

The effect of rotation in face processing

Andrew J. Edmonds

Supervisors: Michael B. Lewis and Tom Freeman

**Thesis submitted to Cardiff University, School of Psychology
for the Degree of Doctor of Philosophy, December 2006.**

UMI Number: U584188

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U584188

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

| | |
|---------------------------|----------------|
| STUDENT ID NUMBER: | 0226616 |
|---------------------------|----------------|

SUMMARY OF THESIS:

Inversion has been shown to disrupt face recognition, but relatively few studies have looked at the processes involved in the recognition of faces seen at intermediate angles of rotation. Here we address this issue by looking at the processing of rotated faces in a recognition memory task thought to encourage holistic processing, and in the matching of Thatcherised faces, a task which is seen as an indicator of configural processing.

When faces were equally rotated at learning and test, we found no evidence of holistic processing. When the task was to recognise differently oriented faces, however, performance declined as an approximately linear function of the difference in orientation between the learning and test faces, suggesting that participants may be mentally rotating faces prior to recognition.

Chapter 4 considered the effects of rotation on a same-different matching task thought to encourage configural processing. When the task required the matching of local configural information, the effect of rotation was approximately equal for normal and Thatcherised faces, but Thatcherisation disproportionately disrupted global configural information. The effect of rotation on these forms of information when face pairs contained identical images of the same person, and in an identity-matching task, was also explored.

Chapter 6 looks at the effects of inversion on the detection of configural and featural changes to faces in a visual search task. Similar effects of inversion and search strategies were observed for both types of change at both angles of orientation, suggesting that face processing mechanisms do not extract configural at the expense of featural information from faces in this task.

The implications of these findings for theories of face processing and the nature of the relationship between rotation and face processing are discussed, and the extent to which the mental rotation hypothesis can account for these findings is also considered.

Contents

| | |
|---|----|
| Chapter 1: Introduction..... | 3 |
| 1.1.1. Definitions and History..... | 4 |
| 1.1.2. Theories and Models: The Dual-Mode Hypothesis..... | 6 |
| 1.1.3. Theories and Models: The Holistic Processing Hypothesis | 10 |
| 1.1.4. Spatial Positioning of Featural Information & Theories of Face Processing | 14 |
| 1.1.5. Summary | 22 |
| 1.2. Rotation and Face Perception | 23 |
| 1.2.1. Mental Rotation | 23 |
| 1.2.3. Isolating Configural Processing in Face Perception | 26 |
| 1.2.4. Rotation and the Holistic View of Face Processing | 29 |
| 1.2.5. Summary | 31 |
| 1.3. Summary of Chapter 1 | 31 |
| Chapter 2: Detecting configural and featural changes in a visual search task: evidence for the face inversion effect | 35 |
| 2.1. Experiment 1 | 35 |
| 2.1.1. Method..... | 40 |
| 2.1.2. Results..... | 43 |
| 2.1.3. Discussion..... | 48 |
| Chapter 3: Parts and Wholes in Face Recognition: The Effects of Rotation..... | 50 |
| Overview..... | 50 |
| 3.1. Experiment 2..... | 50 |
| 3.1.1. Method..... | 53 |

| | |
|---|-----|
| 3.1.2. Results and Discussion | 56 |
| 3.2.1. Method | 62 |
| 3.2.3. General Discussion | 65 |
| Chapter 4: Examining the Effect of Rotation on Configural Processing in a Face Matching Task | 69 |
| Overview..... | 69 |
| 4.1. Experiment 4..... | 69 |
| 4.1.1. Method..... | 76 |
| 4.1.2. Results..... | 78 |
| 4.1.3. Discussion..... | 82 |
| 4.2. Experiment 5..... | 85 |
| 4.2.1. Method..... | 86 |
| 4.2.2. Results..... | 88 |
| 4.2.3. Discussion..... | 97 |
| 4.3. General Discussion | 98 |
| Chapter 5: Face rotation in a matching task: exploring the effects of inversion, the similarity of the image and identity matching | 102 |
| Overview..... | 102 |
| 5.1. Experiment 6..... | 102 |
| 5.1.1. Method..... | 105 |
| 5.1.2. Results..... | 106 |
| 5.1.3. Discussion..... | 112 |
| 5.2. Experiment 7..... | 114 |
| 5.2.1. Method..... | 117 |
| 5.2.2. Results..... | 119 |

| | |
|--|-----|
| 5.2.3. Discussion..... | 125 |
| 5.3. Experiment 8..... | 127 |
| 5.3.1. Method..... | 129 |
| 5.3.2. Results..... | 131 |
| 5.3.3. Discussion..... | 135 |
| 5.4. General Discussion | 136 |
| Chapter 6: Conclusions and Directions for Future Research..... | 141 |
| 6.1.1. The Part-Whole Paradigm | 142 |
| 6.1.2. Face Matching Tasks: Simultaneous or Sequential? | 148 |
| 6.1.3. Visual Search | 154 |
| 6.3. Future Research | 170 |
| 7.2. General Conclusions..... | 172 |
| References..... | 173 |

Abstract

Inversion has been shown to disrupt the processing and recognition of faces, and a vast number of studies have been conducted in an attempt to understand how and why this effect occurs. Relatively few studies, however have looked at the processes involved in the recognition or perception of faces seen at intermediate angles of rotation. In this thesis, we have attempted to address this need for a clearer understanding of these effects by looking at the processing of rotated faces in a face recognition memory task using the part-whole paradigm, which is thought to encourage holistic processing, and in the matching of Thatcherised faces, which is seen as an indicator of configural processing.

Chapter 2 looked at the effects of inversion on the detection of configural and featural changes to faces in a visual search task, to investigate the separate effects of inversion on configural and featural information. We found some evidence of a greater detrimental effect of inversion for configural than featural changes in this experiment, suggesting that there may be a role of configural processing in the detection of featural changes. These findings, although only observed in the error data in this experiment, nevertheless appear to provide some support for the dual-mode hypothesis of face processing.

When holistic processing was encouraged in a face recognition memory task, for faces which were equally rotated at learning and test (Chapter 3), there was no effect of inversion for whole faces, and no difference between whole faces and isolated features, suggesting that rotated faces may not be processed holistically in

this task. Consistent with both the dual-mode and holistic hypotheses, the recognition of differently oriented faces in a second experiment, however, declined as an approximately linear function of the difference in orientation between the learning and test faces, suggesting that participants may be mentally rotating faces prior to recognition.

In Chapter 4 we considered the effects of rotation on a same-different matching task with normal and Thatcherised faces. Our findings suggest that when the task requires the matching of featural information, the effect of rotation is approximately equal for normal and Thatcherised faces, but Thatcherisation disproportionately disrupts matching performance when configural information is required. Moreover, the experiments presented in Chapter 5 show that the effect of rotation is independent of face type only when the face pairs contain identical images of the same person, while the information required for the matching of facial identity, and the way this information is affected by rotation and Thatcherisation, is dependent on the decision to be made ('same' or 'different').

In Chapter 6, the implications of these findings for the dual-mode and holistic theories of face processing and for the nature of the relationship between rotation and face processing are discussed. The extent to which the mental rotation hypothesis can account for these findings is also considered, along with a number of methodological issues and avenues for further research.

Chapter 1: Introduction

1.1. Configural and Featural Information in Face Processing

Faces are of great social and biological importance, and contain a wealth of information which humans use to make judgements about the age, gender, race and identity of a person, as well as subjective judgements of attractiveness and/or distinctiveness, for example. Our understanding of how we process faces, particularly in terms of how we recognise that a face is (or is not) that of a known individual, has developed considerably over the course of several decades, and a number of models have been proposed to account for a wide range of effects and findings within the face processing literature (Bruce & Young, 1986; Burton, Bruce & Johnston, 1990; Burton, Bruce & Hancock, 1999).

Arguably the most dominant idea during this time has been that there are two key modes of processing involved in the encoding of facial information and development of face representations which enable such accurate recognition. While it is generally accepted that information about individual features, such as a persons eyes, nose and mouth accounts for one of these modes of processing, researchers have disagreed over the exact nature and definition of non-componential information which is also important in face perception (e.g. Tanaka & Farah, 1993; Searcy & Bartlett, 1996). In this thesis, we define and consider the importance of two forms of processing, configural and holistic, in a variety of face processing tasks. Moreover, we look at the effects of rotation on tasks which are thought to encourage each of these types of processing in an attempt to gain a broader understanding of how these

forms of information are affected by face inversion, and what processes are involved in producing this effect.

1.1.1. Definitions and History

Over many years, authors have used terms such as second-order relational information, configural and holistic information to refer to a number of definitions of non-componential information, such as configural information being the combination of components that make up an individual face (e.g. Sergent, 1984), or the configuration formed by the individual arrangement of facial features (e.g. Bartlett & Searcy, 1993; Diamond & Carey, 1986). These distinctions have been developed in order to understand arguably the most well established and interesting finding in face perception, the face inversion effect; that is, the finding that inverting a face disrupts the recognition of that face, an effect which is disproportionate to that of inverting other objects, such as houses or aeroplanes (Yin, 1969).

Diamond & Carey's (1986) account has been an influential example of one such theory which distinguishes between two important types of information necessary for face processing. They drew a distinction between the first-order and second-order relational information of faces (and objects generally). They suggested that first-order relational information describes the basic configuration of an object, and consists of the spatial relations of the parts of that object and is used to classify an object (e.g. a face as a face). Faces and other objects which share a basic configuration, however, also have second-order relational information, which refers to information about variations in this basic configuration from a prototypical arrangement for that object. It is this second-order relational information which

enables us to identify individuals within a particular object class, in this case identifying a face as that of a particular known person (see also Rhodes, 1988). Diamond & Carey (1986) suggest that people must be experts with a class of stimuli if they are to be able to process second-order relational information. According to this view, then, the inversion effect may simply reflect a general lack of expertise with inverted faces, with inversion disrupting more complex configural processing. There is some evidence to support this expertise argument. Inversion effects of a similar size to that found for faces have also been observed when dog experts were asked to recognise different species of dogs (Diamond & Carey, 1986), while Gauthier & Tarr (2002) observed a similar effect when participants were trained, to expertise level, to recognise a set of previously novel objects. However, recent studies have failed to support a further prediction of the expertise account, that configural processing of inverted faces can be learnt (McKone et al, 2001, Robbins & McKone, 2003).

Tanaka & Farah (1991) tested Diamond & Carey's (1986) view that it is the second-order relational information of faces which is sensitive to inversion. Participants were taught to identify dot patterns which were either variations on a shared basic configuration (i.e. second-order relational information could be used) or had no shared configuration in common. Inversion effects were observed, but there was no difference between the two types of pattern, suggesting that second-order relational information is not more sensitive to inversion than other types of information. If Tanaka & Farah (1991) are to be believed, therefore, it would appear that Diamond & Carey's (1986) definition of second-order relational information may be too specific in encompassing variations from a shared, basic configuration, but, as we discuss further below, this was an influential account of face processing, as many

researchers have also argued that there are two types of information which are important for face perception.

