is relatively simple (intact, entirely scrambled) but when the degree of scrambling is quite subtle and one feature is in the correct place within the face, then the scrambling task appears more difficult and shows clear evidence of being moderated by familiarity. In this case, however, the familiarity of the face makes the task harder rather than easier.

Response Speed for Correct Decisions

Mean speed of correct decisions was calculated again across familiarity, orientation and degree of scrambling. These are summarised in Table 8.2 below. The table confirms that again, performance seems fast and uniform when stimuli are upright, but seems markedly slower when stimuli are inverted. It is again in the inverted condition that differences across level of scrambling and across familiarity appear to be strongest.

The data were analysed by means of a 2 x 2 x 3 repeated measures ANOVA with familiarity, orientation and scrambling as the factors. As with the error data, the results showed main effects of all three factors, together with two-way interactions. Finally, the emergence of a significant three-way interaction modified all previous effects ($F(1, 29) = 19.58, p < .001$).

Table 8.2: *Mean response times (in msec, with standard deviations in brackets) for all conditions, showing fast classification of upright stimuli, and more variable classification of inverted stimuli.*

	Upright			Inverted		
Scrambling	Intact M (SD)	Moderate M (SD)	High M (SD)	Intact M (SD)	Moderate M (SD)	High M (SD)
Familiar	516 (91.3)	543 (63.4)	517 (55.0)	619 (105.3)	704 (117.0)	617 (97.2)
Unfamiliar	521 (75.9)	544 (65.9)	541 (62.9)	646 (113.4)	658 (107.3)	639 (99.3)

The three-way interaction was examined further by means of two separate two-way ANOVAs on the effect of familiarity and scrambling first when stimuli were upright, and then when stimuli were inverted. Taking the upright case first, the results somewhat surprisingly suggested a main effect of familiarity ($F(1, 29) = 9.56, p < .01$) and of scrambling ($F(1, 29) = 26.88, p < .001$). These were modified by an interaction of familiarity and scrambling ($F(1, 29) = 5.19, p < .05$). *Post-hoc* examination of this interaction showed that familiar and unfamiliar faces were classified with equal speed in all conditions except when stimuli were highly scrambled (intact: $F(1, 29) < 1, p > .05$; moderately scrambled: $F(1, 29) < 1, p > .05$). Under the highly scrambled condition, familiar faces were classified with greater speed than unfamiliar faces ($F(1, 29) = 34.42, p < .001$).

Consideration of the inverted stimuli gave a picture that was consistent with the error data. Here, analyses revealed no main effect of familiarity ($F(1, 29) < 1, p > .05$) but a main effect of scrambling ($F(1, 29) = 49.83, p < .001$). This was qualified by the emergence of an interaction between familiarity and scrambling ($F(1, 29) = 37.49, p < .001$). *Post-hoc* examination of the simple main effects showed this to be the result of the faster classification of familiar faces over unfamiliar ones when stimuli were intact ($F(1, 29) = 23.20, p < .001$) and highly scrambled ($F(1, 29) = 12.65, p < .01$) but the markedly slower classification of familiar faces when they were moderately scrambled ($F(1, 29) = 26.96, p < .001$).

Again, the analyses here complement the error data in showing that familiarity influenced the classification task when stimuli were inverted. A small but significant facilitation was evident when stimuli were easy to classify (intact, highly scrambled), but again a marked inhibition was apparent when stimuli were familiar and when they were hard to classify (moderately scrambled). In contrast to the error data where no effects reached significance, the RT data show some influence of familiarity when upright stimuli were classified. At all levels of scrambling, familiarity helped classification of the upright face, but this reached significance when the stimulus was highly scrambled.

Sensitivity

The non-parametric index of sensitivity – A' – was calculated as a measure of discrimination of old from new within each condition of the design. Mean sensitivity values are provided in Table 8.3 below. From this, it is clear that the sensitivity data mirror the pattern provided by the accuracy data. Performance was very good indeed and showed little variation when stimuli were presented upright. However, when inverted, performance was seen to vary across familiarity, especially when faces were only moderately scrambled.

Table 8.3: *Mean sensitivity values (with standard deviations in brackets) for all conditions (N = 30).*

	Upright			Inverted		
Scrambling	Intact M (SD)	Moderate M (SD)	High M (SD)	Intact M (SD)	Moderate M (SD)	High M (SD)
Familiar	.98 (.01)	.98 (.01)	.98 (.01)	.97 (.02)	.96 (.03)	.98 (.01)
Unfamiliar	.98 (.01)	.98 (.01)	.98 (.01)	97 (.03)	.97 (.03)	.98 (.02)

Analysis by means of a 2 x 2 x 3 repeated measures ANOVA showed no main effect of familiarity ($F(1, 29) < 1, p > .05$). However, a significant effect of orientation ($F(1, 29) = 10.6, p < .01$) and a significant effect of scrambling ($F(1, 29) = 14.4, p < .01$) emerged. These effects were modified by significant interactions between familiarity and scrambling ($F(1, 29) = 16.2, p < .001$) and between orientation and scrambling ($F(1, 29) = 9.2, p < .01$). Unexpectedly, and in contrast to the previous analyses, no 3-way interaction emerged ($F(1, 29) = 2.59, p > .05$).

Exploration of the interaction between familiarity and scrambling, by means of a two-way ANOVA with data collapsed across orientation, showed a main effect of level of scrambling ($F(1, 29) = 14.34, p < .001$) but no main effect of familiarity ($F(1, 29) < 1, p > .05$). The effect of scrambling was modified by the emergence of a significant interaction between familiarity and scrambling ($F(1, 29) = 16.16, p < .001$) and examination of the simple main effects here showed no difference across familiarity when faces were intact ($F(1, 29) < 1, p > .05$) but significant differences across familiarity when stimuli were both moderately scrambled ($F(1, 29) = 6.81, p < .025$) and highly scrambled ($F(1, 29) = 7.21, p < .025$). The means confirmed that when moderately scrambled, the familiar faces were at a marked disadvantage, but when highly scrambled, they were at a small but significant advantage.

In the same way, the interaction between orientation and scrambling was examined by means of a two-way ANOVA with data collapsed across familiarity. The results showed a significant main effect of orientation ($F(1, 29) = 10.57, p < .01$) and a significant main effect of scrambling once again ($F(1, 29) = 14.34, p < .01$). These were qualified by a significant interaction ($F(1, 29) = 9.19, p < .01$) and simple main effects for this interaction showed a difference across orientation when stimuli were intact ($F(1, 29) = 10.73, p < .01$), and moderately scrambled ($F(1, 29) = 13.06, p < .01$) but not when highly scrambled ($F(1, 29) = 3.05, p > .05$). The means suggested that when intact and when moderately scrambled, performance was better when upright. However, when highly scrambled, performance whilst better when upright, was not significantly better.

These results sit somewhat awkwardly with those from the analysis of error rates. However, the data may be examined on the basis of an *a-priori* prediction that the effect of familiarity across scrambling would be evident when inverted but not when upright. What emerges then is that upright performance is largely uniform across familiarity, with a small but significant advantage for familiar faces when highly scrambled ($F(1, 29) = 4.71, p < .05$). However, when inverted, the stimuli show the anticipated pattern, with a large and significant disadvantage for familiar faces when stimuli are moderately scrambled ($F(1, 29) = 8.20, p < .01$), but no effect of familiarity otherwise. The results of these analyses are then entirely compatible with those of the error analysis presented earlier.

In summary, the A' data approximate, but do not entirely mirror, the error data presented earlier. Results show that performance was better when stimuli were upright, and when stimuli were either intact or highly scrambled. Familiarity played a role inasmuch as familiar and unfamiliar faces were responded to equivalently when intact, but the familiar faces were at a slight advantage when highly scrambled, and at a distinct disadvantage when only moderately scrambled. These effects are most clearly seen when only inverted stimuli are considered. The results also suggested that upright stimuli were responded to with greater accuracy than inverted stimuli. This, however, was only significant in two of the three scrambling levels (intact, moderately scrambled).

Response Bias

Finally, B''_D was calculated as a measure of response bias. Negative values indicate a tendency to respond 'intact' while positive values indicated a tendency to response 'scrambled'. These values are summarised in Figure 8.2. From this it appears that there is an overall tendency to say 'intact' no matter what the image looks like. This is strongest when stimuli are inverted, and when they are familiar, and reaches its peak for familiar, inverted, moderately scrambled faces. In contrast, there is little or no bias evident when faces are unfamiliar and inverted.

Analysis by means of a 2 x 2 x 3 repeated measures ANOVA presents a complex picture. The results confirm a main effect of scrambling only ($F(1, 29) = 7.70, p < .05$), with no main effect of either familiarity ($F(1, 29) = 3.12, p = .088$) or orientation ($F(1, 29) = 1.36, p > .05$). Two 2-way interactions emerged to qualify these results: familiarity x orientation ($F(1, 29) = 9.82, p < .01$) and familiarity x scrambling ($F(1, 29) = 9.36, p < .01$). No interaction between orientation and scrambling was evident ($F(1, 29)= 3.99, p > .05$). Somewhat surprising, no 3-way interaction was apparent ($F(1, 29) < 1, p > .05$).

Examination of the two-way interaction between familiarity and scrambling showed a trend only towards a main effect of familiarity ($F(1, 29) = 3.12, p = .09$) but a significant main effect of scrambling ($F(1, 29) = 7.70, p < .01$). A significant interaction was also apparent ($F(1, 29) = 9.36, p < .01$). *Post-hoc* examination of the simple main effects here indicates that whilst familiar stimuli show a greater tendency to elicit an 'intact' response all the time, this reaches significance only for the moderately scrambled case ($F(1, 29) = 9.66.\ p < .01$).

Examination of the two-way interaction between familiarity and orientation showed again a trend only towards a main effect of familiarity ($F(1, 29) = 3.12, p =$ .09, and no effect of orientation ($F(1, 29) = 1.36, p > .05$). However, a significant interaction between familiarity and orientation was present ($F(1, 29) = 9.82, p < .01$). *Post-hoc* examination of the simple main effects of this confirmed no effect of familiarity when stimuli were upright ($F(1, 29) < 1, p > .05$), but a significant effect of familiarity when stimuli were inverted ($F(1, 29) = 13.38, p < .001$).

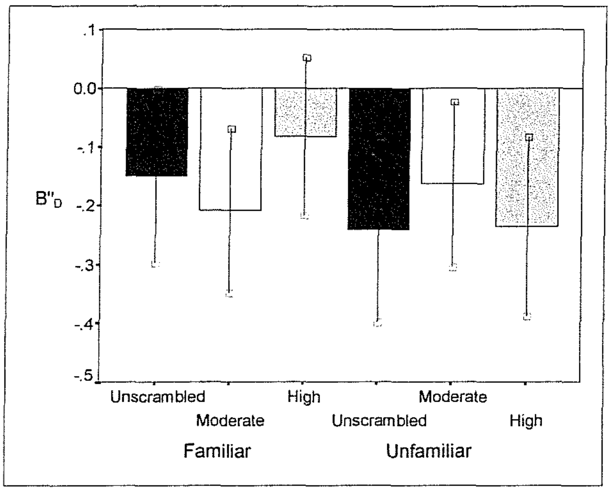

Figure 8.2. The bias to respond 'intact' across familiarity when faces are presented upright.

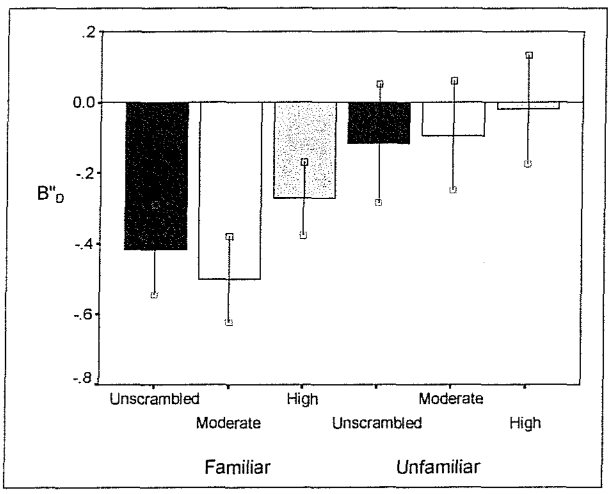

Figure 8.3. The bias to respond 'intact' to familiar faces is more pronounced than for unfamiliar faces when faces are inverted.

In sum, the pattern of response bias suggested that participants tended to respond 'intact' more often than not. Somewhat surprisingly, this bias was most evident not when faces were indeed intact but when faces were moderately scrambled.

This held across both upright and inverted cases, but was strongest in the familiar, inverted, moderately scrambled condition. Consequently, the pattern of response bias obtained here illustrated that the moderately scrambled condition again provided the perceiver with a difficult and error-prone task.

Summary of Results

Accuracy, response times and A' values all indicate a consistent pattern of results: Performance is good and fairly uniform when stimuli are upright. However, when inverted, an influence of familiarity is seen such that in the hardest scrambling condition provided (moderately scrambled), performance for familiar stimuli is impaired relative to performance with unfamiliar stimuli. The impairment seen here is surprising given the previous studies showing a preservation of performance for familiar stimuli. This impairment is, however, consistent with the view provided by the robust representation hypothesis.

Discussion

As the final study in this thesis, the aim was to continue the trend of previous studies and examine the effect of familiarity on face processing. A face-classification task was used. In contrast to the previous two experiments, the present task required a traditional face/non-face decision rather than a normal/odd decision. As in previous experiments though, task difficulty was increased by using stimulus inversion in order to explore the role of familiarity under conditions of normal and increased mental load. On the basis of previous results within the thesis, it was anticipated that the results would favour the robust representation viewpoint.

The results showed a complex set of effects, some anticipated and some not. First, all measures verified that the task was relatively easy when stimuli were upright. Performance was relatively error-free and response times were consistently fast. However, when stimuli were inverted, task difficulty increased substantially. Analysis also showed that it was precisely when task difficulty was great (and stimuli were inverted) that familiarity showed its effect most clearly. Up to this point, the results of the present study support those of all previous valid experiments within the thesis. Where the present results deviate from the previous ones, however, is that instead of performance being facilitated by familiarity under difficult conditions, performance here was actually impaired. Specifically, when stimuli were inverted and moderately scrambled, familiar faces elicited performance that was slower and more error prone than for unfamiliar faces. How should these results be interpreted?

Consideration of the task demands suggests that the effect of different levels of scrambling plays a crucial part in explaining these results. When stimuli are intact, the task of verifying whether a stimulus is a face or a non-face is trivial. Performance is rapid and error-free. Indeed, both the error data and the response time data indicate that, when the task is made difficult through stimulus inversion, the familiarity of a stimulus nevertheless serves to facilitate performance with intact faces. In other words, participants are quicker and more accurate at saying 'intact' to an inverted stimulus if it is familiar than if it is not. This is important because it echoes the results of previous experiments within the thesis. Familiarity has helped under conditions where performance is otherwise made difficult. Facilitation occurs here because the result of processing identity (that is Matt le Blanc) confirms the message from the processing of faceness (that is a face). In essence, the fact that Matt le Blanc *can* be identified provides a strong cue that the stimulus must indeed be intact. Without a robust representation for the familiar face, that identity signal is absent and so there is no mechanism to facilitate the primary task of face classification.