Returning to other definitions of this non-componential information important for face processing, one model of face perception argues that faces exist within a multidimensional space, and are located at different points along configural and featural dimensions (Valentine, 1991) and, as such, configural information refers to information about the deviation of faces from the spatial average, while a more extreme proposal is that featural and configural information is encoded as an 'unparsed' whole (Farah, Wilson, Drain & Tanaka, 1998; Tanaka & Farah, 1993; Tanaka & Sengco, 1997), so the configural information here is the whole facial representation. Face perception tasks can therefore be completed, with varying degrees of success, depending on the extent to which facial representations can be broken down into their composite parts (Collishaw & Hole, 2000). We will discuss each of these theories in more detail later in this paper.

1.1.2. Theories and Models: The Dual-Mode Hypothesis

The different definitions of configural information have led to a number of accounts and theories as to how these different types of information are integrated in face processing, which, crucially, make a number of predictions as to how these types of information will be affected by face inversion. One such theory of face processing is the dual-mode account, which is characterised by the dominant view outlined above, that there is a distinction between two types of information processing, with one mode specialised for encoding configural information, and the other dealing with featural encoding. According to this view, configural and featural information are

processed independently, and are both available when a face is seen (Searcy & Bartlett, 1996), but, critically, the configural information processing mode is dependent upon the face being seen in an upright orientation. Inversion should therefore be of greater detriment to the encoding of configural information of a face than to featural information, which may remain largely intact when a face is inverted, according to this theory (Searcy & Bartlett, 1996).

The dual-mode hypothesis of face processing was developed in part to attempt to explain the Thatcher illusion (Thompson, 1980), where the eyes and mouth are inverted within an upright face, resulting in a grotesque image. The illusion is that if the whole image is then inverted, the changes made to the face are disrupted, and the image no longer appears abnormal or gruesome. Bartlett & Searcy (1993) provided support for the dual-mode hypothesis by testing a number of competing explanations of this illusion. They found that inversion reduced ratings of the grotesqueness of Thatcherised faces and spatially distorted faces (where the mouth was moved further down the face, or the eyes further apart, for example), but ratings of posed grotesque-expression faces remained much the same for upright and inverted faces. Similarly, spatially distorted faces were rated as being more similar to unaltered, original versions of that face when they were inverted than when they were upright, while again, inversion did not affect the perceived similarity of grotesque-expression faces. Bartlett & Searcy (1993) suggested that this was because the expression faces reflected changes in the individual features, not the spatial configuration of the features as in the Thatcherized and spatially distorted faces, and the Thatcher illusion can therefore be explained by the disruption of configural information, rather than by the disruption of facial expression, as suggested by Valentine (1991). The Thatcher

illusion has since become widely used as a research tool with which to investigate configural information and processing, and we shall review this research and a number of other explanations of this illusion in more detail later in this thesis.

Cabeza & Kato (2000) also provide evidence to suggest that configural information is disproportionately disrupted by inversion, compared to featural information. They used the prototype effect – the finding that, having studied a stimulus set, participants incorrectly judge a prototype (or composite) of the stimulus as being part of the original stimulus set. Prototypes emphasising either featural or configural processes were created and presented together with the studied faces and non-studied faces (and their prototypes). When shown upright, the studied and non-studied prototypes were equally likely to be incorrectly identified as faces they had previously seen, but when inverted, the configural prototypes were less likely to be mistaken for old (seen before) faces than featural prototypes, suggesting that inversion had a much more dramatic effect on configural information. Inversion had no effect on featural prototypes, however, and as such participants continued to think these were faces they had seen previously.

Similarly, Leder & Bruce (2000) showed that configural information provides important information for the processing of upright faces. In one experiment, participants learned a number of upright face identities, and were tested on their recognition memory for these faces in upright or inverted orientations. These faces were created such that they differed from each other only in terms of their featural information, their relational information, or a combination of relational and featural information. At test, only the faces differing in their relational information showed an

effect of inversion on recognition memory performance, suggesting that it is the processing of configural information which is disrupted by inversion.

An early finding from a different face perception paradigm also appears to provide some support for the dual-mode view. Young, Hellowell & Hay (1987) found that when the top half of a famous face was paired with the bottom half of a different face, recognition of the top face was considerably slower than if it was shown in isolation. This was only true, however, if the two face halves were joined together to make a 'new' face; if the halves were misaligned, recognition performance was at a similar level to that when the top half was presented in isolation. This difference between aligned and misaligned composite faces disappeared when the face composites were inverted, however. It has been argued that this effect occurs because in order to recognise the top half of an upright chimeric face, one needs to be able to ignore the bottom half, but the configural information processing involved in upright face perception makes this difficult. Inverting the whole face, however, disrupts this configural information and enables the face parts to be correctly identified. We note here, however, that recent studies have suggested that there may be more than one form or type of configural information which is useful for face processing. Carey & Diamond (1994) suggest that the chimeric face effect may reflect an early form of holistic encoding which is able to recognise that the object is a face, while Hole, George & Dunsmore (1999) also suggest that holistic information might form an initial part of the face encoding process.

1.1.3. Theories and Models: The Holistic Processing Hypothesis

It would therefore appear that there is a good deal of support for the dual-mode hypothesis, with a variety of studies showing that configural information is disproportionately disrupted by inversion compared to featural information. The other major viewpoint in this area is that face recognition involves little part decomposition relative to other objects; that is, it relies to a lesser extent on representations of its composite parts, such as the eyes, nose and mouth (Farah et al, 1998; Tanaka & Farah, 1993). These authors acknowledge that such representations do exist under certain circumstances, enabling the recognition of isolated face parts, but generally speaking their hypothesis is that face recognition relies on representations of the whole face over and above any other kinds of information (Farah et al, 1998).

Tanaka & Farah (1993) showed participants a series of faces with an associated name, and other non-face stimuli as belonging to other names (e.g. 'This is Andrew', 'This is John's house). They were then asked to identify isolated features and whole faces in a two-alternative forced-choice task, where, for the whole faces, the distractor face differed from the target by just one feature (e.g. 'Which is Andrew's mouth?', 'Which is John?'), for both face and non-face stimuli. Features were recognised with greater accuracy when seen in the context of the whole face than when seen in isolation, even though the two faces differed by only one feature. However, if the face was inverted or scrambled during the learning phase (and again at test), the advantage for seeing the whole face disappeared; features were recognised with the same accuracy, whether they were seen in isolation or in the face context. Tanaka & Farah (1993) argued that participants formed explicit representations of the

face parts for inverted and scrambled faces (because configural information was not accessible), but not when learning the upright faces.

Furthermore, Tanaka & Sengco (1997) point out that these findings are not evidence of part-based processing in which configural information is most important. If it *was* evidence of part-based processing, then changes in configural information would affect overall face recognition, while that of the individual parts should remain intact. Tanaka & Sengco (1997) found the opposite pattern of results; when participants were shown one of two face configurations and asked to recognise the features at test, they were recognised with the greatest accuracy in the same configuration as they had previously seen them, and with the least accuracy when seen in isolation. Indeed, changes in the configural location of the eyes not only affected recognition of the eyes themselves, but also other features whose spatial locations were not directly affected.

Farah, Tanaka & Drain (1995) have argued that the idea that face recognition involves little part decomposition is also able to account for the face inversion effect. In an experiment looking at the holistic encoding of non-face stimuli, participants learned to identify a number of dot patterns which were either all one colour, encouraging the grouping of dots into a single, whole pattern, or different colours, encouraging the grouping of similar colours together to form different subparts to the overall pattern (Farah et al, 1995). The patterns were learned upright, and seen in upright and inverted orientations at test. An inversion effect was observed when a holistic encoding strategy was encouraged, but not when these patterns were grouped into parts. This appears to suggest that holistic encoding is orientation sensitive, and,

therefore one might expect similar inversion effects for faces if faces are indeed processed in a holistic manner. In their second experiment, participants learned half of a set of faces in a part-based manner, where the features of the face were presented separately, and the other half as normal, whole faces. All faces were seen normally at test in upright and inverted orientations. Faces which had been presented as separate features earlier showed no effect of inversion on recognition, but those which had been seen as wholes in the learning phase showed an inversion effect in the later recognition task. This appears to support the belief that face perception is largely holistic in nature, and that this holistic information is highly orientation sensitive (Farah et al, 1998; Tanaka & Farah, 1991; 1993).

There is also some support for this holistic view of face processing from more recent studies which have looked at the role of featural information in the early stages of face recognition. Carbon & Leder (2005) noted that for inverted Thatcherised faces, while the eyes and mouth are seen in the correct orientation, the Gestalt of the overall image is not coherent. Normal inverted faces, in contrast, have a coherent Gestalt, but the eyes and mouth are inverted, which may hinder featural processing of these faces. They suggest that if featural information is important in the early stages of face recognition, inverted Thatcher faces should be processed faster than inverted normal faces, because the key featural information (the eyes and mouth) is already correctly oriented. If the early stages of face processing are holistic in nature, however, one might expect inverted normal faces to be processed faster, because they have a more coherent holistic Gestalt than inverted Thatcher faces.

Carbon & Leder (2005) tested these predictions in a number of recognition tasks, using short presentation times to look at the nature of early face processing. In Experiment 1, participants were asked to name a set of nine famous faces from an unaltered image of that person. In each trial of the test phase, they were asked if each of the faces was an original picture of one of the nine celebrities (the name of the target celebrity was shown, and on half of the trials the two identities did not match), and each face was shown, inverted, as both a normal and a Thatcherised version of that person. Faces were presented for either 26ms (short) or 200ms (long), after which participants were encouraged to respond as quickly and accurately as possible, but there was no time limit on their responses. The findings showed that inverted Thatcher faces were recognised faster than inverted normal faces in the short presentation condition, but this advantage for Thatcher faces disappeared in the long presentation condition, with inverted *normal* faces showing faster recognition times. These findings may reflect an initial, very early stage of face encoding where contextual and feature information are processed separately, before being combined into a more holistic representation (Carbon & Leder, 2005).

Carbon & Leder (2005) also conducted a similar experiment with upright normal and Thatcherised faces. The task was again to decide whether each of the faces was an original image of one of the celebrities, with each face shown, upright, as both a normal and a Thatcherised version of that person. While Thatcher faces were readily detected as being distorted, as indicated by a sensitivity measure, reaction times to correctly recognise upright Thatcher faces were significantly longer than those for normal faces. These findings lead Carbon & Leder (2005) to conclude that more integrated, holistic processing is involved in the recognition of briefly presented

upright faces, whereas the early encoding of inverted faces is of a more featural nature.

The findings from Carbon & Leder (2005) are arguably quite important ones, as they seem to suggest that the early stages of face processing may involve different processes. As discussed elsewhere in this thesis, however, there are a number of studies which appear to suggest that it is the disruption of configural information that explains the Thatcher illusion. Proponents of the dual-mode hypothesis may therefore interpret Carbon & Leder's (2005) Experiment 1 in terms of holistic information playing a potential role in early face processing, before giving way to configural information in upright faces; holistic information may be more important than previously envisaged in the early encoding of faces, but this soon gives way to configural information as the most important source of information in upright face processing.

1.1.4. Spatial Positioning of Featural Information & Theories of Face Processing

A related paradigm here, which has also informed this debate between the two major theories of face processing, as well as informing our understanding of the face inversion effect, is that of studies which have looked at the *type* of information being processed which differentiates upright from inverted faces (e.g. Bartlett & Searcy, 1993; Leder & Bruce, 2000; Searcy & Bartlett, 1996). These studies have looked primarily at the effects of changes to the spatial position of features in an attempt to consider the exact nature of 'configural' or 'featural' information, and how face processing is affected by changes made to these types of information. Searcy & Bartlett (1996) considered the effect of inversion on a simultaneous comparison task

with spatially distorted and featurally distorted faces. Inversion significantly hindered participants ability to decide, within a given time frame, that a pair of spatially distorted faces were the same or different, but this effect of inversion was not found with featurally distorted pairs, and responses made within this time frame (3 secs) were longer for detecting configural differences than for detecting featural changes. This appears to suggest that inversion is more disruptive to the processing of configural information than to that of featural information, a view which can readily be accounted for by the dual-mode hypothesis (see also, for example, Freire, Lee & Symons, 2000; Mondloch, Le Grand & Maurer, 2002).

Yovel & Kanwisher (2004), however, have argued that these studies did not manipulate the shape of the features as such, but their lower-level properties, such as the colour of the eyes or teeth (e.g. Le Grand, Mondloch, Maurer & Brent, 2001; Searcy & Bartlett, 1996), which may not be a valid test of the extent to which the face system processes configural and featural information. Yovel & Kanwisher (2004) addressed this issue by comparing configurally and featurally altered faces and houses in a sequential matching task. Faces and houses differed either in terms of the configural relations between features, or the shape of the features themselves. In the latter condition the relational information between features was retained as much as possible. They looked at the neural response of the fusiform face area (FFA), an area which is involved in, and indeed thought to be necessary for, face perception (e.g. Barton, Press, Keenan & O'Connor, 2002), and the effect of inverting these faces on face matching performance. Participants simply decided whether two sequentially presented faces (or houses) were the same or different.

The results were somewhat surprising. Houses produced no inversion effects, and the FFA did not respond differently to configural or featural changes with these stimuli. Faces on the other hand produced consistent inversion effects, but surprisingly the FFA also responded equally to configural and featural changes in faces. These findings appear to suggest that the activity of the FFA, and the inversion effect itself, is domain specific for the processing of faces.

Yovel & Duchaine (2006) followed up this study by considering the extent to which performance on this task can be accounted for by the dual-mode and holistic accounts of face processing. They also note that in previous research, two different types of stimuli have been used, yielding very different results. 'Alfred' faces (as used in Yovel & Kanwisher, 2004) produced similar inversion effects for configural and featural changes in normal subjects, while 'Jane' faces, used in experiments such as those of Le Grand et al (2001, 2003) have shown a larger inversion effect for configural than featural changes. Yovel & Duchaine (2006) compared performance on their earlier task using these two types of stimuli in order to examine where the critical difference between these faces lies. In the study, they compared developmental prosopagnosics, individuals who have failed to develop normal face processing skills, and controls on the matching task reported in Yovel & Kanwisher (2004), with both 'Alfred' and 'Jane' faces, and houses. If the holistic processing hypothesis is correct, one would expect that prosopagnosics will be impaired at detecting configural and featural changes with faces, but will perform normally with houses in both conditions. If the dual-mode hypothesis is supported, however, prosopagnosics would be impaired at processing configural changes, but perform normally with featural changes, and for both types of changes with houses.

For the 'Alfred' faces, controls were found to detect configural and featural changes with equal levels of accuracy, for both face and house stimuli. The developmental prosopagnosics, however, showed a similar level of performance for houses, but accuracy levels were much lower for face stimuli. Importantly, however, there was also no significant difference between configural and featural discriminations for these individuals. With the 'Jane' faces, however, controls were better at detecting featural than configural changes and this effect was even more pronounced for the developmental prosopagnosics with these stimuli.

One key difference between the two types of faces is that the 'Alfred' faces differed in terms of the shape parts, whereas featural changes with the 'Jane' faces also included changes to the brightness and contrast of those parts. To investigate this difference further, in a second experiment, Yovel & Kanwisher (2006) created a new stimulus, Ann, which was manipulated in three ways; configural changes to the distance between the eyes, nose and mouth; featural changes alone, where the eyes and mouth of four different faces replaced the original features (the 'Alfred-like' change), and featural changes (as above) differing in both shape and contrast (e.g. wearing lipstick, different colour eyes; the 'Jane-like' change).

Yovel & Duchaine (2006) hypothesised that face mechanisms may not be necessary to discriminate between faces which differ in terms of low-level visual properties, such as colour or contrast, and there should therefore be no effect of inversion for normal participants for discriminations which can be made on these kinds of information. Face mechanisms should, however, be engaged for the

discrimination of features that differ in shape as they are for faces that differ in their configural information, and therefore normal participants should show an inversion effect in both cases. Prosopagnosics, who only saw upright faces, should also show a similar impairment in the discrimination of configural and featural changes.

The findings with the new 'Ann' face were consistent with these predictions; an inversion effect was observed for normal participants for configural and 'Alfred-like' changes, but not with the 'Jane-like' faces (see also Leder & Bruce, 2000), and this pattern of performance was also true of prosopagnosics with upright face discriminations. Taken together, these findings appear to suggest that face perception mechanisms extract configural and featural information together, supporting the holistic account of face perception, and that this operates through a domain-specific (i.e. face only) processing system.

Riesenhuber, Jarudi, Gilad & Sinha (2004) have also highlighted a further methodological shortcoming of studies which have suggested that inversion is more disruptive to the processing of configural information than to that of featural information (e.g. Freire et al, 2000 (experiments 3 & 4); Mondloch et al, 2002). That is, they have used blocked designs, where either the trials were blocked by the type of change, or change type was a between-participants variable. Knowing the type of change that will be seen may enable participants to adopt change-type strategies accordingly and, particularly for featural changes, such strategies may engage a low-level processing strategy rather than the face processing system itself.

In their study, participants performed a same-different matching task, in which two faces differed by either a configural or a featural change, in both upright and inverted orientations. When the trials were blocked, there was no overall difference in performance between configural and featural trials in the 'configural first' group (all configural pairs shown first, then featural pairs), but in the 'featural first' group, there was a dramatic difference between these two types of changes, suggesting that blocking trials *can* induce different processing strategies. Further, when the trials were unblocked, there was a significant main effect of orientation, but the size of the inversion effect was similar for both types of change, consistent with the research of Yovel and colleagues (2004, 2006) and the holistic account of face processing.

Schwaninger, Lobmaier & Collishaw (2002) have noted that changes such as replacing features such as the eyes and mouth with those of another face may also disrupt the spatial locations (configural information) of these features. Similarly, configural changes to the inter-eye distance, for example, may also produce featural changes, such as in the appearance of a wider bridge of the nose (e.g. Rakover, 2002). Schwaninger et al (2002) used psychophysical techniques, incorporating scrambling and blurring, in an attempt to investigate the roles of configural and featural information separately. Participants saw ten intact faces, with each face shown for ten seconds and shown twice in the learning phase. At test, twenty scrambled faces, twenty blurred faces or twenty scrambled and blurred faces were shown, ten of which were targets (seen before) and ten distractors (previously unseen). Participants were simply required to decide as quickly and accurately as possible whether the face was new or old, by means of a key press. Recognition performance was measured by calculating a d' prime value for each participant, which was then averaged across each

group (scrambled, blurred, scrambled and blurred). Faces were found to be recognised significantly above chance for scrambled faces, but no different from chance when faces were blurred and scrambled at test. This appears to suggest that featural information was encoded in the learning phase, enabling participants to recognise the scrambled versions of those faces at test. Indeed, faces were also recognised significantly above chance for blurred images, suggesting the independent encoding of configural information in the learning phase.

In a second experiment, Schwaninger et al (2002) replicated these findings with both unfamiliar and familiar faces, suggesting that featural and configural information play an important role in both of these types of face recognition. These findings suggest that featural and configural information can be encoded independently from faces, and that these are both important sources of information in face recognition.

Leder & Carbon (2006) also addressed the idea that featural changes may well disrupt the spatial relations of these features in an attempt to investigate whether configural processing is required (or used) when faces differ in terms of their featural information. They looked at faces varying in three different types of information; firstly, faces differed in terms of just their colour, with each face being allocated a unique combination of colour values, with the features and the relations between these features kept the same. A second set of faces differed only in the spacing of their composite features (with features themselves and colour held constant), and a final set differed only in the individual features, with, as far as was possible, the spatial relations between them remaining constant. Names were assigned to each of these

artificial faces, and participants learned each of the faces and their associated name before the test phase began. In each trial, one face was presented beneath the list of names, and participants were simply asked to press the number corresponding to the name of the face shown. Each face was shown twice in each orientation (upright and inverted), and was shown on the screen for either 2 or 8 seconds, depending on the presentation time condition.

Following Leder & Bruce (2000), no inversion effects were observed when faces differed only in terms of their colour information, but those which differed in their configural information showed large effects of inversion on recognition accuracy performance. Of most interest for this study, faces which differed in featural information showed smaller but significant effects of inversion, leading Leder & Carbon (2006) to suggest that making featural changes to faces also affects configural information, and it is this, rather than the featural information itself which produces the inversion effects observed here. Leder & Carbon (2006) suggest that these findings indicate that features also consist of orientation-sensitive configural information, and that these should be distinguished from purely 'local' features, such as colour, which show no effects of inversion.

Yovel & Kanwisher (2004), however, would interpret these findings differently. They note that any change to configural information caused by making featural changes to a face will be much more subtle than that when configural information is deliberately altered. If participants rely on configural information to make featural discriminations, therefore, performance should be much lower for featural than for configural distinctions, a finding they failed to observe in their

studies (see also Yovel & Duchaine, 2006). Indeed, while the inversion effect was still significant for featural changes in Leder & Carbon's (2006) study, performance was higher, not lower, for featural than configural discriminations. Thus, the findings with regard to the relative roles of configural and featural information in the inversion effect show a number of discrepancies, not least in the sizes of the inversion effect found for configural and featural discriminations. In Chapter 2 we address this issue by looking at the effects of configural and featural changes to faces in a visual search task, in which participants search for configurally or featurally distorted faces in arrays of otherwise identical faces.