In a similar way, when all features are in the wrong position, as in the case of a highly scrambled stimulus, then the task is again simple. One might consider that through our expertise with the broad class of faces, we have a very strong (or robust) representation of what constitutes a correct anatomical face. An intact stimulus will easily be confirmed as a face an equally, an entirely scrambled face will easily be denied as a face. Facial familiarity has little effect on the ease of decision to highly scrambled stimuli, and this is perhaps not surprising given that identity will be difficult (though not impossible) to determine from a face in which all features have been moved.

The case of the moderately scrambled face presents a different problem though. If the view of Donnelly *et al.* (1994) and Collishaw and Hole (2000) is to be taken then scrambled faces are processed in a serial, feature-by feature fashion in order to determine whether each feature is in the correct place. If no features are in the correct place then the non-face decision can be taken quickly. However, if one feature is in the correct place then the stimulus requires further checking. As a consequence, response times are slowed and performance may be more error-prone if participants are conscious of time-pressure in the task. Why then, is performance on these moderately scrambled faces influenced by familiarity, and why is familiarity a hindrance rather than a help?

Here, the notion of a robust representation formed through familiarity with a stimulus is again important. It is argued that possession of a robust representation allows the identity of a face to be accessed with great speed. One might argue that this process appears automatic and appears to be the primary purpose of the face recognition system. It may also be expected that such a robust representation can be accessed even though an input may be partial. For example, the correctly arranged top half of the face, or the correctly arranged bottom half of the face may be sufficient to access the stored robust representation of a familiar face. Intuitively, this situation accounts for the fact that we can often recognise a familiar individual even though they may be occluded or angled away from us. The literature also shows a correlate of this situation in the demonstration of priming of face recognition from earlier presentation of only a part of that face. For this to happen, the whole representation must have been accessed initially even though only a part of the face was presented. Similarly, the literature has shown priming of a part of the face by prior presentation of a non-overlapping other part of the same face (Ellis, Burton, Young, & Flude, 1997). Again, the whole representation must be activated given presentation of only a part at both study and at test.

Within the current experiment, it might be considered that the moderately scrambled face effectively provides the face processing system with an intact part of a familiar face and that this is sufficient to provoke some activation of the robust representation held for a familiar individual. As such, the moderately scrambled face presents the perceiver with a contradiction. It signals an identity (that is Matt le Blanc) whilst at the same time signalling that it is a non-face. The consequence is that the stimulus requires further processing in order for the contradiction to be resolved. The stronger the mentally held representation for a face – the more familiar that face is – the less evidence in the form of matching input is required before activation of the representation occurs. As such, familiarity will hinder the classification of a moderately scrambled face because of the partial activation of the identity to which it is associated. This effect is seen on examination of accuracy and response time data, and is echoed to some degree by analysis of the A' data and the response bias (B''_D).

In summary, these results may be wholly rationalised using the robust representation account. When faces are intact, familiarity and intactness provide a consistent message for the perceiver and thus performance is fast and accurate. When faces are highly scrambled, the oddness of the face is quickly identified (one merely needs to attend to a single feature to elicit this information) and this is so even in the inverted condition. However, when faces are moderately scrambled, robust representations are of little help in classification tasks. In fact they make performance worse, as there is a conflict between the fact that an identity can (weakly) be determined from a stimulus that is not a complete face. In this case, the difficulty in establishing the intactness of the face is compounded by the automatic attempt to access the identity of the face. As a consequence, the moderately scrambled stimulus does not immediately appear scrambled, and this time delay in determining faceness is enough to initiate the identification process, which then interferes with the primary task of classification. This interference was shown to preserve performance in the previous classification tasks (Experiments Five and Six), where the identity of the face was not compromised (That's Matt le Blanc, and I recognise him even with his eyes and mouth altered). Nevertheless, in a scrambling task where the internal facial features are out of position, and the primary task is deciding whether a face is indeed a face, partial recognition (I think that might be Matt le Blanc, but I'm not sure) hinders primary task performance.

Conclusions

This study has provided confirmatory evidence for the influence of familiarity at a classification level. However, in contrast to the previous classification studies, familiarity has not been shown to advantage face processing in all conditions. On the contrary, for the case of moderately scrambled faces, knowing a face has been shown to be a disadvantage. This contradictory finding can, nonetheless, be explained in terms of robust representations. Both the notion of robust representations, and the notion of automatic processing of familiarity feeding back to earlier stages of visual processing, require that the theoretical models describing a sequence of face processing stages need to be revisited. This will be the focus of the General Discussion within this thesis.

CHAPTER 9

General Discussion

The aim of this thesis was to provide an empirical demonstration of the influence of familiarity on face processing. Throughout the thesis three hypotheses have been tested. These have arisen through interpretation of the existing empirical findings. The first hypothesis was that familiarity would have an influence on face processing because it accesses expert operations such as configural processing (Rhodes *et al.*, 1989). Disrupting access to configural information by means of inversion or negation means that performance on familiar faces should have been degraded if the expertise hypothesis is correct.

The second hypothesis was that familiarity would have an effect on face processing because it leads to a robust representation of face, which preserves recognition under difficult viewing conditions (Tong & Nakayama, 1999). This hypothesis holds that performance should be better for familiar than unfamiliar faces when any manipulation is used.

Finally, and in the light of previous findings (Collishaw & Hole, 2000; Scapinello & Yarmey, 1970; Yarmey, 1970), the null hypothesis stated that there would be no qualitative differences in performance between familiar and unfamiliar faces. This outcome would support the idea that while familiarity allows faster, more accurate recognition of faces, it does not qualitatively affect the process used.

These hypotheses were tested in recognition and classification tasks by inverting or negating stimuli, which are known to disrupt the use of configural processing, and by comparing results against the recognition of unmodified stimuli. The first set of experiments involved recognition tasks, which were used to test for differences in performance on upright/positive and inverted/negated stimuli across familiarity. The second set of experiments presented within this thesis was designed to examine a lower level of processing, and as such took the form of face classification tasks. The point of interest here was whether familiarity would exert any influence at an earlier stage of processing than required for recognition. The same predictions were made in each of the studies. That is, familiarity would either hinder processing (expertise hypothesis), help processing (robust representations hypothesis) or have no effect (null hypothesis). In some senses, the second set of experiments was the more interesting in that one might anticipate a familiarity advantage to recognition studies, but there is no reason one might expect familiarity to have any influence in a task which ostensibly requires processing prior to familiarity being signalled. In this chapter, the findings of the experiments at the two different task levels will be compared and then integrated to describe the role of familiarity in face processing. First, however, the methodologies used will be summarised with an eye to their strengths and to areas of improvement.

Summary of methods used

All participants were either undergraduate, or postgraduate students within the university. Participants volunteered for, and took part in, one experiment only. This was to ensure that each was naïve to the recognition task (where appropriate) and to the set of unfamiliar faces.

Stimuli used in the experiments comprised personally familiar faces, publicly familiar faces, and unfamiliar faces. Personally familiar faces were all members of the Psychology Department, and therefore well known to the participants. Publicly familiar faces were celebrities (actors, pop-stars, and sports personalities) and were downloaded from various Internet web-sites. Stimuli for each experiment came from within the same celebrity set. Unfamiliar faces comprised a mix of staff and students from other university departments, and images from the Stirling Pics database. Although it might be said that the celebrity images were more glamorous than the unfamiliar faces, and therefore not truly matched, the results of the recognition studies using these images (Experiments Three and Four) are not substantially different to those using personally familiar faces (Experiment Two). This suggests that any cosmetic differences in the appearance of the images were not enough to have influenced findings.

All images were prepared in exactly the same way, and to the same dimensions in order to minimise variations between stimulus sets. Presentation of stimuli was uniform across all studies, and indeed all studies were run on the same PC thus ensuring stability of resolution and luminance, etc. Experiments 1-4 used a specially designed program written in C++, while Experiments 5-7 were written using Superlab experimental software.

The task used in the first experiment was a speeded familiarity task, while Experiments Two, Three and Four used an old/new recognition task. The final three experiments used face classification tasks, the idea being to seek converging evidence from two different levels of performance. In all experiments speed and accuracy of response were the measured variables. Additionally, sensitivity and bias were calculated in all tasks requiring a yes/no response (all but Experiment 5).

Non-parametric measures of sensitivity (A') and bias (B''_D) were used due to the small number of experimental trials in each condition. Analyses on the differences between these measures, however, used parametric tests as the number of participants was great enough to allow this, and data were approximately normally distributed.

Summary of main findings

In Experiment One, personally familiar and unfamiliar faces were used in a simple speeded familiarity task to test the hypothesis that familiarity would influence the effect of stimulus manipulation. The manipulation used was inversion. The rationale for this assumption was that inversion has been used extensively to test for expertise effects in the face recognition literature. It was found that familiarity *did* have an effect on the role of inversion. Recognition of inverted faces was more error prone when they were familiar, than when they were unfamiliar, and as such, these results supported the expertise hypothesis. However, this finding was found to be the result of a response bias to call all inverted faces 'unfamiliar'. Furthermore, the design of the experiment was not optimal to test the hypotheses as "familiar" decisions could be made only for the familiar faces, (only familiar faces could appear familiar), and so the task demands across familiar and unfamiliar stimuli were unbalanced. Therefore the results are treated with extreme caution and the following experiments were designed to overcome these methodological problems.

The second experiment used the same stimulus set as Experiment One and aimed to correct the problems above, by employing an old/new recognition task with upright and inverted faces. This study demonstrated the influence of familiarity in terms of accuracy, sensitivity and bias. In contrast to Experiment One, performance on inverted stimuli was more accurate when faces were familiar and both the accuracy and the sensitivity data support the idea of robust representations being formed for familiar faces. These representations are capable of withstanding changes in viewing conditions such that a smaller disruption to processing results. However, the upright stimuli elicited ceiling performances for the familiar faces, and this calls into question the reliability of the results. In addition, there was a bias to call familiar faces 'old' and unfamiliar faces 'new'. This *false fluency* effect was present whether faces were upright or inverted and suggests that representations for familiar faces are likely to be recalled easily, and sometimes, inappropriately.

Experiment Three replicated the previous experiment but used a different stimulus set of familiar faces in order to increase the difficulty of the task. A set of publicly familiar (celebrity) faces was used with the result that accuracy was reduced slightly. However, the findings *are* consistent with those in Experiment Two in terms of accuracy. The interaction between familiarity and inversion for the sensitivity data did not, unfortunately, reach significance. Nevertheless, these two experiments taken together give support to the hypothesis that familiarity influences processing by preserving recognition under inversion. No *false fluency* was seen in this experiment and it is argued this may be due to participants having a less robust representation of publicly, than personally familiar faces. Alternatively, it may be due to the incapacity to engage in semantic priming for celebrity faces that lack strong association both to one another and to the context of a psychological experiment.

The final recognition study (Experiment Four) used a different manipulation to investigate the influence of familiarity. The manipulation used was image negation. The rationale behind this manipulation change was that inversion and negation might have similar effects in terms of disrupting configuration. The same publicly familiar and unfamiliar faces, as used previously, were presented in photographic positive and negative in an old/new recognition task. The results were in line with those in Experiments Two (and partially with Experiment Three) with better performance for familiar than unfamiliar negated faces. Although performances on positive images were at ceiling for familiar faces, the uniformity of results across these three studies suggests the effect of familiarity is reliable.

Throughout these experiments, the use of the same image at inspection and at test may be criticised on the grounds that participants could use a pictorial code, rather than a structural one, to identify faces (Bruce & Young, 1986). Bruce and Young have suggested that structural codes are used to identify known faces, whereas unfamiliar faces are as likely to be identified by a general pictorial code as a structural one. This means that rather than focussing on the properties of a face (e.g., large eyes, flared nostrils, wide mouth, etc.) attention is directed to a property of the shown example of an individual face. Therefore, recognition at 'time two' may be accomplished without recourse to the facial features which mark out an individual, but by relying instead on a particular aspect of the example shown (e.g., hair sticking up on end, distinctive shadows, etc.). Whilst this is acknowledged as a problem (indeed, with hindsight different views should have been shown at inspection and test) it could be argued that the results would be affected only by degree given this change. Using the same pictures at inspection and test allows familiar face recognition to proceed either structurally or pictorially. Unfamiliar faces, on the other hand, will not have a stored representation (or only a very weak one from inspection) and so only the pictorial code is available for identification. Therefore processing differences should become apparent in the response time/accuracy data. Using different views at inspection and test will still allow recognition of familiar faces to proceed structurally (representation available) but not pictorially. Moreover, although unfamiliar face identification can proceed, it can only do so based on a weak representation, so again, a processing difference should be seen.

Thinking this issue through at a more practical level permits a clearer picture to emerge. In the present context, repetition of an image at study and test allows a perfect match between the stored representation of an unfamiliar face (based on pictorial and facial cues from one instance) and the test image for that face. Therefore, performance should be good. On the other hand for familiar faces, the representation resulting from study may be a composite of the shown exemplar <u>plus</u> all previous instances (or an abstracted prototype of them). Thus, the perfect mapping of study to test image possible for unfamiliar faces may <u>not</u> be possible for familiar ones. The implication of this is that the results would bias against the robust representation hypothesis, because superior performance in the unfamiliar case, despite any image manipulations, might be anticipated. That the results did not show this pattern of superiority, and in fact showed the contrary pattern of superiority when familiar, may actually be taken as a strength of the robust representation viewpoint, despite what may be considered, in hindsight, as a methodological weakness. The consistent pattern of results across the thesis testifies to the strength of the effect shown. This is especially so given the classification studies where images are presented once only and so the question of pictorial vs. structural codes is not an issue.

Experiment Five was the first of the face classification studies to be used. This two-alternate forced choice paradigm required participants to decide which of two presented images was 'odd' where one image was normal, and one had been 'Thatcherised'. The stimulus sets were as used in Experiments Three and Four and comprised publicly familiar and unfamiliar faces. The results showed the expected pattern in that Thatcherisation was easy to detect when upright, but difficult to detect when inverted. Moreover, analysis of both accuracy and response time showed that even though detection was difficult when inverted, classification was better and quicker when faces were familiar. This result is as unexpected as it is interesting. Familiarity with a face is not required to complete this task, however, it appears that knowing a face makes the Thatcher illusion easier to spot. A confound here is that participants may perform the task in different ways. They may determine which face is normal, and select the other face as odd, or they may determine that one face looks less normal (or more odd) than the other. Finally they may determine the odd face without recourse to the relative strategies used in the other two decisions. In order to remove this ability to use a relative, rather than an absolute strategy, the task was changed to a yes/no decision task (Experiment Six). The results were the same. Familiarity was demonstrated as preserving performance in a classification task thus supporting the robust representation hypothesis.

Although the Thatcher Illusion has been used extensively as a means of determining configural processing in face recognition research, it cannot be called a true classification task. In the Thatcher illusion, the face remains a face. The only manipulation is in the inversion of the eyes and mouth. For this reason, a face-scrambling task was selected as the final test of the influence of familiarity on processing (Experiment Seven). The same stimuli were used, with a yes/no decision task where participants decided whether the face was intact (face), or scrambled (non-face). Faces were presented either intact, moderately scrambled (one feature in its correct position) or highly scrambled (position of all features changed). The analyses presented some interesting findings. When faces were intact or highly scrambled (all features changed), performance on familiar faces was faster and more accurate than unfamiliar faces. However, when faces were moderately scrambled, the opposite was the case. It is proposed that this paradox is due to the formation of a robust representation for familiar faces, which makes it easy to decide whether a highly scrambled face is indeed a face or not. Contrastingly, when faces are moderately scrambled, there is more difficulty in deciding whether a face is a face, especially if the face is known to the perceiver and a partial match to a robust representation is activated.

These three classification studies have confirmed the findings of the recognition tasks: familiarity leads to a preservation in recognition performance under difficult viewing conditions. Therefore, at two processing levels, and across two sets of stimuli, two manipulations, and three different tasks, the empirical studies presented in this thesis have supported the notion that knowing a face allows the formation of a robust representation for that face which influences the way the face is processed. The remainder of this chapter will discuss the theoretical implications of this position, and discuss the future direction to take this research.

Theoretical implications

This thesis began by presenting the theoretical models and processes that have been influential in describing the likely mechanisms thought to be fundamental in processing familiar faces. How do the findings from this thesis fit within these models? The first of these, Bruce and Young's (1986) Information Processing model, based on Hay and Young's (1982) earlier version, described a hierarchical sequential-staged structure with parallel processing assumed for recognition of identity, expression, and facial speech. However, the model is unable to account for the findings from either the recognition or the classification tasks in this thesis.

Taking the recognition tasks first, the issue at hand is the finding that recognition performance was better in familiar than unfamiliar faces when stimuli were manipulated. Bruce and Young's model, however, was essentially a familiar face processing model, and as such offered no processing route for unfamiliar faces. In order to explain the present recognition results, one would have to assume that an FRU can be generated after a single presentation of a face, and that being based on a single, perhaps idiosyncratic view of the person, it might be activated less effectively in a later recognition task. Whilst the first assumption is possible, the second counters Bruce and Young's belief that the FRU worked in an all-or-nothing fashion. As such, it is difficult to explain the present recognition data without modification of the basic model.

The fact that facial familiarity affected performance on the present classification tasks is also difficult to account for within Bruce and Young's model. This is because the model assumes a sequential-staged process, where one stage must be completed prior to moving on to the next stage. In terms of signalling faceness and familiarity, the model presumes that faceness is signalled through structural encoding, and that this must be completed before the FRUs can be activated to indicate familiarity. Furthermore, the model does not specify any processing prior to familiarity being signalled, and, as such, provides no account for how faces are dealt with in the earlier stages of processing. The finding that familiarity can help a process which is completed before familiarity is registered seems at odds with the model, as there is no mechanism for feedback from the FRUs to structural encoding. This leads to two possible explanations. Either familiarity is detected earlier within the model than has been suggested, or, there is some mechanism for automaticity and feedback, such that familiarity is accessed prior to the completion of the classification task, and so can inform the perceptual experience.

The first explanation, that familiarity is detected earlier than at the FRU level seems unlikely, given the data from timing studies, error diaries, and neurological conditions to suggest otherwise. This leaves the idea of automaticity and feedback. There is no facility within the Bruce and Young model to include the feedback of familiarity in order to aid performance on classification tasks. Furthermore, although Valentine's (1991) MDFS model can help us to understand structural encoding, there is no mechanism for feedback within this system either, so this also fails to explain the findings of this thesis. However, the modelling of familiar face recognition using IAC (Burton, Bruce & Johnston, 1990) does offer an explanation to account for the findings of the recognition *and* classification studies.

IAC consists of a set of units connected by links which carry weights. The weight of a link determines the strength of that link and so determines the proportion of activation that can be propagated forward and the speed with which this can be achieved. Within IAC, it is assumed that link strength builds up over repeated exposures: effectively the more times a link is used, the stronger it gets. In this sense, familiar faces by necessity will be represented by units connected with stronger links compared to unfamiliar faces. This is important in explaining the present recognition results because, as graded signallers of resemblance (rather than all-or-nothing operators) the units within IAC can be activated by degrees. Given their relative link strengths, it should be clear that the familiar face units will always be activated more than the unfamiliar ones, hence the familiar superiority. In terms of the classification data presented here, a mechanism of feedback is required in order to explain the influence of familiarity at this level. IAC can provide this too with reference to the bi-directional links between information pools. This means that a familiarity signal from the PIN may propagate backwards to inform an 'earlier' stage of processing. Therefore, both sets of findings may be accounted for within the IAC architecture.

Process Level Explanations

So far, the findings within this thesis have been explained in terms of the available theoretical models. However, these models do not make explicit which processes are used in recognition and classification of faces. At a process level, the dominant view is that faces are processed with reference to their configuration; that is, the spatial layout of features within a face. Moreover, it has been proposed that experts are more likely to rely on configural than featural processing. If observers can be described as being more expert with a stimulus class with which they have had more exposure, then familiar faces are one such class. The empirical evidence presented in Chapter One suggests that configural processing is the hallmark of expertise, and that this processing may be disrupted by either inverting or negating the stimuli. If experts rely on configural processing, then disrupting the ability to use this form of processing should result in the degradation of expert performance (the expertise hypothesis). The results of the studies presented here do not support this position. Familiar faces, with which participants were deemed to be more expert, were better recognised/classified than were unfamiliar faces, both when inverted and when negated. Therefore, the familiarity effect seen here *cannot* rest on better configural processing.

Instead the alternate hypothesis, based on the notion of *robust representations*, that familiar stimuli would be less affected than the unfamiliar stimuli by manipulation, was supported. The question now is what does the concept of robust representations give us? Is the term 'robust representation' just another way of labelling familiarity, or are there any processing implications? This issue will now be addressed.

Defining Robust Representations and their Function

What is a robust representation? It is possible that robust representations may function in a similar way to facial *prototypes*. The *prototype effect* refers to the tendency to abstract an unseen version of face that is perceptually related to a number of seen faces (Solso & McCarthy, 1981). This effect has been demonstrated using face recognition tasks (e.g., Bruce *et al.,* 1991; Cabeza, *et al.*, 1999; Solso & McCarthy, 1981), as well as object classification tasks (e.g., Homa, Goldhardt, Burruel-Homa, & Carson Smith, 1993). Classically, this effect has been demonstrated using *identikit* faces (Solso & McCarthy, 1981) or schematic faces (Bruce *et al.*, 1991). More recently, however, the robustness of the prototype effect has been demonstrated using colour photographs of faces which were presented from different views (Cabeza *et al.*, 1990), and under different lighting conditions (Tong & Nakayama, 1997). This suggests a generalisability of the effect and makes it a contender for the basis of a robust representation of a face.

How might a robust representation preserve performance? It is tempting to imagine that the existence of a robust representation will lead to different (and better) processing of faces. This would, inevitably, lead one to speculate that it aids configural processing, given that this is the way most theorists believe recognition is enhanced. However, it has been established that inversion/negation reduces the ability to process faces configurally. Therefore, whatever process is being aided by the use of robust representations to perform the recognition or classification tasks under manipulation in these studies, it is unlikely to be configural processing. Cabeza and Kato (2000) support the notion of dual coding and assert that featural information is as important an aid to recognition as configural information, a view confirmed by Rakover and Teucher (1997) and Collishaw and Hole (2000). It is possible that other processes are involved as well, but the evidence for this is somewhat vague. As such, the proposal that robust representations permit the use of a different or better processing strategy and thus aids performance, doesn't take us forward unless we can identify what the process might be.

More parsimonious is the idea that robust representations lead to faster, more accurate, and seemingly automatic processing. Consider the following hypothetical situation (Figure 9.1 below): performance on a primary task (T_1), takes a certain amount of time (t_1), performance on a primary task where viewing condition are made difficult (T_2) takes longer (t_2, where t_2 is greater than t_1). A secondary task, recognition (R), is automatic and proceeds in parallel with the primary task. When faces are familiar (R_F) and the primary task is easy (T_1), the primary task is completed well before the face is identified ($t_1 < r_f$) and the familiarity of the face has been of no help. However, when the primary task is difficult (T_2), such as in the case of inverting or negating the stimuli, and faces are familiar, robust representations enable the secondary task, recognition, to be complete prior to T_2 ($r_f < t_2$) and therefore information in the form of the face's identity is fed back to aid completion of T_2. When faces are unfamiliar, and there are no robust representations to speed recognition, the secondary task (R_U) takes longer than the primary task, and so there is no additional help available to speed performance on T_2.

This idea of automaticity and feedback is not without precedent. There is empirical evidence that feedback within the system has contributed to improved performance in expression detection (Baudouin *et al.,* 2000), gender identification (Bruce, 1987; Roberts & Bruce, 1988), and mental rotation (Ruthruff & Miller, 1995). The findings from Baudouin *et al.* and Ruthruff and Miller suggest that when a task is made difficult, such that processing is delayed, familiarity may be signalled prior to task completion, and this can aid task performance.

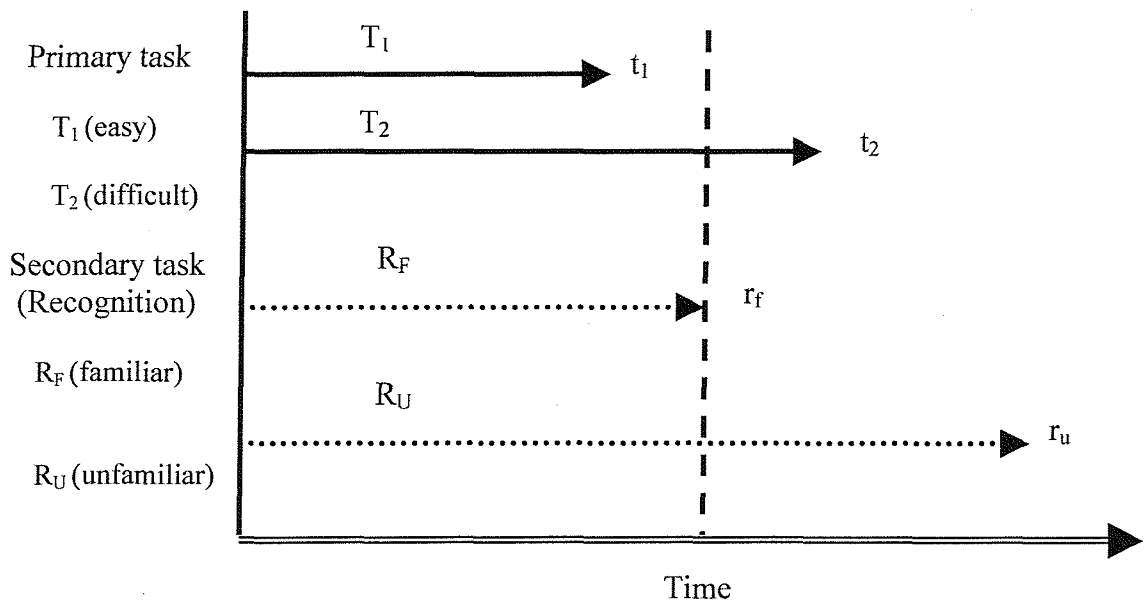

Figure 9.1: Illustration describing the processing advantage gained by familiar faces when viewing conditions are difficult.

At this stage it is important to step back and view these ideas in terms of what the literature considers as a qualitative versus a quantitative effect. The distinction between a qualitative and a quantitative difference may be explained through analogy to a school sports day. In race A, Jane runs and Susan hops. Jane wins. Jane's performance is better because she used a better strategy – running. In race B, Peter runs and Ben runs as well. Peter wins. The two children are using the same strategy, but Peter is two years older and so is taller, stronger and quicker. Race A describes a qualitative difference in performance, while race B describes a quantitative difference. How do these levels of explanation relate to the current findings? A qualitative difference may be indicated if a manipulation affects one strategy and so the processing of one set of stimuli, more than the other. In the present results a qualitative difference in processing would follow if it were believed that the robust representation, held for familiar faces, allowed a different and better processing strategy to be used, which was not affected by stimulus manipulation. On the other hand, the more likely explanation that the robust representation for familiar faces allows the faster use of common face strategies describes a quantitative effect. The evidence for a quantitative difference in processing is supported by a wealth of empirical evidence, and by connectionist modelling. IAC accounts for more efficient processing in terms of stronger links between pools and greater activation within units.

Although, at face value, a quantitative shift in processing from familiar to unfamiliar faces may appear to be a relatively uninteresting conclusion, what underlies this assumption is the part played by robust representations, without which, familiar faces would be processed in the same way as unfamiliar faces. It is proposed that robust representations are the tools which enable the rapid and automatic processing of familiar faces, but this alone is not enough to account for the findings presented herein. Automaticity then, is necessary but not sufficient. The mechanism required to complete the explanation is feedback. Together, but independent of each other, automaticity and feedback operate to preserve performance under difficult viewing conditions when faces are familiar.

The incorporation of automaticity and feedback into Bruce and Young's model would produce something rather like Figure 9.2 below. This shows the parallel processes of classification and familiarity, along with the facility for feeding back from the FRUs into structural encoding. Although, initially, classification occurs prior to familiarity being assessed, the two processes continue in parallel when viewing conditions render classification more difficult. Identifying familiarity at the FRU is an automatic process, which, once initiated, carries on to completion. Information from the FRU feeds back to inform the process of classification.

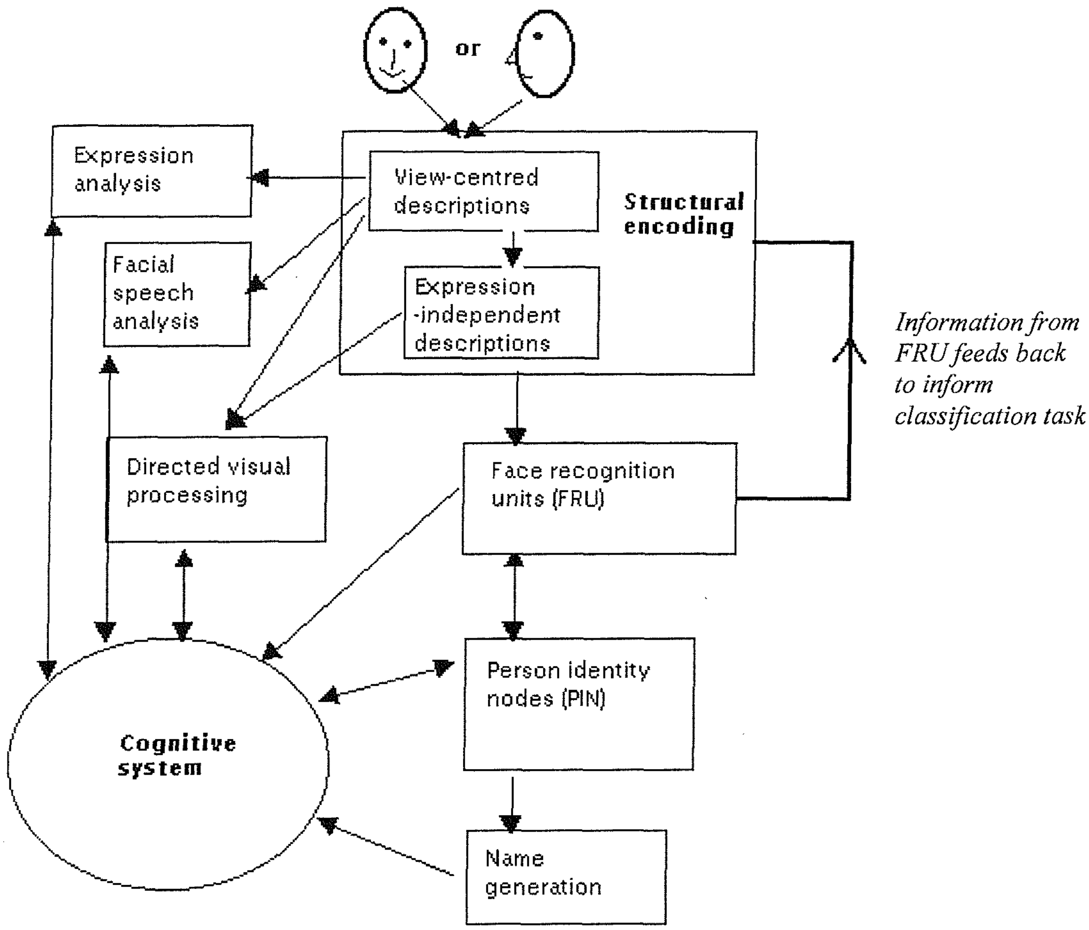

Figure 9.2: Bruce and Young's (1986) Information Processing model (adapted) incorporating a pathway for information from the FRU to feed back into structural encoding.