1.1.5. Summary

We have provided a variety of evidence to support each of these two major theories of face processing, and this is by no means an exhaustive review of the research in this general area. These studies have been informative as to the conditions under which different types of processing may occur. As we have shown in this review, recent research has acknowledged the methodological shortcomings of previous studies which have suggested a clear division between configural and featural processing. Studies such as those of Yovel & Duchaine (2006), Riesenhuber et al (2004) and Carbon & Leder (2005) have all suggested that holistic processing is required for tasks where previous research has appeared to support the dual-mode view of processing. In this thesis, we attempt not to distinguish between these two views of face processing as such, but to look closely at their assumptions and predictions with specific regard to the effects of rotation on face perception. This is of particular interest in studies which encourage holistic processing, as few studies have looked at the effects of rotation on performance on such tasks. In Chapters 4 and 5,

we move on to look at the effect of rotation on a matching task with Thatcherised faces, and draw on a revised definition of configural information, proposed by Boutsen & Humphreys (2003) to examine closely the effects of rotation on a task which is known to disrupt configural information.

1.2. Rotation and Face Perception

While many studies have considered the detrimental effect of inversion on face perception, until recently few had looked at the effect of other (intermediate) angles of rotation. These are important and interesting questions for a number of reasons. First, considering the effects of rotation allows us to examine more closely the roles of configural and featural information in face processing and inversion, and provide an insight into where, if at all, there may be a switch between one kind of processing and another. A further question is whether the effect of rotation on a number of measures of performance is of a linear decline, or follows more of a non-linear pattern, possibly with a more definite switch in processing.

1.2.1. Mental Rotation

A number of early theories of the effect of inversion on face recognition drew on the concept of mental rotation as an important component of being able to recognise inverted faces. Rock (1973), for example, suggested that inverted faces were particularly hard to recognise because we are unable to simultaneously reorient the features of faces. The need to rotate each of the features one at a time, therefore, means that configural information about their spatial relationships cannot be accessed. Carey & Diamond (1977) also suggested that the recognition of inverted faces may depend on the processing of isolated features, as configural (or second-order

relational) information is disrupted, and these features may need to be reoriented one at a time to be recognised. Indeed, in an object recognition study, Shepard & Metzler (1971) actively encouraged the mental rotation of three-dimensional shapes; participants had to decide whether two simultaneously presented shapes were the same or different. Reaction times were found to increase linearly as a function of the difference in orientation between the two shapes.

Schwaninger & Mast (2005) have recently provided a direct test of Rock's (1973) mental rotation hypothesis. They asked participants to detect featural changes (changing the eyes and mouth for those from another face) and configural changes (the interpupillary distance, the distance between the pupils and the lowest part of the nose, and the nose to mouth distance) in a same/different matching task at seven angles of orientation from upright to inverted (0° , 30° , 60° , 90° , 120° , 150° , 180°). The mental rotation hypothesis argues that rotated faces can only be processed by the mental rotation of features, so rotation should affect the accuracy of configural but not of featural changes. This mental rotation process is time consuming, however, and therefore reaction times will be affected by rotation for both types of change (featural and configural). The findings were consistent with these predictions. In terms of accuracy, the detection of featural changes was not affected by rotation of the face, but configural changes were dramatically affected by face orientation. Moreover, the reaction time data showed that both featural and configural changes were significantly affected by rotation of the face, with no difference between these two conditions.

Schwaninger & Mast (2005) did not, however, find a linear relationship between matching performance and angle of rotation. Analysis of the error data

revealed that participants made the greatest number of errors for faces rotated 90° and 120° from upright, with this number dropping for 150° rotated and inverted faces. Indeed, a similar, albeit much less pronounced, pattern was found for the reaction time data, with 90° and 120° rotations yielding longer reaction times than 150° or inverted faces. Schwaninger & Mast (2005) did not look at these findings in more detail with statistical tests, but did suggest that a mental ‘flipping’ strategy, proposed by Corballis, Zbrodoff, Shetzer & Butler (1978), may be being used for inverted faces, whereby participants flip inverted faces to match them to a stored representation of that face. Such an explanation would also account for the slightly faster reaction times for detecting configural changes in inverted faces.

The findings of Schwaninger & Mast (2005) can also be accounted for by the dual-mode hypothesis, although this theory is unable to predict an increase in reaction times as a function of rotation for both configural and featural changes. Nevertheless, these findings provide further evidence against a purely holistic view of face processing. Featural information appears to be explicitly represented in the face processing system, suggesting that faces are not encoded as ‘unparsed’ wholes, as previous authors have argued (e.g. Tanaka & Farah, 1993). On the basis of these results, it may be that the dual-mode hypothesis is (largely) able to explain the physical nature of the effects of face rotation (and the inversion effect), whereas the mental rotation hypothesis is able to account for the means by which face recognition tasks are accomplished.

1.2.3. Isolating Configural Processing in Face Perception

A number of studies have used a range of techniques to isolate configural processing, in an attempt to look at the relationship between rotation and face recognition performance in more detail. Collishaw & Hole (2002) used blurring as a means of distorting the information provided by local features to isolate configural processing, and found evidence of a linear relationship between rotation and face recognition performance on a familiar/unfamiliar discrimination task. Collishaw & Hole (2002) suggested that the linear relationship found here was due to the progressive detrimental effect of rotation on the processing of configural information. McKone (2004) has suggested that methodological issues may account for these findings, however. She notes that Collishaw & Hole (2002) used blurring to isolate configural information in this study, which removes higher spatial frequencies from faces. The familiar/unfamiliar discrimination task may also contain a confound, in that the faces were either of English or Dutch origin. McKone (2004) suggests that this may have enabled participants to complete the task using relatively coarse differences between faces, rather than a fine detail analysis of the configural information in the face.

In her research, McKone (2004) attempted to design a task would isolate configural processing completely. In her first experiment, McKone (2004) presented faces in one of 9 locations; centrally and four horizontal fixation points to the left or right of central, and each face was shown for 150ms to ensure that participants could not move their eyes to look at the face. Participants were trained to identify two same sex faces in the central position, at each of six angles of orientation between upright and inverted (0°, 45°, 67.5°, 90°, 112.5°, 180°). In the test phase, in each trial, a

fixation cross was presented, then replaced by one of the two faces at one of the nine possible locations and in one of the six possible orientations. Participants were simply asked to decide which of the two faces was shown, by means of a key press.

McKone (2004) predicted that at an unknown distance from the central fixation point, a difference in recognition accuracy would emerge between upright and inverted faces, with configural processing remaining largely intact for upright faces, but not for inverted faces. This was indeed the case, with lower recognition accuracy the greater the distance from fixation, and better performance for upright than inverted faces.

The data was collapsed across the two experimental faces and left versus right of fixation, and isolation (of configural processing) points were calculated for 17 of the 21 participants. An isolation point was where accuracy for upright faces was at least 70% (between 70% and 95%) and accuracy for inverted faces was no better than 60% (between 40% and 60%). Recognition accuracy was then extracted for all orientations at that position. The effect of rotation was curvilinear, rather than a linear decline in performance as a function of rotation. In contrast, in a similar experiment with isolated features (in this case the nose), no peripheral effect of inversion was found. This appears to rule out an explanation of the findings observed with whole faces that the peripheral inversion effect may occur “for any stimulus with a preexperimentally defined upright orientation” (McKone, 2004, pg. 188). It also appears to suggest that inversion effects do not always emerge in the periphery when accuracy levels fall below ceiling.

A similar pattern of rotation effects with whole faces was found in McKone's (2004) second experiment, using Mooney faces (Mooney, 1957). These are high contrast photographs with the face formed of white surfaces and black shadows. Participants who were able to perceive the Mooney face strongly when upright but not at all when inverted took part in the experiment, and were shown the Mooney face in 72 different orientations (formed by rotating the image in 5° steps from 0°), and in each orientation 8 times. Participants were asked to judge how strongly they saw the face in the image and how three-dimensional the face appeared (if the face can be seen, it appears three-dimensional rather than flat). Scores were collapsed across these two measures, as no differences were observed between them.

Consistent with the findings of Experiment 1, there was a clear curvilinear effect of rotation on configural processing, producing a bell-shaped curve from -180° to 180°. The breadth of orientation tuning (how far from upright the image can be rotated before configural processing disappears) was quite wide between participants, ranging from 73° to 140°. McKone (2004) suggests that these differences may account for the different accounts of breadth of tuning reported in previous studies, with relatively broad tuning observed when a large number of participants are tested, whereas studies using a smaller number of participants (e.g. McKone et al, 2001) have tended to observe a narrower range.

The important findings to note here, then, are that McKone (2004) has attempted to isolate configural processing using two unique tasks, and found strong effects of inversion on performance for each. Moreover, these findings suggest that

when configural information is isolated, the effect of rotation on this information is one of a curvilinear, rather than linear, decline.

1.2.4. Rotation and the Holistic View of Face Processing

Although they only studied performance across three different angles (0° , 90° , 180°), Lewis & Glenister (2003) provided an alternative account of the effects of rotation on face perception. These authors expanded on Tanaka & Farah's (1993) work, which found that when a set of faces were learned in the upright view, there was an advantage at test for discriminating between two faces to identify a particular individual, compared with deciding which of two isolated features were belonged to an individuals face. However, this whole-face advantage disappeared when the faces were inverted in the learning phase and at test. Lewis & Glenister (2003) also considered the effect of 90° rotations on performance in this forced-choice discrimination task to assess how configural encoding, which is disrupted by inversion, is affected by rotation. A similar effect of rotation was observed for isolated features and whole faces for upright and 90° rotated faces, suggesting that features and whole faces are equally affected by acute rotations. Whole face recognition was also affected by the second 90° of rotation, but isolated feature recognition remained intact following this transformation. Lewis & Glenister (2003) suggest that the first 90° of rotation affects configural and featural processing in similar ways, while the second 90° rotation disrupts highly configural processing of faces.

Lewis & Glenister (2003) also found that isolated features were similarly affected by inversion, when compared to upright features. They suggest that, for high quality images, configural encoding of the individual features may take place as well

as configural coding of the whole face, and inversion disrupts this processing compared with upright features. Moreover, the findings of a whole-face advantage over isolated features for 90° rotations here is problematic an expertise account of the face inversion effect, particularly if it is doubtful that we are experts in processing faces rotated to this degree.

Whether the effect of rotation on the processing of configural information is of a linear or a nonlinear decline in face recognition performance, we suggest that either pattern of results provides problems for the notion that we are poor at recognising inverted, or rotated, faces simply because we have less expertise in dealing with them than we do with upright faces. The linear detrimental effect of rotation found by Collishaw & Hole (2002), for example, presents clear problems for such a theory, suggesting that there is not a point at which face recognition suddenly becomes more difficult as faces are rotated beyond the level of our expertise.

Indeed, a similar argument can be used against the expertise account even if a nonlinear effect of face rotation is observed. A number of studies have suggested that a change in processing occurs somewhere between 90° and 120° rotations (Schwaninger & Mast, 1999; Murray et al, 2000; Sturzel & Spillmann, 2000). However, it is difficult to convincingly suggest that we are likely to have more experience of seeing faces rotated up to 105° from upright (to take a midpoint of between 90° and 120°) than those rotated beyond this angle. We often see faces tilted in everyday life, for example in conversations, but it is unlikely that we have much experience or expertise with seeing faces beyond a small angle of rotation (Collishaw & Hole, 2002).