Furthering the findings from this thesis.

The findings described in this thesis have, in some sense, been surprising, and for this reason a number of future projects are required to confirm and extend the results. Perhaps the most important of these is to replicate the recognition tasks using different views at inspection and test. This would establish whether a structural code is indeed being used to advantage the familiar faces. Of course, this would have no bearing on the findings from the classification studies where only one instance of each image was presented.

The results have shown familiarity effects at an item level. Investigation of subsets of faces, with which participants are familiar, would help to clarify the mixed results from studies exploring expertise effects at a sub-group level, for example, using own-race and other-race faces (e.g., Valentine & Bruce, 1986 vs. Rhodes, *et al.*, 1989). Race is one obvious sub-group across which familiarity may vary. In addition, investigating those who have expertise with a certain age group, such as babies or the elderly could be profitable, as clear expertise predictions across expert and novice participants could be explored. Babies' faces may be particularly useful, as a stimulus set, as they are highly homogenous, and therefore very difficult to differentiate unless one is familiar with a particular baby (Chance, Goldstein, & Andersen, 1986). Furthermore, de Roiste & Bonnet (1995) have shown that multiparity (having more than one child) does not confer an advantage in recognition of infants' faces over primiparity (first-time mother). This may not be so surprising, given that maternal expertise would a) take time to develop, and b) is specific to their *own* offspring. However, comparing recognition between novice and expert groups (e.g., nursery nurses, midwives, etc.) where exposure to a huge number and variety of exemplars may have facilitated the development of a robust representation for experts, might be more informative. Such an experiment is already being planned.

A further area for development is the investigation of physiological measures such as galvanic skin responses (GSR) or EEG patterns in classification tasks. Of particular interest would be a difference in response to familiar and unfamiliar moderately and highly scrambled faces. It was proposed that the partial recognition of moderately scrambled familiar faces might have contributed to the slower and more errorful performances. By examining the pattern of physiological responses across scrambling and familiarity, the influence of familiarity could be assessed physiologically, in order to complement the findings in this thesis.

Finally, modelling the findings from the classification tasks would provide exploration of the proposal that automaticity and feedback are the independent and essential requirements allowing the preservation of performance under difficult viewing conditions. This would also add support to the presented adaptation of Bruce and Young's model described above.

Conclusions

The experiments conducted within this thesis have provided evidence for a difference between the processing of familiar and unfamiliar faces. This difference has been shown in both recognition and classification tasks. The expertise hypothesis, that performance on familiar faces would be degraded compared to unfamiliar faces, and the null hypothesis, that there would be no difference in performance across familiarity, were not supported. The hypothesis that performance under difficult viewing condition would be preserved where faces were familiar is accepted.

It is proposed that performance is preserved by the development of robust representations for familiar faces, and that this allows fast, accurate and relatively effortless processing. Furthermore, this is facilitated by the notion of feedback, whereby familiarity is signalled and feeds back into structural encoding to aid primary task performance. Automaticity and feedback are suggested as being independent mechanisms, both of which are required in order to allow a superior level of processing of familiar faces to proceed whether viewing is impaired, or not. In other words, the robust representation gives the familiar face a buffer to cope with life's unpredictable conditions and so maintain performance. The goal of the system is identification, and the robust representation ensures this is achieved.

APPENDIX A

Stimuli used in Experiment One

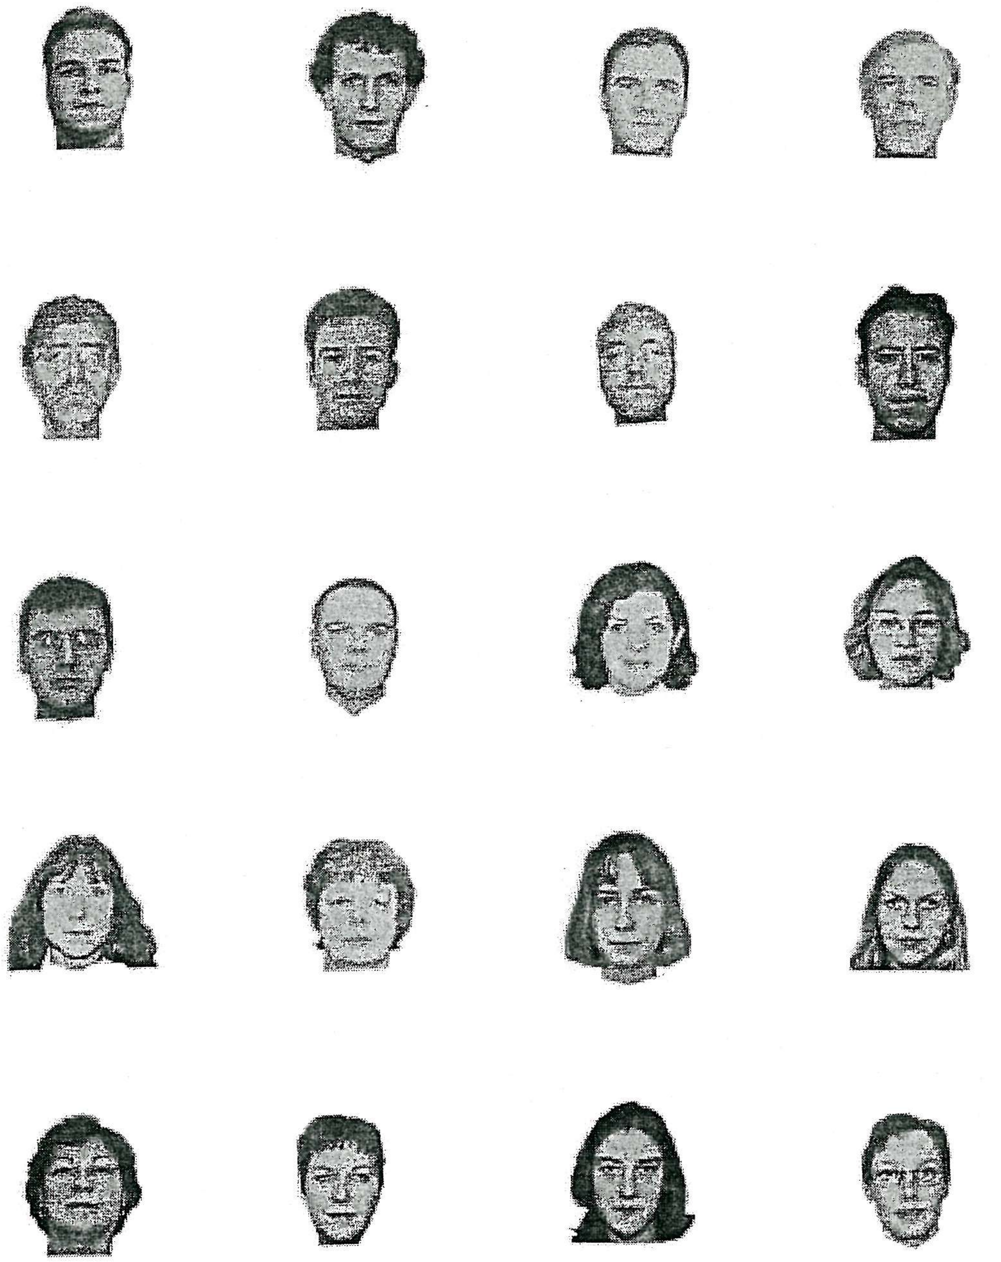

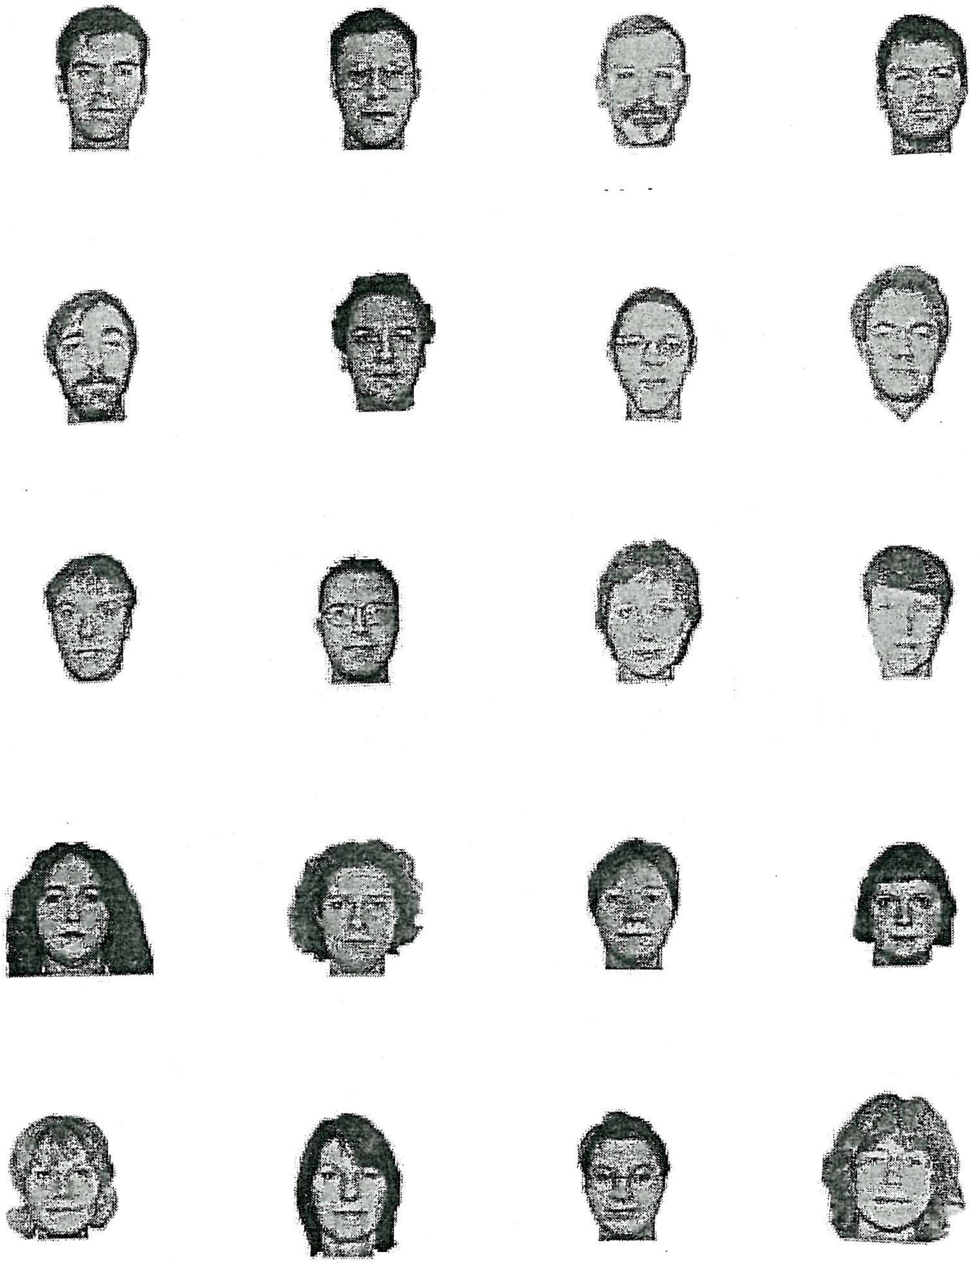

APPENDIX B

Full set of personally familiar stimuli

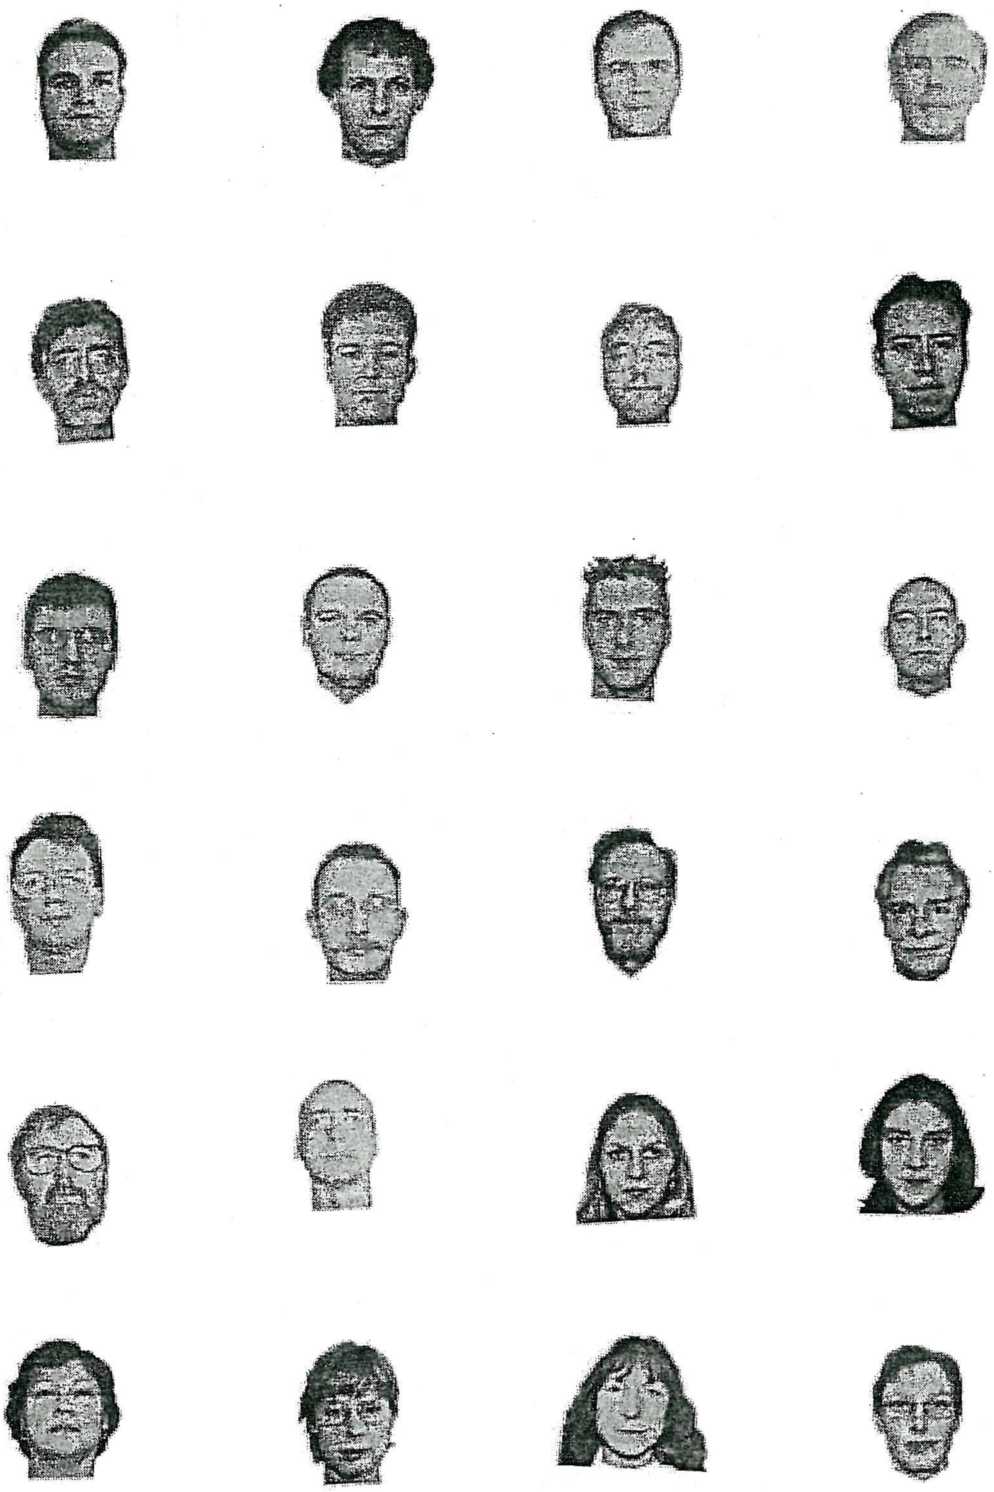

APPENDIX C

Full set of unfamiliar stimuli

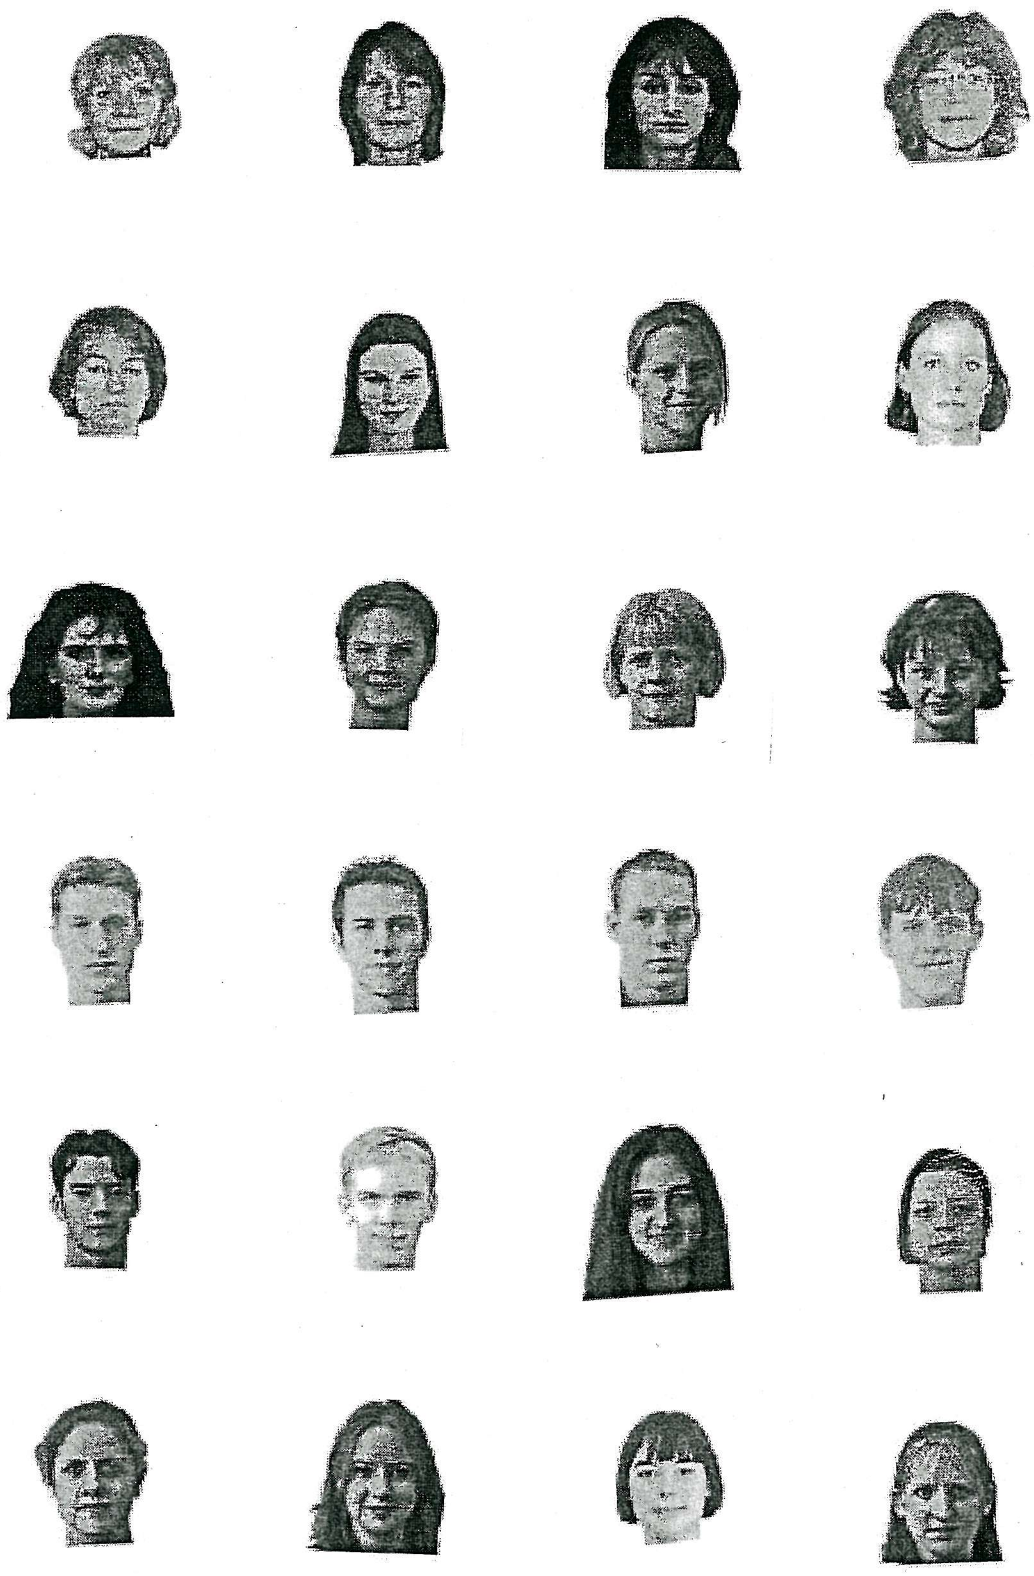

APPENDIX D

Full set of publicly familiar stimuli

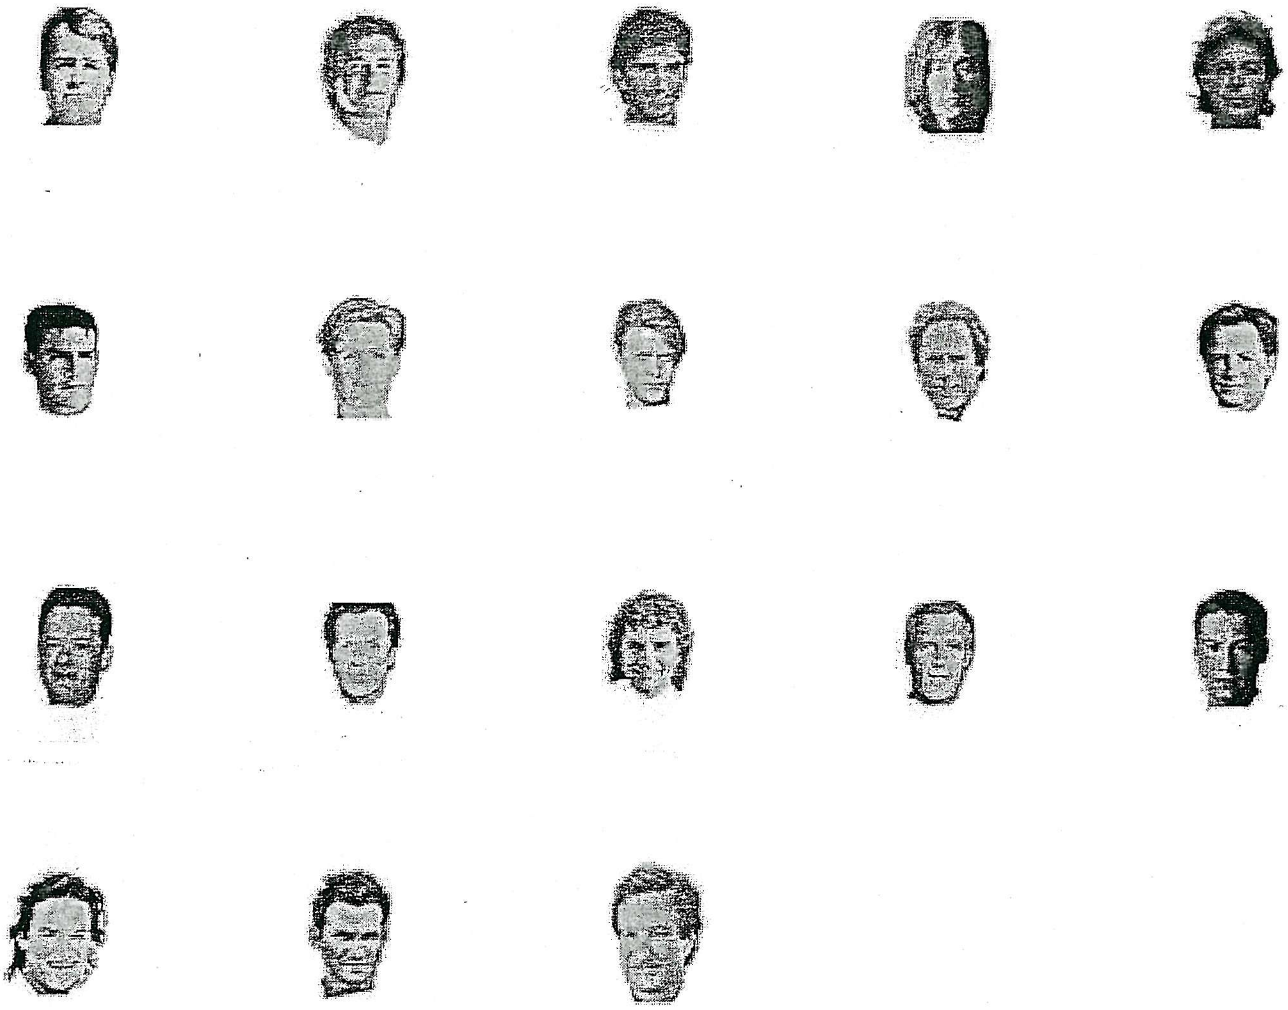

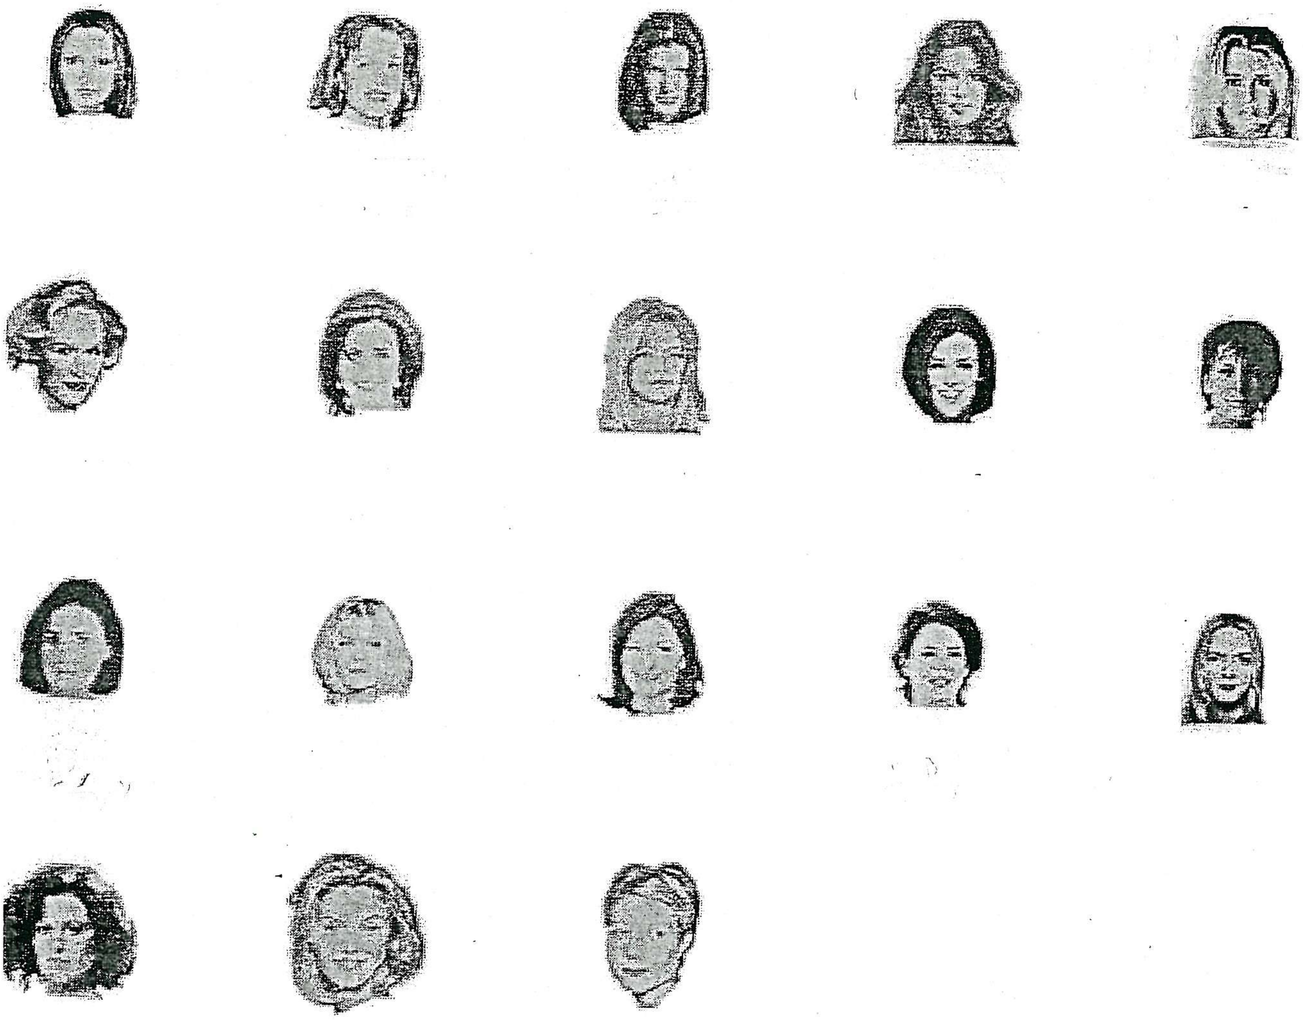

APPENDIX E

Negative stimuli used in Experiment Four

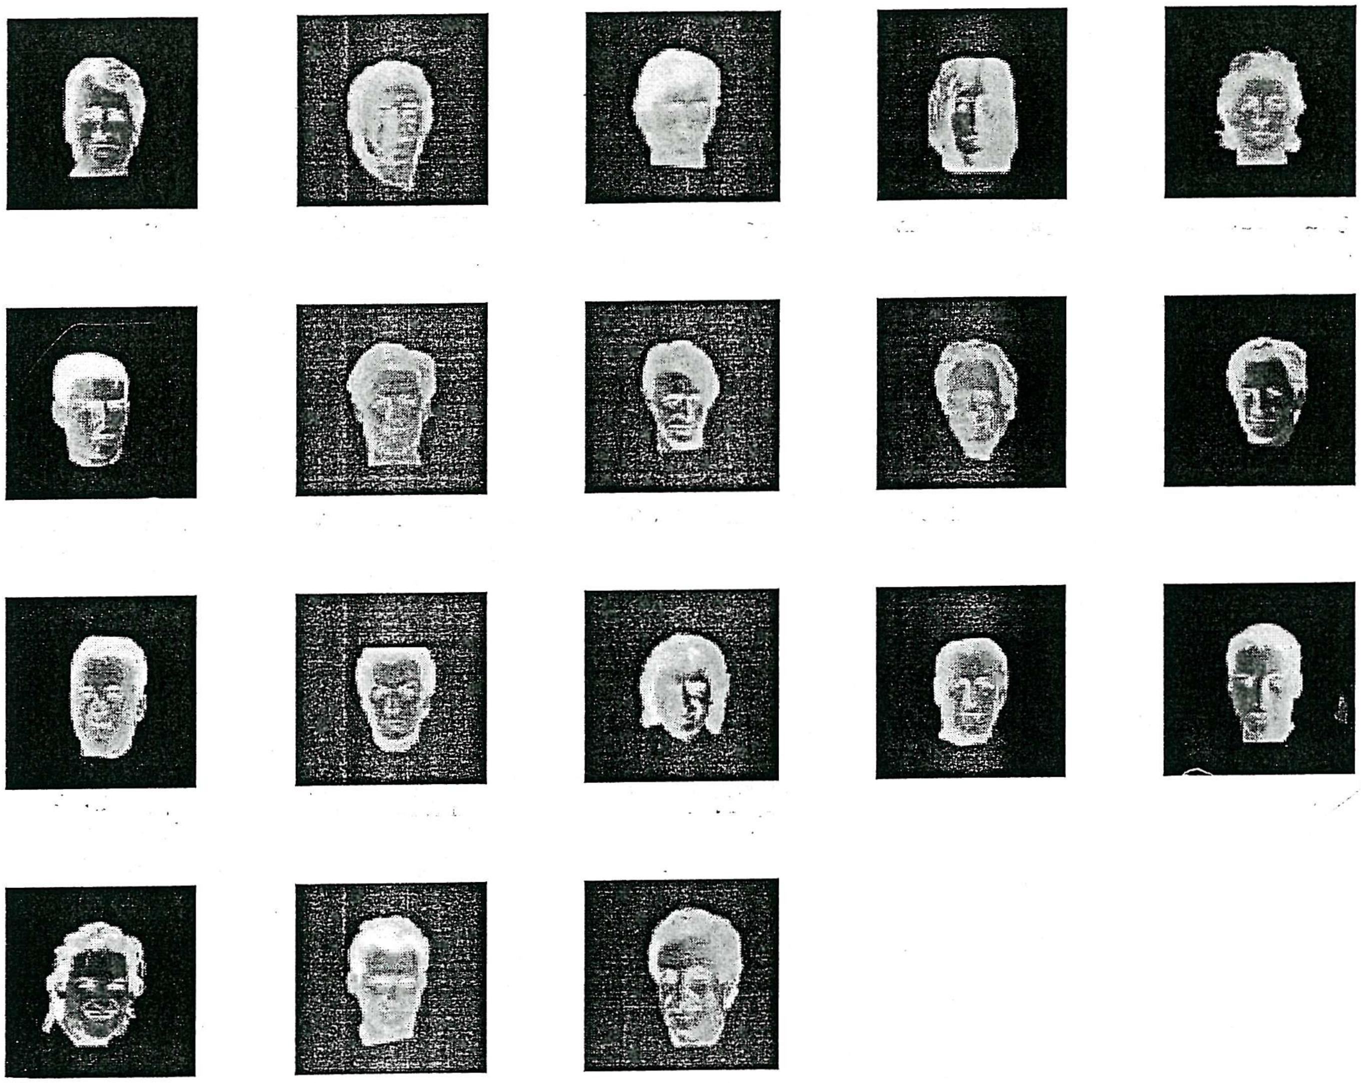

APPENDIX F

Thatcherised stimuli used in Experiments Five and Six

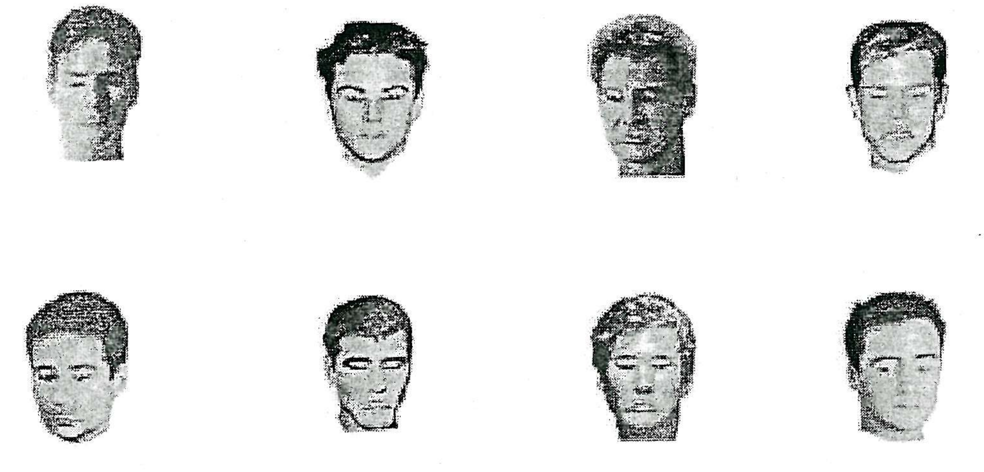

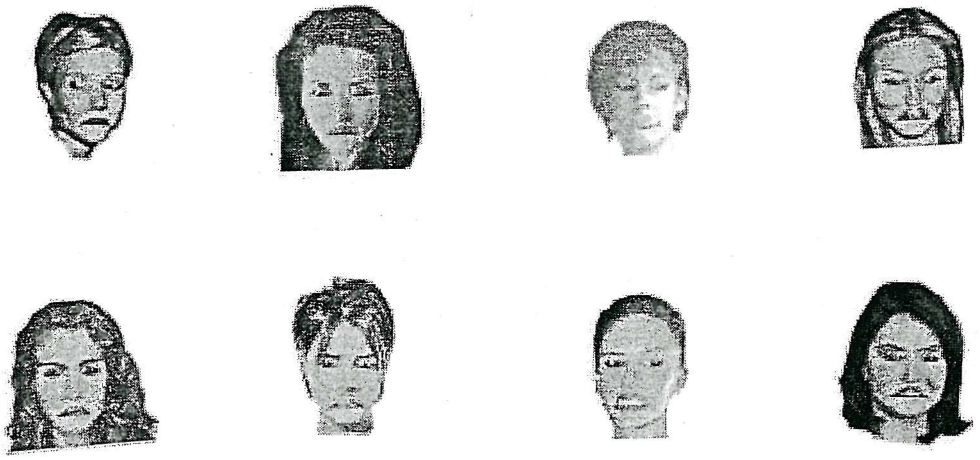

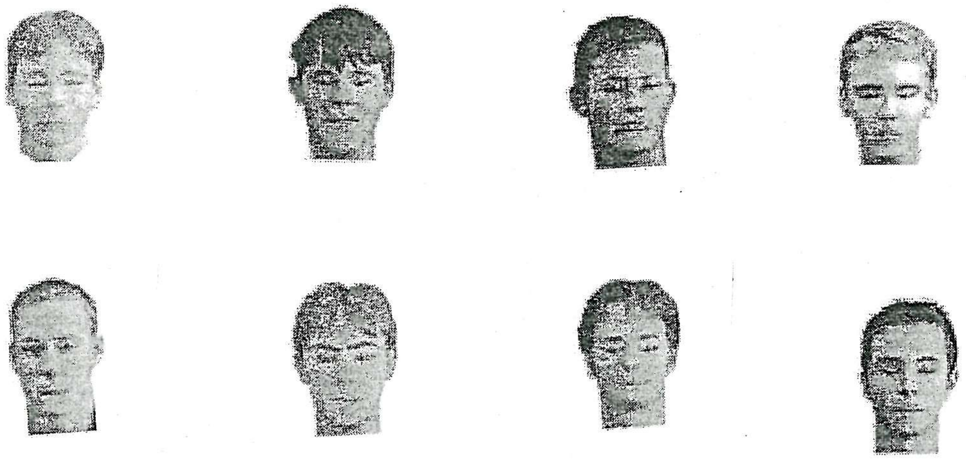

APPENDIX G

Scrambled stimuli used in Experiment Seven

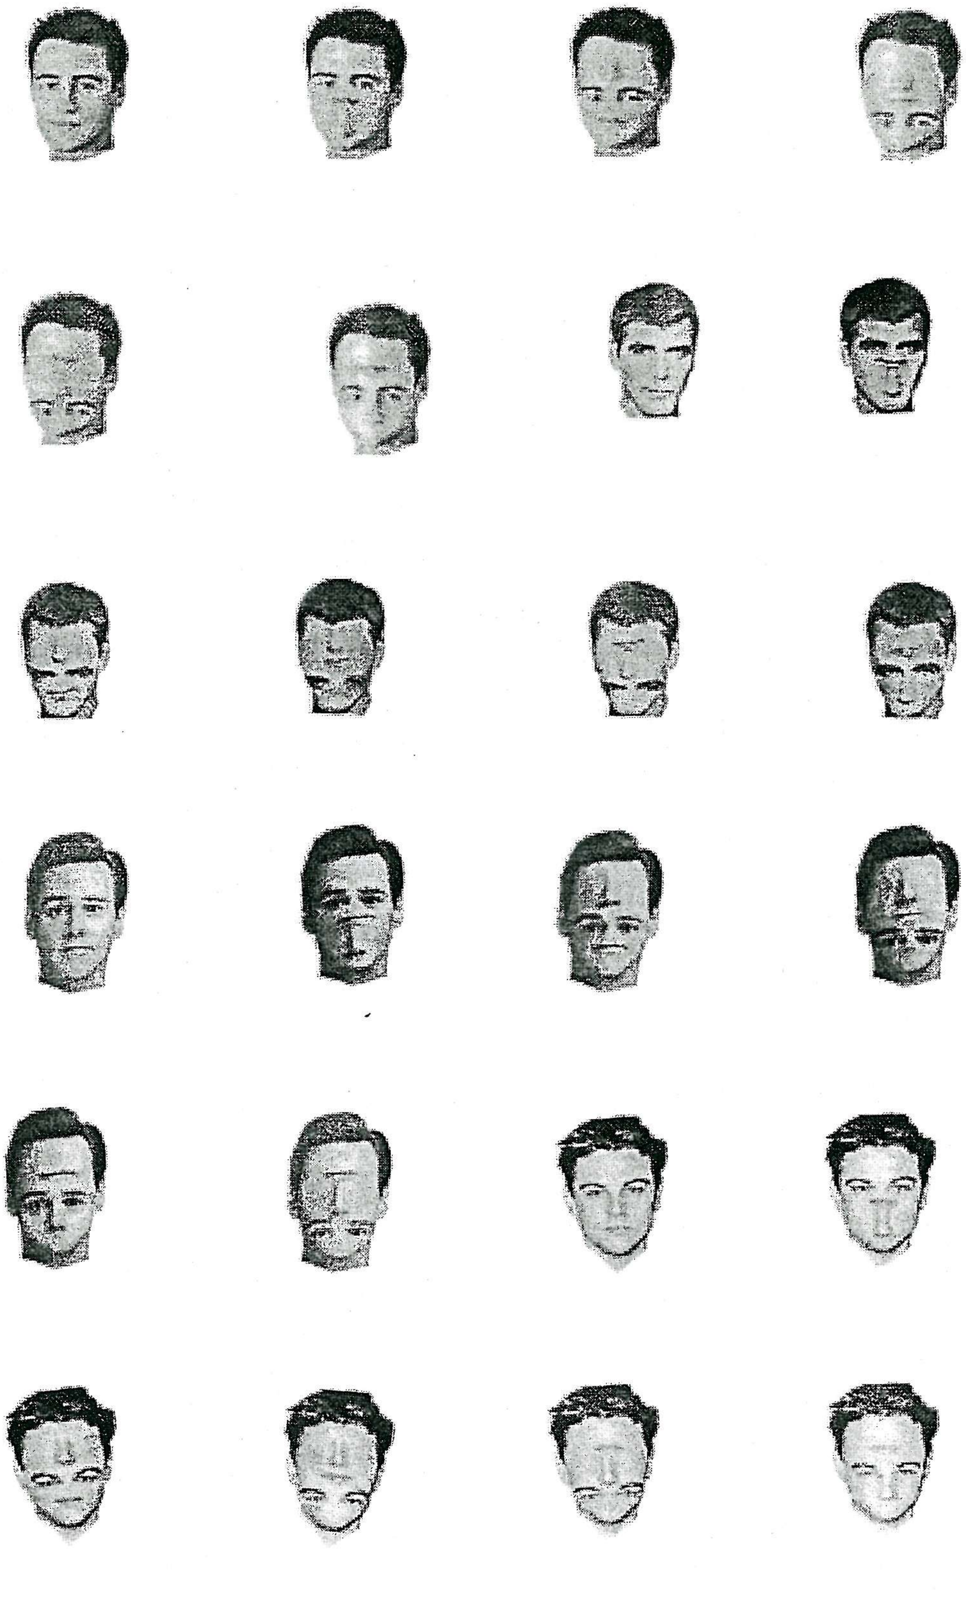

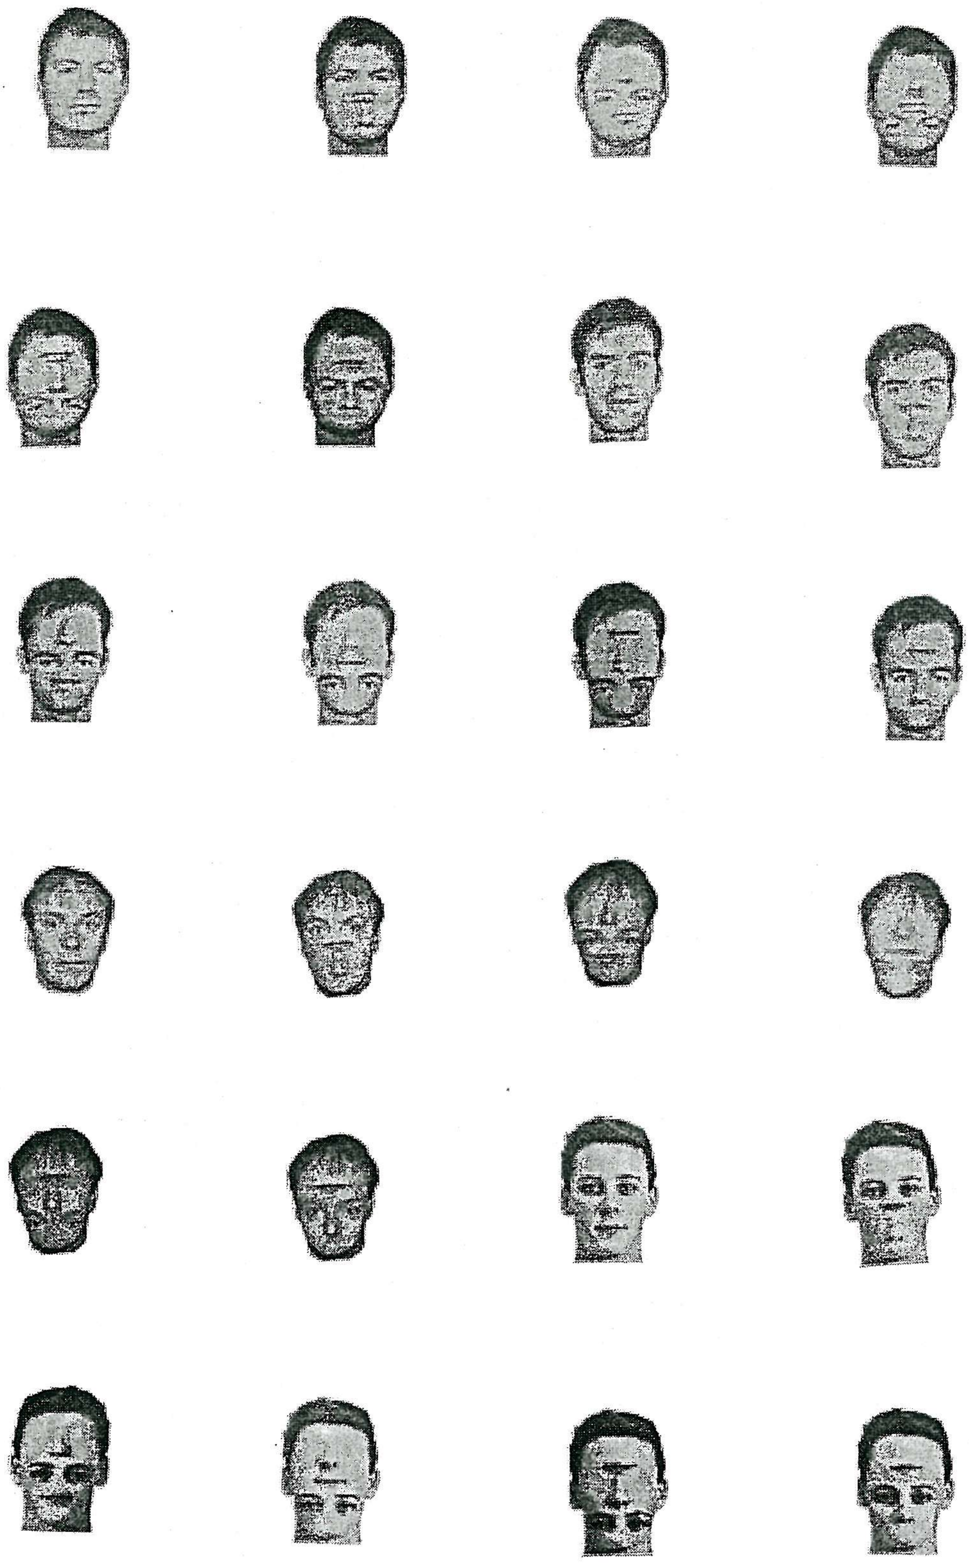

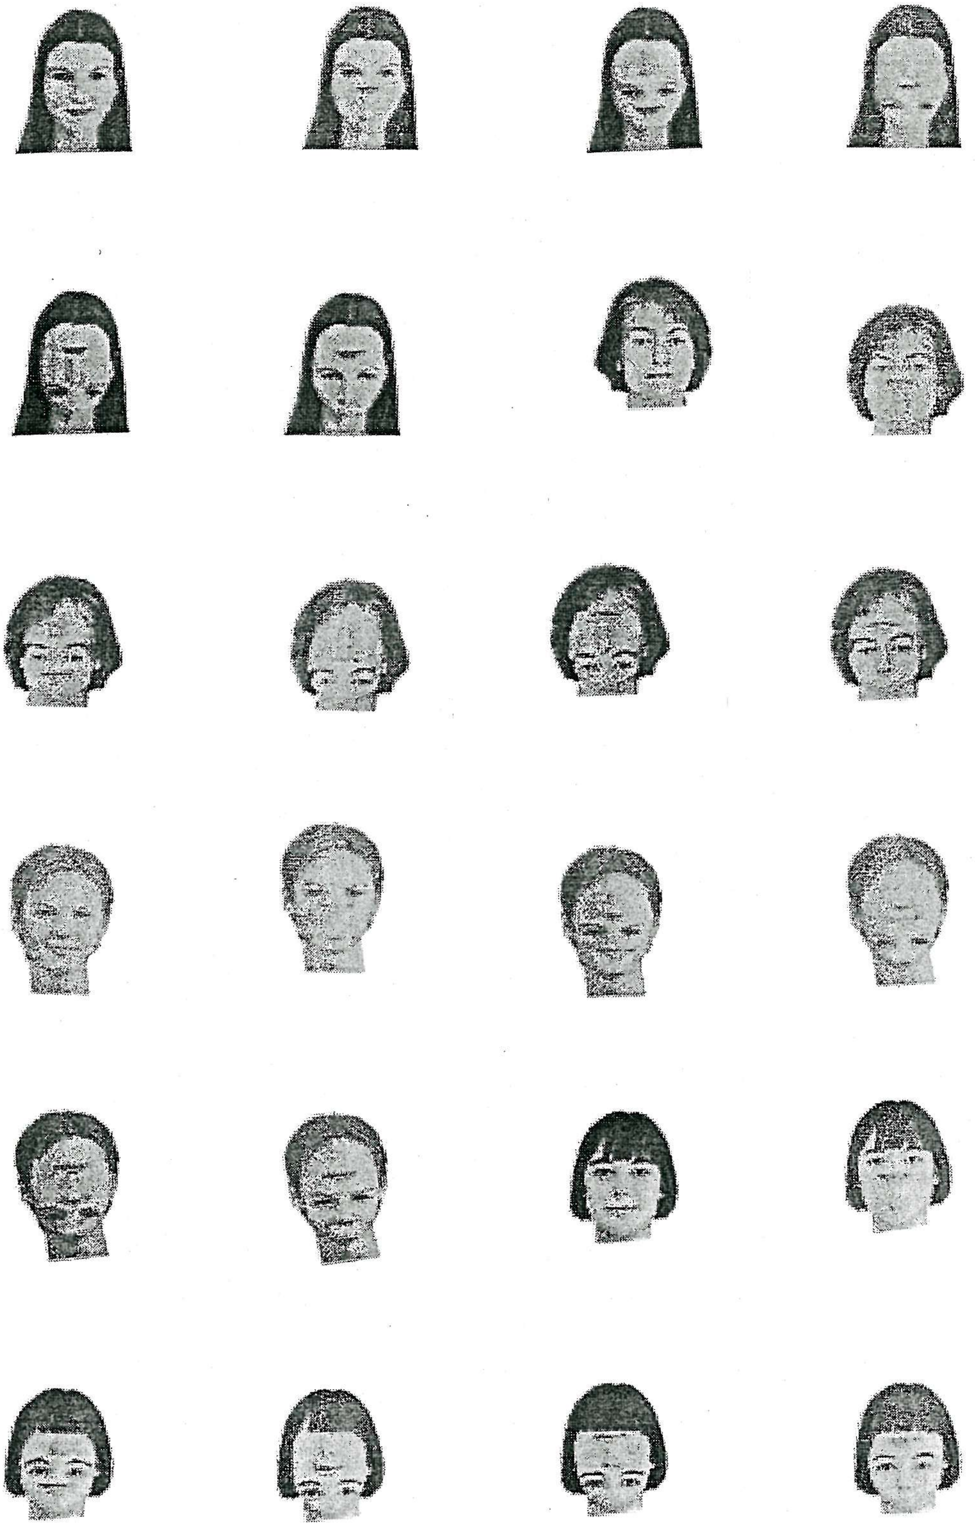

APPENDIX H

Formulae for A' and B''_D

H	FA	A'	Bd
0.5	0.5	0.5	0

A'

IF(A2>=B2,0.5+((A2-B2)*(1+A2-B2))/(4*A2*(1-B2)),0.5-((B2-A2)*(1+B2-A2))/(4*B2*(1-A2)))

B''_D

((1-A2)*(1-B2)-A2*B2)/((1-A2)*(1-B2)+A2*B2)

REFERENCES

Aha, D.W., & Goldstone, R.L. (1990) Learning attribute relevance in context in instance-based learning algorithms. *Proceedings of the Twelfth Annual Conference of the Cognitive Society*. pp 141-148, Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Allison, T.M., Puce, A., Spencer, D.D., & McCarthy, G. (1999). Electrophysiological studies of human face perception 1: Potentials generated in occipitotemporal cortex by face and non-face stimuli. *Cerebral Cortex, 9,* 415-430.

Archer, J., Hay, D.C., & Young, A.W. (1992). Face processing in psychiatric populations. *British Journal of Clinical Psychology, 31,* 45-61.

Baddeley, A.D., & Hitch, G.J. (1974). Working memory. *In* G. Bower (Ed.), *The Psychology of Learning and Motivation Vol. 8.* New York: Academic Press.

Bartlett, J.C., & Searcy, J. (1993). Inversion and Configuration of Faces. *Cognitive Psychology, 25,* 281-316.

Baudouin, J.Y., Sansone, S., & Tiberghien, G. (2000). Recognizing expression from familiar and unfamiliar faces. *Pragmatics and Cognition, 8 (1),* 123-146.