1.2.5. Summary

We have reviewed a wide range of literature on the effects of rotation on face processing here, highlighting why this forms such an important question in face perception research. There is a wide variety of evidence to suggest that it is configural, rather than holistic information which is affected by face rotation (e.g. Collishaw & Hole, 2002; McKone, 2004; Schwaninger & Mast, 2005). Moreover, a number of studies have suggested that the relationship between face recognition performance and rotation is of a linear nature, but these studies have been criticised on methodological grounds, not least because they have employed tasks which are unable to discriminate between configural and featural information. Linear rotation effects may therefore be the product of contributions of both types of information (Martini et al, 2006). Rather, the emerging view appears to be that the effect of rotation is nonlinear in nature, and this has been shown in a wide variety of tasks, and also in studies which have attempted to isolate configural processing in a systematic way (e.g. McKone, 2004; see also Martini et al, 2006). In this thesis, we attempt to expand on these findings to look at the effects of rotation in recognition memory studies (which research has not considered to date) and with Thatcherised faces, where previous research has produced inconsistent findings regarding the nature of the relationship between face processing and rotation.

1.3. Summary of Chapter 1

In this chapter, we have reviewed the existing literature on the two main areas of this thesis; the roles of configural, holistic and featural information in face processing, and the effects of inversion and rotation in face perception. We argue that

the fact that the dominant views in each of these areas have been challenged in recent years (particularly with regard to the issues of configural vs. holistic processing and linear vs. non-linear relationships respectively) suggests that research in these areas is still very much alive, with many questions still to be addressed. In this thesis, we begin by looking at the effects of inversion on the detection of configural and featural changes to faces in a visual search task, a question which has also not been addressed in previous research, in order to look at the separate effects of inversion on (changes to) configural and featural information, and to consider the suggestion that there may be a role of configural information in the processing of featural information (Chapter 2). If faces are processed in a holistic manner in this task, we would expect to find inversion effects of approximately similar size for both types of change. If configural and featural changes are differentially affected by inversion, however, we would expect configural changes to be affected by inversion, while featural information should remain unaffected, consistent with previous research. Further, however faces are affected by inversion on this task previous research (e.g. Murray, 2004) suggests that we should expect to find effects of array size consistent with a serial search strategy for both types of changes.

Having considered the effects of inversion in our own face processing experiment, we then attempt to address the need for a clearer understanding of the effects of rotation, and the processes involved in the recognition and perception of these faces, by looking at the effects of rotation in a face recognition memory task, an area which has not previously been considered (Chapter 3). If holistic processing is important for upright faces on this task, previous studies suggest that we would expect to find a whole-face advantage when faces are seen upright in the learning phase and

at test, but this advantage should disappear for inverted faces (Tanaka & Farah, 1993; Leder & Carbon, 2005; Lewis & Glenister, 2003). A whole-face advantage would also be expected for upright faces if configural information is important for completion of this task, but for inverted stimuli recognition accuracy may be similar for whole faces and features, or there may be a disadvantage for whole faces, depending on whether configural information is disproportionately disrupted by inversion relative to featural information. We also consider the mental rotation hypothesis in a second experiment in this chapter, whereby if participants are mentally rotating faces to match a stored memory representation, recognition accuracy should decline as a function of the difference in orientation between the learned faces and those seen at test.

We then move on to look at the effect of rotation on the processing of Thatcherised faces (Chapters 4 and 5), a task which is seen as an indicator of configural processing, and where inconsistent findings regarding the face processing-rotation relationship have been reported. Previous research suggests that we should expect to find a greater detrimental effect of rotation for normal faces than for Thatcherised faces on this type of task, regardless of whether the images are same-person or different-person pairs (Boutsen & Humphreys, 2003) and, moreover, that these findings can be explained by a distinction between local and global configural processing. If featural and configural processing strategies are required for the matching of same and different identity face pairs respectively, however, we would expect to find a greater effect of rotation for Thatcherised faces, but only for different-person pairs. In a second set of experiments (Chapter 5), we suggest that showing two different images of the same person may also reveal a configural processing strategy, in which case we would again predict a greater effect of rotation for Thatcherised than

for normal faces, but for different-image pairs only. Finally, if featural information is used to match faces for identity irrespective of type, we might expect to find significant effects of rotation for normal and Thatcherised faces, but no difference between these types of face pairs. The implications of our findings for the dual-mode and holistic theories of face processing for the nature of the relationship between rotation and face processing are discussed, and the extent to which the mental rotation hypothesis can account for these findings is also considered.

Chapter 2: Detecting configural and featural changes in a visual search task: evidence for the face inversion effect

In this chapter, we look at the effect of inversion on the detection of configural and featural changes to faces in a visual search task, in which either the faces are all identical, or one face is configurally or featurally altered. Our findings indicate that correct 'same' or 'different' responses are the product of serial, rather than parallel processing. Moreover, we found some evidence of a greater detrimental effect of inversion for configural than featural changes in this experiment, suggesting that there may be a role of configural information in the processing of featural information.

2.1. Experiment 1

While there is clearly considerable evidence for the inversion effect in the face perception literature, many of these studies, such as those discussed in Chapter 1, have been unable to look at the separate effects of inversion on (changes to) configural and featural information. One way of addressing this issue is to consider the effects of inversion on a task in which configural or featural changes can be made to a stimulus, and compared with performance with normal, unaltered faces. We suggest that a visual search paradigm, in which participants search for a target stimulus among an array of distractors, provides such a task. This task has also been used to inform the debate between the two major theories of face processing and our understanding of the face inversion effect, again looking at the type of information being processed which differentiates upright from inverted faces (e.g. Kuehn & Jolicoeur, 1994; Murray, 2004).

The main aim of a visual search experiment is to determine how efficient search can be for an item when the size of the array is increased. Also of interest are the various manipulations made to the stimuli, and how these interact with the size of the array. While the distinction between two types of search strategy, parallel and serial, may not be as simple as it appears (see Chapter 7), we suggest that it nevertheless provides a useful method of describing the efficiency of visual search. Performance on visual search tasks is usually measured in one of two ways, which can be used to indicate whether parallel or serial processing is taking place. First, the increase in reaction time to detect an item per additional distractor is used, and where this increase is greater than 10 ms per item, search is deemed to be serial, whereas where it is less than this it is parallel, and item can be said to 'pop out' of a visual array (Wolfe, 1998a). Second, a self-terminating search hypothesis predicts that a serial search will have a 2:1 ratio between target absent and target present slopes (Treisman & Gelade, 1980). A ratio greater than 2:1 can be used to indicate pop-out. In this experiment we look at the question of whether search is parallel or serial in the detection of configural and featural changes to faces in a visual search array. Moreover, we also look at the effects of inversion on the search strategies employed, for different features and, in an overall analysis, for each type of change (configural or featural).

A number of studies have used visual search tasks to assess the role of holistic encoding in face perception. A further prediction of the holistic view (see for example Schwaninger et al, 2002) is that perception of a feature will be influenced by the configural context in which it is found, and therefore changes to configural information should affect the recognition of individual features. Mermelstein et al

(1979) tested this prediction in their study, and found that the detection of a target shape, such as a diamond, in a visual search task was impaired when presented within a face-like configuration (for example as a nose). Mermelstein et al (1979) argued that this was evidence of holistic processing taking precedence over featural information, preventing perception of the feature without reference or access to the configuration in which it occurred. They predicted, however, that this 'face-inferiority' effect could be reversed if the presence of the target feature could be correlated with a feature of the face, derived from holistic information, such as an assessment of attractiveness (e.g. grotesqueness), and this feature discriminates the target from the distractor. In this case, the face would be encoded in a holistic manner, and would result in the rapid perception of the face as being grotesque. The grotesqueness of the face would therefore be associated with the presence of the target feature, which should facilitate discrimination of the target from the distractors.

Murray (2004) attempted to test this prediction, using photographs of real faces and by looking at the effects of inversion on the holistic encoding of faces. Participants searched for an inverted mouth in four different contexts; normal upright faces, normal inverted faces or scrambled faces, in which the configural information is disturbed (upright and inverted). The distractor faces contained an upright mouth, but were otherwise identical to the target face in each of the face contexts. The four face contexts were presented in separate blocks, to ensure that the location of the mouth within the face would be known on each trial, and arrays were comprised of between 1 and 6 items.

For upright faces, Murray (2004) predicted that if holistic processing of configural information precedes featural processing, the rapid perception of grotesqueness, indicating the presence of the target feature, should lead to faster target discrimination (and therefore faster search rates) in upright than in scrambled faces. For inverted faces, however, the perception of grotesqueness created by inverting facial features is almost completely lost when the face is inverted, as has been shown in the case of the Thatcher illusion (e.g. Sturzel & Spillmann, 2000). If inversion disrupts holistic information, inversion should impair target discrimination compared with upright faces, while search rates should not differ between inverted and scrambled faces.

The results were largely consistent with these predictions. Correct target-present responses were faster for upright faces, where the inverted mouth produced a grotesque appearance, than for scrambled or inverted faces, which did not induce holistic encoding. Holistic encoding was also disrupted by inversion, but in this case the original 'face inferiority' effect was observed; search for an inverted mouth was slower in inverted faces than in upright scrambled *and* inverted scrambled faces, suggesting that the loss of configural information in inverted faces is not complete.

Murray (2004) suggests that upright faces are processed through the holistic encoding of configural information, and this holistic information can be used to establish that a face is a face (i.e. that it has a coherent Gestalt), and can also convey information about any emergent features of the face (such as grotesqueness). If this emergent feature is associated with the presence of a target feature, it can be rapidly perceived, and a fast discrimination decision made. This configural information is

disrupted for inverted faces, however, meaning that information about emergent features of the face is also disturbed. Inversion may not disrupt holistic encoding as such, but the holistic processing required to determine that an inverted face is a face may slow discrimination of the individual features, as evidenced by the face-inferiority effect with inverted faces. These findings appear to suggest that faces may be processed in a holistic manner in visual search tasks, even when configural changes are made to those faces, while inversion severely disrupts the holistic encoding of configural information in these tasks.

Thus the evidence reported here from visual search tasks appears to suggest that faces may be processed in a holistic manner when participants are searching for a target face among an array of distractors, even when configural changes are made to the face, while inversion disrupts configural information, impairing the discrimination of individual features (Murray, 2004). Evidence from face matching tasks suggests that changes to featural information are also detrimentally affected by inversion (see for example Yovel & Kanwisher, 2004; Yovel & Duchane, 2006). In this experiment we attempt to combine some of the methodologies reported here to consider the separate effects of inversion on (changes to) configural and featural information. Specifically, we look at the effect of these changes to faces in a visual search task, in which participants are required to decide whether all of the faces are the same (identical), or whether one face is different in some way, i.e. has been altered in either a configural or featural manner.