Bauer, R.M. (1984). Autonomic recognition of names and faces in prosopagnosia: A neuropsychological application of the guilty knowledge test. *Neuropsychologia, 22,* 457-469.

Beale, J., & Keil, F. (1995). Categorical effects in the perception of faces. *Cognition, 57,* 217-239.

Benson, P.J., & Perrett, D.I. (1991). Perception and recognition of photographic quality facial caricatures: Implications for the recognition of natural images. *European Journal of Psychology, 3,* 105 – 135.

Benson, P.J., & Perrett, D.I. (1994). Visual processing of facial distinctiveness. *Perception, 23,* 75-93.

Bentin, S. Allison, T., Puce, A., Perez, A., & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. *Journal of Cognitive Neuroscience, 8,* 551-565.

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

Blanc-Garin, J. (1986). Faces and non-faces in prosopagnosic patients. *In* H.D. Ellis, M.A. Jeeves, F. Newcombe & A. Young (Eds.) *Aspects of face processing,* pp. 273-278. Dordrecht, Martinus Nijhoff.

Bornstein, M.H., & Korda, N.O. (1984). Discrimination and matching within and between hues measured by reaction times: Some implications for categorical perception and levels of information processing. *Psychological Research, 46 (3),* 207-222.

Brennan, S.E. (1985). The caricature generator. *Leonardo, 18,* 170-178.

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. Per*ception, 15,* 387-397.

Bruce, V. (1988). *Recognising faces.* Lawrence Erlbaum, London.

Bruce, V., Carson, D., Burton, A.M., *&* Kelly, S. (1998). Prime time advertisements: Repetition priming from faces seen on subject recruitment posters. *Memory & Cognition, 26 (3),* 502-515.

Bruce, V., Doyle, T., Dench, N., & Burton, M. (1991). Remembering facial configurations. *Cognition, 38,* 109-144.

Bruce, V., Green, P.R., & Georgeson, M.A. (1997). *Visual Perception. Physiology, Psychology, and Ecology. 3rd Edition.* Psychology Press Ltd.: UK.

Bruce, V. & Langton, S. (1994). The use of pigmentation and shading information in recognising the sex and identities of faces. *Perception, 23,* 803-822.

Bruce, V. *&* Valentine, T. (1985). Identity priming in the recognition of familiar faces. *British Journal of Psychology, 76,* 363-383.

Bruce, V. *&* Young, A. (1986). Understanding face recognition. *British Journal of Psychology, 77,* 305-327.

Burton, A.M., Bruce, V., *&* Johnston, R.A. (1990). Understanding face recognition with an interactive activation and competition 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. *The Quarterly Journal of Experimental Psychology, 51A (3),* 475-483.

Cabeza, R. Bruce, V., Kato, T., *&* Oda, M. (1999). The prototype effect in face recognition: Extension and limits. *Memory & Cognition, 27 (1),* 139-151.

Campanella, S., Chrysocoos, AQ., *&* Bruyer, R. (2001). Categorical perception of facial gender information: Behavioural evidence and the face-space metaphor. *Visual Cognition, 8,* 237-262.

Campanella, S. Hanoteau, C., Depy, D., Rossion, B., Bruyer, R., Crommelinck, M., & Guerit, J.M. (2000). Right N170 modulation in a face discrimination task: An account for categorical perception of familiar faces. *Psychophysiology, 37,* 796-806.

Campbell, R., Coleman, M., Walker, J., Benson, P., Wallace, S., Michelotti, J., & Baron-Cohen, S. (1999). When does the inner-face advantage in familiar face recognition arise and why? *Visual Cognition, 6 (2),* 197-216.

Carey, S. (1992). Becoming a face expert. *Philosophical Transactions of the Royal Society of London B, 335,* 95-103.

Cavanagh, P. & Leclerc, Y.G. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 3-27.

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 (4),* 257-260.

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

Davidoff, J. & Donnelly, N. (1990). Object superiority: A comparison of complete and part probes. *Acta Psychologia, 73, 225-243.*

DeRenzi, E. (1986). Current issues in prosopagnosia. *In* H.D. Ellis, M.A. Jeeves, F. Newcombe & A. Young (Eds.) *Aspects of face processing* pp. 243-252. Dordrecht, Martinus Nijhoff.

deRoiste, A., & Bonnet, G. (1995). Are multiparous mothers more skilled than primiparous mothers in recognising their newborn? *The Irish Journal of Psychology, 16 (1),* 29-37.