If faces are processed in a holistic manner in this task, we would expect to find inversion effects for both types of change, and these should be of approximately

similar size for both types of change. If configural and featural changes are differentially affected by inversion, in line with the dual-mode view of face processing, however, we would expect configural changes to be affected by inversion, while featural information should remain unaffected, consistent with previous research (see, for example, Leder & Bruce, 2000). Further, however faces are affected by inversion on this task, we would expect to find effects of array size consistent with a serial search strategy for both types of changes; Murray (2004), for example, found no evidence to suggest that faces ‘pop out’ of an array when participants are searching for a face among very similar face distractors.

2.1.1. Method

2.1.1.1. Participants

Seventeen undergraduates from Cardiff University with normal or corrected-to-normal vision received a small financial reward for their participation in this experiment.

2.1.1.2. Materials

A single Caucasian greyscale face was generated using the computer-based face-reconstruction system Faces 3.0 (published by IQ Biometrix, Inc.). This was the only face used in the current experiment. From this face, configural and featural changes were made to the face in the following ways. For featural changes, the eyes, nose and mouth were simply replaced with two alternative features in each case (two different sets of eyes, two different noses etc.) For configural changes, the eyes of the original face were moved 4 pixels apart and 4 pixels closer together from the original face. The nose was moved 4 pixels up and 8 pixels down from its original position,

and the mouth was moved 4 pixels up and 6 pixels down from its original position. This was done to ensure that moving the nose up did not look the same as moving the mouth down, and moving the mouth up did not look the same as moving the nose down. Examples of the stimuli, as shown in the two sizes of arrays (taken as screenshots from the experiment itself), can be seen in Figures 2.1 and 2.2.




Figure 2.1. An example of an array of four faces seen in Experiment 1. This shows a ‘different’ trial, with the eyes moved closer together in the face on the right of the top row.

The positioning of the faces was chosen in an attempt to prevent participants simply matching along a line of faces to complete the task. In ‘different’ trials each altered face ($n = 12$) appeared in each possible location in the visual array ($n = 8$),

making 96 trials for each of the array size and orientation independent variables (384 trials in total), with the same number of 'same' trials also used. All images were presented on a computer using Superlab.




Figure 2.2. An example of an array of eight faces seen in Experiment 1. This array shows a featural change, with the original mouth replaced with an alternative mouth on the bottom row, the second face from the left.

2.1.1.3. Design

The four within-participants independent variables were the orientation of the faces (upright or inverted), the region of the face which was changed (eyes, nose or mouth), the type of change made (configural or featural) and the size of the array (4 or 8 items). The dependent variables were the latency and accuracy of responses to the 'same' and 'different' arrays.

2.1.1.4. Procedure

Before the experiment began, participants were shown the face that would appear in the experiment, and each of the twelve ways in which one of the faces in the array may differ (if at all). Each trial began with a fixation cross, shown for 200ms, at the centre of the screen. This was followed by one of the visual search arrays, and participants were simply asked to decide whether all of the faces were the same or one was different in some way, by means of a key press on a keyboard. Participants were asked to respond as quickly and as accurately as possible, and the array remained on the screen until the participant made a response, followed by a 1 second blank interval before the next trial.

2.1.2. Results

While the effects of inversion on the detection of changes for each of the different features are undoubtedly of interest (see for example Barton, Keenan & Bass, 2001; Malcolm, Leung & Barton, 2005), in light of the main aim of this experiment, we restrict ourselves to an overall analysis of the effects of inversion on the detection of configural and featural changes. The RTs to make correct 'same' or 'different' decisions to the face arrays were recorded. Separate analyses were performed for the same and different array trials. Only RTs for correct responses and for those not exceeding 10000ms were analysed. This cut-off point was chosen because of the particularly difficult nature of the task. The data was collapsed across the feature changed and a series of ANOVA's were conducted to look at the effects of inversion on changes to configural and featural information. The mean combined RTs and percentage of errors for 'different' pairs are shown in Figures 2.3 and 2.4 respectively. For these trials, an error was a 'same' response to an array of faces

containing an altered face, while for 'same' trials, an error was a 'different' response to an array of faces which were all the same.

A repeated measures ANOVA of the combined RT data shown in Figure 2.3 revealed significant effects of rotation ($F(1, 16) = 8.058; p < 0.05$), type of change ($F(1, 16) = 29.240; p < 0.05$) and array size ($F(1, 16) = 27.992; p < 0.05$), but none of the interactions were significant ($p > 0.05$ in all cases).




Figure 2.3. Mean reaction time for correct 'different' decisions in Experiment 1 as a function of feature, type of change and array size for upright and inverted faces. Error bars show ± 1 standard error.




Figure 2.4. Percentage of errors for 'different' decisions in Experiment 1 as a function of type of change and array size for upright and inverted faces. Error bars show ± 1 standard error.

A repeated measures ANOVA of the combined error data shown in Figure 2.4 also revealed significant effects of rotation ($F(1, 16) = 44.234; p < 0.05$) and type of change ($F(1, 16) = 72.888; p < 0.05$), but no effect of array size ($F(1, 16) = 2.308; p > 0.05$). This analysis, however, also showed that the interaction between rotation and type of change was significant ($F(1, 16) = 9.989; p < 0.05$). This appears to suggest that the effect of inversion was greater for configural than for featural changes. None of the other interactions were significant, however ($p > 0.05$ in all cases). Investigating this interaction, two separate repeated measures ANOVAs were conducted, the first on the data for upright faces. This showed significant effects of array size ($F(1, 16) = 6.515; p < 0.05$) and type of change ($F(1, 16) = 31.201; p < 0.05$), but no significant interaction between these variables ($F(1, 16) = 3.943; p > 0.05$). The analysis of the data for inverted faces, however, found an effect of type of change ($F(1, 16) = 37.770; p < 0.05$) but no effect of array size ($F(1, 16) = .160; p > 0.05$) or interaction between array size and type of change ($F(1, 16) = 1.317; p >$

0.05). This indicates that the effect of size was significant for upright but not inverted faces, and this is responsible for the interaction found here.

As no changes were made to the faces in the 'same' arrays, the variables in the analysis of this data are simply array size and rotation. As can be seen in Figure 2.5, reaction times were much longer to make correct 'same' decisions. The repeated measures ANOVA conducted on this data showed a significant effect of array size ($F(1, 16) = 85.465; p < 0.05$) but not of rotation ($F(1, 16) = 2.333; p > 0.05$) and there was no significant interaction between array size and rotation ($F(1, 16) = .114; p > 0.05$). These results are indicative of a serial search to decide that all of the faces were the same and, moreover, in contrast with 'different' decisions, inversion had no effect on RTs to decide that all of the faces were the same.




Figure 2.5. Mean reaction time for correct 'same' decisions in Experiment 1 as a function of array size and orientation. Error bars show ± 1 standard error.




Figure 2.6. Percentage of errors for 'same' decisions in Experiment 1 as a function of array size for upright and inverted faces. Error bars show ± 1 standard error.

A repeated measures ANOVA was also conducted on the error data for correct 'same' decisions shown in Figure 2.6. Consistent with the RT data shown above, there was a significant main effect of array size ($F(1, 16) = 10.736; p < 0.05$), but not of rotation ($F(1, 16) = .767; p > 0.05$), and there was no significant interaction between array size and rotation ($F(1, 16) = 2.943; p > 0.05$).

Finally, slope sizes were calculated for each condition of the 'different' RT data, and found gradients of between 64.8 and 192.6 for upright faces, and between 57.4 and 158.2 for inverted faces. Array size had a much greater effect on 'same' responses, with search gradients of 907.3 and 865.7 found for upright and inverted faces respectively. The ratio of the 'different' versus 'same' slopes, therefore, was between 4.7 and 13.9 for upright faces, and between 5.4 and 15.0 for inverted faces, which is significantly greater than 2 for all of the conditions ($p < 0.05$ in all cases).

This indicates that a serial, rather than parallel search strategy was used in this experiment.

2.1.3. Discussion

In this experiment we have looked at the effects of inversion on the detection of configural and featural changes to faces in a visual search task. Previous visual search studies addressing the effects of inversion on configural and featural information have argued that faces can be processed in a holistic manner if the change to the configural information in the face brings out an emergent feature of the face, such as grotesqueness, and this can be associated with the presence of the target feature (Murray, 2004). Evidence from face matching tasks, however, suggests that changes to featural information are also detrimentally affected by inversion (Yovel & Kanwisher, 2004; Yovel & Duchane, 2006).

In the analysis of the RT data, we found significant effects of array size, inversion and type of change, but crucially, there were no significant interactions between these variables, suggesting that the effects of inversion were the same for both types of change, and also that similar search strategies were employed in the detection of both types of change, and at both angles of orientation, as evidenced by the similar search slopes sizes. An interaction *was* observed in the error data, however; there was a larger inversion effect for configural than for featural changes. We note that errors are particularly high in this experiment, especially for inverted faces. We also note that Yovel & Kanwisher (2004) found a larger inversion effect for *featural* than configural discriminations in their error data, but this pattern of performance did not generalise to the reaction time data.

While differences in the task demands may account for the different pattern of inversion effects between these studies, this nevertheless suggests that the error data may provide an important insight into the roles of configural and featural information in the inversion effect. Indeed, these findings are at least in part consistent with those of Leder & Carbon (2006), who also found a larger inversion effect for discriminations of configural than featural changes in a face recognition memory task. These findings may suggest that making featural changes to faces also affects configural information, and we suggest that this may account, in part at least, for the larger inversion effects for discriminations of configural than featural changes observed in this study.

Thus, in this experiment, we have shown a clear effect of inversion on face processing, and we have found some evidence for a role of configural processing in the detection of featural changes, at least when the task is to decide whether an array of faces is all the same or one is different in some way. Having established a face inversion effect, in the remainder of this thesis we turn our attention to the processes involved in the recognition and perception of faces seen at intermediate angles of rotation. We use these studies to help inform the debate between the holistic and dual-mode hypotheses of face processing, and to consider the nature of the relationship between rotation and face processing. We begin by looking at the processing of rotated faces in a face recognition memory task using the part-whole paradigm, which is thought to encourage holistic encoding.

Chapter 3: Parts and Wholes in Face Recognition: The Effects of Rotation

Overview

In this chapter we use the part-whole paradigm to look at the effects of rotation on a face recognition task which is thought to encourage holistic processing. We consider the implications of our findings for competing theories of face processing (e.g. the dual-mode hypothesis, the holistic processing hypothesis) and accounts of the inversion effect. Our findings show a number of differences with previous research (e.g. Tanaka & Farah, 1993), and a number of competing explanations of these results, such as the mental rotation hypothesis, are offered. In a second experiment, we consider whether the orientation in which a face is initially encoded is important in determining the effects of rotation observed in this task. Moreover, we attempt to test the hypothesis that rotated faces are processed by the mental rotation of features, but inverted faces may be processed by a different 'flipping' strategy. A number of explanations for the different effects of rotation and inversion observed in these experiments are offered, and the findings discussed in relation to the dual-mode and holistic theories of face processing.