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, (2),* 107-117.

Donnelly, N. & Davidoff, J. (1999). The mental representations of faces and houses: Issues concerning parts and wholes. *Visual Cognition, 6 (3/4),* 319-343.

Donnelly, N., Humphreys, G.W., & Sawyer, J. (1994). Stimulus factors affecting the discrimination of faces. *Acta Psychologica, 5,* 219-234.

Dubois, S., Rossion, B., Schiltz, C., Bodart, J.M., Michel, C., Bruyer, R., & Crommelinck, M. (1999). Effect of familiarity on the processing of human faces. *NeuroImage, 9,* 278-289.

Eimer, M. (2000*a*). Event-related brain potentials distinguish processing stages involved in face perception and recognition. *Clinical Neuropsychology, 111,* 694-705.

Eimer, M. (2000*b*). 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., Burton, A.M., Young, A.W., & Flude, B.M. (1997). Repetition priming between parts and wholes: Tests of a computational model of familiar face processing. *British Journal of Psychology, 88,* 579-608.

Ellis, H., Shepherd, J., & Davies, G. (1979). Identification of familiar and unfamiliar faces from internal and external features: Some implications for theories of face recognition. *Perception, 8,* 431-439.

Ellis, A.W., Young, A.W., & Hay, D.C. (1987). Modelling the recognition of faces and words. In (ed. P.E. Morris) *Modelling Cognition,* pp 269-297. Wiley, Chichester.

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 & Young (Eds.), *Aspects of face processing.* pp 53-58. Dordrecht: Martinus Nijhoff.

Enns, J.T. & Shore, D.I. (1997). Separate influences of orientation and lighting in the inverted-face effect. *Perception and Psychophysics, 59,* 29-31.

Etcoff, N.L., & Magee, J.J. (1992). Categorical perception of facial expressions. *Cognition, 44,* 227-240.

Farah, M. J., Tanaka, J. W., & Drain, H. M. (1995). What Causes the Face Inversion Effect? *Journal of Experimental Psychology: Human Perception and Performance, 21, (3),* 628-634.

Flude, B.M., Ellis, A.W. & Kay, J. (1989). Face processing and name retrieval in an anomic aphasic: Names are stored separately from semantic information about familiar people. *Brain & Cognition, 11,* 60-72.

Galper, R.E., & Hochberg, J. (1971). Recognition memory for photographs of faces. *American Journal of Psychology, 84,* 351-354.

Gauthier, I,. Skudlarski, P., Gore. J., & Anderson, A. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience, 3 (2),* 191-197.

Gauthier, I., & Tarr, M. J. (1997). Becoming a 'greeble' expert: Exploring the face recognition mechanisms. *Vision Research, 37,* 1673-1682.

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

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 (6),* 574-580.

Goldstein, A.G. (1965). Learning of inverted and normally oriented faces in adults and children. *Psychonomic Science, 3,* 447-448.

Goldstein, A. G., & Chance, J. E. (1981). Laboratory studies in face recognition. *In* G. Davies, H. Ellis, & J. Shepherd, (Eds.) *Perceiving and Remembering Faces.* London. Academic Press.

Goren, C. C., Sarty, M., & Wu, P. Y. K. (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Pediatrics, 56 (4),* 544-549.

Harnad, S. (1987). *Categorical Perception: The Groundwork of Cognition.* Cambridge UK: Cambridge University Press.

Hasselmo, M.E., Rolls, E.T., & Baylis, G.C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behavioural Brain Research, 32,* 203-128.

Hay, D. C. & Young, A. W. (1982). The human face. *In* A.W. Ellis (Ed.), *Normality and Pathology in Cognitive Functions.* pp. 173-202. Academic Press, London.

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

Hécaen, H. & Angelergues, R. (1962). Agnosia for faces (prosopagnosia). *Archives of Neurology, 7,* 92-100.

Hill, H. & Bruce, V. (1996). Effects of lighting on the perception of facial surfaces. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 986-1004.

Hill, H, Schyns, P., & Shigeru, A. (1997). Information and viewpoint dependence in face recognition. *Cognition, 62,* 201-222.

Hines, D., & Braun, J. (1990). Order of feature recognition in familiar and unfamiliar faces. *Brain & Cognition, 14,* 165-184.

Hochberg, J. & Galper, R.G. (1967). Recognition of faces: 1. An exploratory study. *Psychonomic Science, 9,* 619-620.

Hole, G.J., George, P., & Dunsmore, V. (1999). Evidence for holistic processing of faces viewed as photographic negatives. *Perception, 28 (3),* 341-359.

Homa, D., Goldhardt, B., Burruel-Homa, L., & Carson Smith, J. (1993). Influence of manipulated category knowledge on prototype classification and recognition. *Memory & Cognition, 21,* 529-538.

Homa, D., Haver, B., & Schwartz, T. (1976). Perceptibility of schematic face stimuli: Evidence for a perceptual gestalt. *Memory & Cognition, 4,* 176-185.

Jacoby, L.L., Kelley, C.M., & Dywan, J. (1989). Memory attributions. *In* H.L. Roediger & F.I.M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honor of Endel Tulving.* Hillsdale, New Jersey: Erlbaum.

Johnston, A., Hill, H., & Carman, N. (1992). Recognising faces: Effects of lighting direction, inversion and brightness reversal. *Perception, 21,* 365-375.

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

Kendrick, K. M., & Baldwin, B. A., (1987). Cells in the temporal cortex of conscious sheep can respond preferentially to the sight of faces. *Science, 238,* 448-450.

Kendrick, K. M., Da Costa, A. P., Leigh, A. E., Hinton, M. R., & Pierce, J. W. (2001). Sheep don't forget a face. *Nature, 414,* 165-166.

Klatzky, R., & Forrest, F. (1984). Recognising familiar and unfamiliar faces. *Memory & Cognition, 12 (1),* 60-70.

Kohler, W. (1940). *Dynamics in Psychology*. New York: Liveright.

Kreiman, G., Koch, C., & Fried, I. (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. *Nature Neuroscience, 3 (9),* 946-953.

Lawson, R. & Jolicoeur, P. (1998). The effects of plane rotation on the recognition of brief masked pictures of familiar objects. *Memory & Cognition, 26 (4),* 791-803.

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

Levin, D.T., & Beale, J.M. (2000). Categorical perception occurs in newly learned faces, other-race faces, and inverted faces. *Perception & Psychophysics, 62 (2),* 386-401.

Lewis, M.B. & Johnston, R.A. (1997). The Thatcher illusion as a test of configural disruption. *Perception, 26,* 225-277

Marr, D. & Nishihara, K. (1978). Representation and recognition of the spatial organisation of three-dimensional shapes. *Proceedings of the Royal Society of London, B200*, 269-294.

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

Meissner, C.A. & Brigham, J.C. (2001). Thirty years of investigating the own-race bias in memory for faces. *Psychology, Public Policy, and Law, 7 (1),* 3-35.

Mermelstein, R., Banks, W., & Prinzmetal, W. (1979). Figural goodness effects in perception and memory. *Perception & Psychophysics, 26 (6),* 472-480.

Newell, F. N. (1998). Stimulus context and view dependence in object recognition. *Perception, 27,* 47-68.

Newell, F. N., Chiroro, P., & Valentine, T. (1999). Recognising unfamiliar faces: The effects of distinctiveness and view. *The Quarterly Journal of Experimental Psychology, 52A (2),* 509-534.

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39-57.

O'Donnell, C. & Bruce, V. (2001). Familiarisation with faces selectively enhances sensitivity to changes made to the eyes. *Perception, 30,* 755-764.

Paller, K. A., Gonsalves, B., Grabowecky, M., Bozic, V. S., & Yamada, S. (2000). Electrophysiological correlates of recollecting faces of known and unknown individuals. *Neuroimage, 11,* 98-110.

Palmer, S., Rosch, E., & Chase. P. (1981). Canonical perspective and the perception of objects. *In* H. Long & A.D. Baddeley (Eds.), *Attention and Performance IX.* New Jersey: Lawrence Erlbaum Associates Inc. pp 135-151.

Parks, T.E. (1983). Letter to the editor. *Perception, 12,* 88.

Parks, T.E., Coss, R.G., & Coss, C.S. (1985) Thatcher and the Cheshire cat: Context and the processing of facial features. *Perception, 14,* 747-754.

Payne, J.W., Bettman, J.R., & Johnson, E.J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology, 43, 87-131.*

Rakover, S.S. (1999). Thompson's Margaret Thatcher illusion: When inversion fails. *Perception, 28,* 1227-1230.

Rakover, S.S. & Teucher, B. (1997). Facial inversion effects: Parts and whole relationship. *Perception & Psychophysics, 59 (5),* 752-761.

Rhodes, G., Brake, S. & Atkinson, A. (1993). What's lost in inverted faces? *Cognition, 47*, 25-37.

Rhodes, G., Brake, S., Taylor, K & Tan, S. (1989). Expertise and configural coding in face recognition. *British Journal of Psychology, 80*, 313-331.

Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representations of faces. *Cognitive Psychology, 190,* 473-497.

Rhodes, G., & McLean, I. (1990). Distinctiveness and expertise effects with Homogenous stimuli: Towards a model of configural coding. *Perception, 19*, 773-794.

Rhodes, G. & Moody, J. (1990). Memory representations of unfamiliar faces: Coding of distinctive information. *New Zealand Journal of Psychology, 19,* 70-78.

Roberts, T. & Bruce, V. (1988). Feature saliency in judging the sex and familiarity of faces. *Perception, 17,* 475-481.

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

Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382-439.

Rumelhart, D. E. & McClelland, J. L. (1981). An interactive activation model of the effect of context in perception, Part 1. An account of basic findings. *Psychological Review, 88, 375-406.*

Ruthruff, E. & Miller, J. (1995). Can mental rotation begin before perception finishes? *Memory & Cognition, 23 (4),* 408-424.

Scapinello, K. & Yarmey, D. (1970). The role of familiarity and orientation in Immediate and delayed recognition of pictorial stimuli. *Psychonomic Science, 21, (6),* 329-331.

Schneider, W. & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review, 84,* 1-66.

Searcy, J. & Bartlett, J. (1996). Inversion & processing of component and spatial-relational information in faces. *Journal of Experimental Psychology: Human Perception & Performance, 22 (4),* 904-915.

Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology, 75,* 221-242.

Smith, E.E., & Medin, D.L. (1981). *Categories and Concepts.* Cambridge, MA: Harvard University Press.

Solso, R.L. & McCarthy, J.E. (1981). Prototype formation of faces: A case of pseudo-memory. *British Journal of Psychology, 72,* 499-503.

Stevenage, S. (1995*a*). Can caricatures really produce distinctiveness effects? *British Journal of Psychology, 86,* 127-146.

Stevenage, S. (1995*b*). Demonstration of a caricature advantage in children. *Cahiers de Psychologie Cognitive, 14,* 325-341.

Stevenage, S. (in prep). The effect of familiarity on within and cross-modality repetition priming.

Stürzel, F., & Spillman, L. (2000). Thatcher illusion: Dependence on angle of rotation. *Perception, 29,* 937-942.

Schweinberger, S.R., Burton, A.M., & Kelly, S.W. (2001). Priming the access to names of famous faces, *British Journal Of Psychology, 92*, 303-317.

Tanaka, J., & Farah, M. (1993). Parts and wholes in faces recognition. *Quarterly Journal of Experimental Psychology, 46A*, 225-245.

Tanaka, J., & Gauthier, I. (1997). Expertise in object and face recognition. *The Psychology of Learning and Motivation,* 36, 83-125.

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

Tanaka, J.W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology, 23,* 457-483.

Tarr, M.J. & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience, 3 (8),* 764-769.

Thompson, P. (1980). "Margaret Thatcher: a new illusion" *Perception, 9*, 483-484.

Tong, F. & Nakayama, K. (1999). Robust representations for faces: Evidence from visual search. *Journal of Experimental Psychology: Human Perception & Performance, 25 (4),* 1016-1035.

Troje, N.F., & Kersten, D. (1999). Viewpoint-dependent recognition of familiar faces. *Perception, 28*, 483-487.

Turk, M., & Pentland, A. (1991) Eigenfaces for recognition. *Journal of Cognitive Neuroscience, 3 (1)*, 71-86.

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

Valentine, T. (1991).A unified account of the effects of distinctiveness, inversion and race in face recognition. . *Quarterly Journal of Experimental Psychology, 43A (2)*, 161-204.

Valentine, T. & Bruce, V. (1986*a*). The effect of race, inversion and encoding activity upon face recognition. *Acta Psychologica, 61*, 259-273.

Valentine, T. & Bruce, V. (1986*b*). The effects of distinctiveness in recognising and classifying faces. *Perception, 15*, 525-535.

Whittlesea, B.W.A., & Williams, L.D. (1998). Why do strangers feel familiar, but friends don't? A discrepancy-attribution account of feelings of familiarity. *Acta Psychologica*, 98, 141-165.

Williamson, P.D., Thadani, V.M., Darcey, T.M., Spencer, D.D., Spencer, S.S., & Mattson, R.H. (1992). Occipital lobe epilepsy: Clinical characteristics, seizure spread patterns, and results of surgery. *Annals of Neurology, 31*, 3-13.

Wiser, A., Andreason, N., O'Leary, D., Crespo-Facorro, B., Boles-Panto, L., Watkins, L., & Hichwa, R. (2000). Novel vs. well-learned memory for faces: A positron emission tomography study. *Journal of Cognitive Neuroscience, 12 (2)*, 255-266.

Yarmey, A.D., (1971). Recognition memory for familiar 'public' faces: Effects of orientation and delay. *Psychonomic Science, 24 (6)*, 286-288.

Yarmey, A.D., (1973). I recognize your face but I can't remember your name: Further evidence on the tip-of-the-tongue phenomenon. *Memory & Cognition, 1 (3),* 287-290.

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

Yin, R.K. (1970). Face recognition: A dissociable ability? *Neuropsychologia, 8,* 395-402.

Young, A.W. (1984). Right cerebral superiority for recognising the internal and external features of famous faces. *British Journal of Psychology, 75*, 161-169.

Young, A.W., & Bruce, V. (1991). Perceptual categories and the computation of 'Grandmother'. *European Journal of Cognitive Psychology, 3 (1),* 5-49.

Young, A.W., Hellawell, D., & De Haan, E.H.F. (1988). Cross-domain semantic priming in normal subjects and a prosopagnosic patient. *Quarterly Journal of Experimental Psychology, 40A*, 561-580.

Young, A.W., Hay, D.C., & Ellis, A.W., (1985). The faces that launched a thousand slips: Everyday difficulties and errors in recognising people. *British Journal of Psychology, 76*, 495-523.

Young, A.W., Hay, D.C., McWeeny, K., Flude, B., & Ellis, A.W., (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception, 14,* 737-746.

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

Young, A.W., McWeeny, K.H., Hay, D.C., & Ellis, A.W. (1986*a*). Matching familiar and unfamiliar faces on identity and expression. . *Psychological Research, 48*, 63-68.