3.1. Experiment 2

In the part-whole paradigm, developed by Tanaka & Farah (1993), participants are shown a series of faces with an associated name, in an upright or inverted orientation. At test, they are required to identify these faces from their isolated features or from their whole face, in a two-alternative forced-choice task. For the facial features, two different facial parts are shown (e.g. two different mouths), while

for the whole faces, the distractor differs from the target by just one feature. Participants decide which of the two faces or features belongs to the named person.

Tanaka & Farah (1993) found that features were recognised with greater accuracy when seen in the context of the whole face than when seen in isolation, even though the two faces differed by only one feature. However, for inverted and scrambled faces, and schematic houses, the advantage for wholes over parts disappeared; facial features (and house features such as doors and windows) were recognised with the same accuracy, whether they were seen in isolation or in the whole (face) context. Tanaka & Farah (1993) argued that participants formed explicit representations of the face parts for inverted and scrambled faces (because configural information was not accessible), but not when learning upright faces. This, they argued, suggests that upright faces are processed in a holistic manner, and that this information is disrupted by inversion.

Many of the studies looking at the effects of rotation on face perception have looked at perceptual matching tasks (Valentine & Bruce, 1986), bizarreness ratings (Murray et al, 2000) or mental rotation paradigms (Schwaninger & Mast, 2005), but have not considered the impact of rotation on face recognition memory tasks. Our recognition of, and memory for, faces is clearly an important part of our face processing abilities, and it is imperative that we are able to understand which forms of information are important for face recognition memory. Only Lewis & Glenister (2003) have previously considered the effect of rotation on face recognition memory. They found a whole-face advantage over isolated features for 90° rotated as well as upright faces in an adaptation of the Tanaka & Farah (1993) experiment. These

authors argued that their findings indicated two separate effects of face rotation; isolated features and whole faces are equally affected by acute rotations (a similar pattern of performance was observed for each between upright and 90°), while whole face recognition is also affected by the second 90° of rotation, but isolated-feature recognition remains unaffected by this transformation.

A related question here is whether the relationship between rotation and face recognition memory performance is of a linear or a nonlinear nature. This is an important question in that it enables us to look at how the processes involved in face recognition are affected by rotation, and furthermore, may provide an insight into the means by which face rotation tasks are accomplished (e.g. mental rotation; Schwaninger & Mast, 2005). Lewis & Glenister (2003) found different effects of rotation beyond 90°, but recent studies have suggested that any change or 'switch' in processing may occur between 90° and 120°, with a possible improvement in performance beyond these angles to inversion. Here we build on Lewis & Glenister's (2003) research by looking at recognition memory performance across a number of angles of rotation (0°, 45°, 90°, 135°, 180°) to discover exactly where the effects of face rotation occur, how they affect face processing, and the implications of our findings for competing explanations of the effects of face rotation.

If upright faces are processed in a holistic manner on this task, we would expect to find a whole-face advantage when faces are seen upright in the learning phase and at test, but this advantage should disappear for inverted faces, consistent with previous research (Tanaka & Farah, 1993; Leder & Carbon, 2005; Lewis & Glenister, 2003). If configural information is important for completion of this task, we

would also expect to find a whole-face advantage for upright faces. For inverted stimuli, however, recognition accuracy may be similar for whole faces and features, or there may be a disadvantage for whole faces, depending on whether configural information is disproportionately disrupted by inversion, relative to featural information (which remains largely unaffected by inversion) in a recognition memory task.

3.1.1. Method

3.1.1.1. Participants

Eighteen undergraduates from Cardiff University with normal or corrected-to-normal vision received a small financial reward for their participation in this experiment.

3.1.1.2. Materials

Five Caucasian greyscale faces were generated using the computer-based face-reconstruction system Faces 3.0 (published by IQ Biometrix, Inc.). The basic features, such as the head shape, jaw shape and hairstyle, were different for each face. From each face, a further set of faces was created by interchanging 3 types of eyes, 3 types of nose and 3 types of mouth to create a set of 27 unique faces, 135 faces in all. All faces were of approximately similar size, and each image was resized to measure 709 x 647 pixels on the screen. An example of the stimuli seen in the learning phase of the experiment is shown in Figure 3.1.




Figure 3.1. An example of the stimuli used in Experiments 2 and 3 in the upright condition for the (from left to right): learning phase, whole-face test and isolated feature test.

Six faces from each set were used in the learning phase, and none of these faces had more than one feature in common with any of the other learning faces. Distractors were taken from the remaining faces from each set, and were selected such that each differed from the target face by only one feature. The isolated feature images were created from the original faces, using a mask to ensure that all images were of the same size for each type of feature (eyes, nose and mouth). An example of these images is also shown in Figure 3.1. Thirty common names were generated to be associated with each of the faces to be seen in the learning phase. All images were presented on a computer using Superlab.

3.1.1.3. Procedure

There were five parts to the experiment, each representing a different angle of rotation (0° , 45° , 90° , 135° and 180°) and each consisting of a learning phase and a test phase. In the learning phase, participants saw each of the six faces from one of the face sets with its associated name, at one of five angles of orientation (upright, 45° , 90° , 135° , 180°). Each face and name was shown on the screen for 5 seconds, and was seen twice in a random order. There was then a short break to remind participants of the task.

In the test phase, the faces were always seen in the same orientation in which they had appeared in the learning phase, and in each trial, a name and either two faces or two isolated features (eyes, noses or mouths) were shown. One of the faces or isolated features was the correct stimuli for the given name. Participants were asked to decide which of the two faces or isolated features was associated with the name shown by pressing one of two keys on the keyboard. For whole-face tests, the distractor face was a face that differed by only one feature from the target face. For isolated-feature tests, the distractor was one of the other two types of that feature used in the prototype set. All features from all six faces were tested twice, once in a whole face and once in isolation, which made 36 tests for each part of the experiment. The order in which the faces were seen was counterbalanced between participants. The experiment was conducted five times with each participant, once for each of the five different angles of rotation. of the two sets of faces, and participants saw both sets of faces at the same angle of rotation. Both the set of faces used in each condition and the order in which the different angles were shown were counterbalanced between participants.

3.1.1.4. Design

A within-participants design was used in this experiment. The two independent variables were orientation (upright, 45°, 90°, 135° and 180°) and the type of stimulus seen at test (whole face or an isolated feature). The task was to decide which of the two faces or features belonged to the given name, and the dependent variable was recognition accuracy (percentage correct).

3.1.2. Results and Discussion

The proportion of correct identifications for faces and isolated features was calculated for each participant for each angle of rotation for both whole-face and isolated-feature targets. The means of these proportions are plotted in Figure 3.2 below.




Figure 3.2. Mean percentage of correct recognition in Experiment 2 as a function of angle of rotation and stimulus type. Error bars show ± 1 standard error.

A two-way ANOVA was conducted on the proportion of correct choices, with the two within-participant independent variables being the angle of rotation of faces during learning and at test (0° , 45° , 90° , 135° and 180°) and the type of stimuli seen at test (whole faces or isolated features). The analysis found a significant effect of orientation ($F(4, 68) = 5.075$; $p < 0.05$) and of test stimulus type ($F(1, 17) = 10.918$; $p < 0.05$), but the interaction between these two variables was not significant ($F(4, 68) = 0.511$; $p > 0.05$). Comparing just the upright and inverted faces also revealed no significant differences for whole faces ($t(17) = .387$, $p > 0.05$) or isolated features (t

(17) = .000, $p > 0.05$). With regard to the whole-face advantage, a series of paired t -tests revealed no significant differences between whole faces and isolated features for upright ($t(17) = 1.892$; $p > 0.05$) or for inverted ($t(17) = 1.343$ $p > 0.05$) faces.

In terms of the regression analyses, linear and cubic functions were fitted to the data for each condition (e.g. whole faces, isolated features). Comparisons of the ability of the linear and cubic functions to account for the data showed that a cubic function provided a better fit to the data than a linear function for whole faces ($F(2, 1) = 479.632$) but not for isolated features ($F(2, 1) = 4.726$). Thus, the effect of rotation appears to follow a curvilinear, rather than a linear function on this task, but only for whole faces. The effect of rotation on the processing of isolated features appears to be one of a roughly linear decline in performance as the face is rotated away from the upright.

Discussion

The results of this experiment have failed to replicate previous findings of a whole-face advantage for upright faces in this task (Tanaka & Farah, 1993; Lewis & Glenister, 2003). These findings appear to suggest that featural information can be encoded and explicitly represented in upright face processing, in contrast to a purely holistic view of face processing. It is important to note, however, that this pattern of findings has also been observed in other recognition memory studies; Schwaninger et al (2002) found that faces previously learnt upright and intact could be recognised when they were scrambled in the test phase, suggesting that explicit representation of featural information does occur for upright faces.

A further observation here is that we failed to find effects of face inversion on this task, for both whole faces and isolated features. For isolated features, we note that this finding may not be entirely surprising. Tanaka & Farah (1993) and Rhodes et al (1993) both found roughly equivalent levels of accuracy for feature tests from upright and inverted faces, suggesting that the recognition of featural information is not affected by inversion, although Lewis & Glenister (2003) have suggested that this may have been due to the poor quality of the facial images used in these studies. We note, however, that a number of other studies have also shown that featural information remains relatively unaffected by inversion. Leder & Carbon (2005; Experiment 1), for example, used features from high quality digital photographs to create artificial faces, and found equivalent levels of recognition accuracy for upright and inverted features in this part-whole paradigm. This appears to indicate that configural encoding of isolated features may not occur even with high quality facial images of male faces, suggesting that the quality of the facial images used may not have been responsible for this effect in previous studies as Lewis & Glenister (2003) suggested. Our study has replicated these findings, and shows that recognition accuracy remains largely constant across intermediate angles of rotation on this task.

In terms of the models of face processing discussed earlier, the findings reported here clearly pose problems for a purely holistic view of face processing. The failure to find an advantage for upright or inverted whole faces in this task appears to suggest that featural information is explicitly represented in the face processing system, suggesting that faces are not encoded as 'unparsed' wholes, as previous authors have argued (e.g. Tanaka & Farah, 1993). The finding that rotation (but not inversion) affected whole faces but not isolated features also provides evidence

against the holistic processing hypothesis. Rather, these findings appear to be consistent with the dual-mode hypothesis, with configural information from whole faces affected more than featural information by rotation from the upright. The improvement in recognition accuracy with inverted whole faces may mark a further switch in processing, with inverted faces mentally ‘flipped’ for recognition, producing better recognition than for faces presented at intermediate angles of orientation (Schwaninger & Mast, 2005).

3.2. Experiment 3

Here we build on the findings of Experiment 2 to test a specific prediction suggested by these results. The recognition of whole faces was affected by vertical rotation up to 135°, but improved when rotated beyond this to being fully inverted, almost returning to the level of upright faces (this was not tested statistically, however). One possible account of this finding is that while rotation impairs face recognition because faces have to be mentally rotated further (to the upright position) the more they are rotated, participants are able to make use of the symmetry of inverted faces, enabling the faces to be flipped for recognition. An empirical test of this hypothesis can be conducted by looking at recognition performance when all the faces are learned in the upright orientation, and seen at different angles of rotation at test. It is important to note, however, that if we fail to find similar results to those reported in Experiment 2 here, this does not necessarily suggest that mental rotation and flipping strategies cannot account for the findings reported in Experiment 2; it may merely suggest that faces are not being mentally rotated or flipped to a stored memory representation of *that* face, but to a memory for faces as upright stimuli in general.

In one of the first studies to look at the effects of rotation on face perception, Valentine & Bruce (1988; Experiment 1) looked at the difference between upright and rotated faces, in a sequential matching task. A face was presented upright, and was then followed by either another upright or a rotated (45°, 90°, 135°, 180°) face, with either a same or different expression, and participants decided whether the two faces were of the same person or of two visually similar (different) people. They found that reaction times increased linearly as a function of the difference in orientation between the two faces, for same and different expressions, on both 'same' and 'different' trials. These findings lead Valentine & Bruce (1988) to argue that there is not a switch, or a qualitative difference, in processing between upright and inverted faces.

In a second experiment they also found a similar pattern of rotation effects in a face familiarity judgement task in which participants simply made judgements about one face in each trial, suggesting that these linear effects may also be true for face recognition tasks. Valentine & Bruce (1988) acknowledge, however, that although the deviations from linearity in the data were not statistically significant, this does not necessarily indicate that the effect of rotation is linear; it is possible that the procedure used in this study was not strong enough to detect any other possible component (e.g. cubic). Further, there may be a change in processing strategy between upright and inverted faces, but this change may not necessarily be reflected in a change in response times on a same-different matching task.

More recently, Schwaninger & Mast (2005) also used a sequential same-different matching task, this time to look at the effect of rotation on the detection of

configural and featural changes. Again, a face was presented upright, and was then followed by either another upright or a rotated face. Consistent with the predictions of the mental rotation hypothesis, they found no effect of rotation for the accuracy of detecting featural changes between consecutive faces ('different' trials), but configural changes significantly impaired matching accuracy. The hypothesis also predicts that reaction times should increase as a function of the difference in orientation between the two faces, because faces are processed by mentally rotating face parts, and this should be the case for the detection of both configural and featural changes. Again the findings were consistent with these predictions. Schwaninger & Mast (2005) did not, however, appear to find a linear relationship between matching performance and angle of rotation as Valentine & Bruce (1988) reported, although a statistical analysis of the functions was not actually conducted. Analysis of the error data revealed that participants made the greatest number of errors for faces rotated 90° and 120° from upright, with this number dropping for 150° rotated and inverted faces. Indeed, a similar, albeit much less pronounced, pattern was found for the reaction time data, with 90° and 120° rotations yielding longer reaction times than 150° or inverted faces.

Here we address the predictions of the mental rotation hypothesis using a recognition memory paradigm, and also consider the implications of our findings for competing theories of face processing (e.g. the dual-mode hypothesis, the holistic processing hypothesis) and accounts of the inversion effect. As discussed earlier, the issue of linearity in the relationship between rotation and face perception is an important one, and the data appears to remain inconclusive to date. We also aim to address this issue in this experiment. If faces are mentally rotated to match a stored

memory representation (from the learning phase) we would expect recognition accuracy to decline as a function of the difference in orientation between the learned faces and those seen at test. If a flipping strategy is evoked for inverted faces, however, we might expect to find an advantage for inverted faces over rotated faces for recognition accuracy on this task. With regard to featural information, consistent with previous findings (e.g. Experiment 2; Carbon & Leder, 2005), we again expect to find no effect of rotation for featural information. This in turn would produce a whole-face advantage for upright but not inverted faces if the effect of rotation is a function of the difference in orientation between the two faces, while a whole-face advantage may be observed for both upright and inverted faces if the flipping strategy for inverted faces is observed.

3.2.1. Method

3.2.1.1. Participants

Twenty undergraduates from Cardiff University with normal or corrected-to-normal vision received course credits for their participation in this experiment.

3.2.1.2. Materials

The five artificial Caucasian greyscale prototype faces described in Experiment 2 were used. As outlined above, each of the five faces had a unique set of features (head and jaw shape, hairstyle etc), and from each prototype, a set of faces was created by interchanging 3 types of eyes, 3 types of nose and 3 types of mouth to create a set of 27 unique faces for each prototype, 135 faces in all. The same 30 common names given to the faces in Experiment 2 were used. All other details of the materials were the same as in Experiment 2.

3.2.1.3. Procedure

In this experiment, the faces were always seen in the upright orientation in the learning phase. As in Experiment 2, there were five parts to the experiment, this time each representing a different angle of rotation of the stimulus at test (0°, 45°, 90°, 135° and 180°) and each consisting of a learning phase and a test phase. Participants saw each of the six faces from one of the prototype sets with its associated name. Each face and name was shown on the screen for 5 seconds, and was seen twice in a random order. At test, the faces (and isolated features) were seen in one of five different orientations. As in both previous experiments, participants made a forced-choice decision as to which of the two faces or isolated features was associated with the name shown by pressing one of two keys on the keyboard. The prototype set of faces used in each rotation condition and the order in which the different angles were shown were counterbalanced between participants. All other details of the procedure were as described in Experiment 2.

3.2.1.4. Design

A within-participants design was used. The independent variables were the type of stimulus seen at test (whole face or an isolated feature) and the orientation of the stimulus at test (upright, 45°, 90°, 135° and 180°). As in Experiment 2, the task was to decide which of the two faces or features belonged to the given name, and the dependent variable was recognition accuracy (percentage correct).

3.2.2. Results

The proportion of correct identifications for faces and isolated features was calculated for each participant for each angle of rotation at test for both whole-face and isolated-feature targets. The means of these proportions are plotted in Figure 3.3 below.




Figure 3.3. Mean percentage of correct recognition in Experiment 3 as a function of angle of rotation at test and stimulus type. Error bars show ± 1 standard error.

A two-way ANOVA was conducted on the proportion of correct choices, with the two within-participant independent variables being the angle of rotation of the stimulus at test (0° , 45° , 90° , 135° and 180°) and the type of stimuli seen at test (whole faces or isolated features). The analysis found a significant effect of stimulus type ($F(1, 19) = 4.630$; $p < 0.05$), but not of orientation ($F(4, 76) = 1.895$; $p > 0.05$). The interaction between these two variables was also not significant ($F(4, 76) = .594$; $p > 0.05$). Comparing just the upright and inverted faces, however, an *inversion* effect was found for whole faces ($t(19) = 3.031$, $p < 0.05$) but not for isolated features ($t(19) = .248$, $p > 0.05$). In contrast with Experiment 2, a whole-face advantage was

observed for upright faces ($t(19) = 2.105$; $p < 0.05$), but there were no significant differences at any of the other angles of rotation ($p > 0.05$ in all cases).

In terms of the regression analyses, linear and cubic functions were fitted to the data for each condition (e.g. whole faces, isolated features). Comparisons of the ability of the linear and cubic functions to account for the data showed that a cubic function provided a better fit to the data than a linear function for isolated features ($F(2, 1) = 5685.583$) but not for whole faces ($F(2, 1) = 0.244$). This is in contrast to analysis of the data from Experiment 2, which found the opposite pattern of results. The effect of rotation appears to follow a curvilinear, rather than a linear function on this task, but only for isolated features. The effect of rotation on the processing of configural or holistic information for the recognition of whole faces, however, appears to be one of a roughly linear decline in performance as the face is rotated away from the upright, although there is some evidence of non-linearities.

3.2.3. General Discussion

In this experiment we considered the effect of rotation of the stimulus in the test phase when faces had been learnt in the upright orientation. Our findings show a number of consistencies with previous research in which this methodology has been employed. (Valentine & Bruce, 1988; Schwaninger & Mast, 2005). Valentine & Bruce (1988) found that reaction times increased linearly as a function of the difference in orientation between the two faces, an effect which is also apparent for whole faces in this study, while Schwaninger & Mast (2005) found that rotation was detrimental to the detection of configural but not featural changes in a similar

matching task, which is consistent with the differential effects of *inversion* for whole faces and isolated features observed here.

Indeed, these differential effects of inversion suggest that an alternative processing strategy might be being employed on this task. Specifically, the inversion effect for whole faces may be due to at least a partial failure of a mental rotation mechanism (e.g. Rock, 1973, 1988), whereby the face is rotated to match the stored memory representation of that image. In this case, we might expect a continuous decline in performance with increasing angles of rotation up to inversion, which should be significantly different from upright. Our analysis suggests that this is a possible explanation of our findings, but only for whole faces. Furthermore, our findings with regard to the issue of linearity between rotation and face recognition appear to be consistent with these explanations; isolated features were found to be affected in a curvilinear fashion by rotation, whereas the processing of whole faces seems to be most consistent with a linear relationship between rotation and recognition memory.

One possible explanation for the different effects of rotation (and inversion) between this experiment and Experiment 2 is that in our earlier study we looked at the effects of rotation on the processing of equally rotated faces; that is, faces were learnt in the same orientation as they would be seen at test. This experiment has specifically considered the processing of differently rotated faces, with faces learned in the upright view and rotated through different angles of rotation at test, to examine a possible mental rotation strategy being employed in this task. Bauml, Schnelzer & Zimmer (1997) also found no effect of face inversion with equally rotated faces, on a task

which employed a comparable memory component to the recognition memory studies discussed here. Specifically, these authors used a similar paradigm to that developed by Goldstein (1965), where participants learn to associate a pair of successively presented faces (at learning) with a neutral response, while discriminating between other visually similar faces from different pairs. At test, a face is presented and participants attempt to recall the response member when presented with the stimulus member of the pair (Bauml et al, 1997). Participants learned and discriminated between upright, 60°, 120° and 180° rotated faces, but the results showed no effect of orientation on face discrimination. This appears to suggest that the effects of rotation on face perception may be dependent on the degree of similarity between the orientation of the faces in the learning and recognition phases.

In contrast with the findings from Experiment 2, the results from this can be accounted for by both the dual-mode and holistic views of face processing. The whole-face advantage reported for upright faces appears to suggest that configural or holistic information is particularly important in enabling participants to discriminate accurately between the two images shown at test. The different effects of rotation and inversion for whole faces and isolated features are also consistent with both of these theories of the face inversion effect. The holistic view argues that faces are processed as ‘unparsed’ wholes, and that inversion is detrimental to this whole-face processing, while the dual-mode argues that it is configural rather than featural information that is adversely affected by face rotation and inversion. With regard to the mental rotation hypothesis, however, we have not found evidence to support a mental flipping strategy being employed for inverted faces in this task. Rather, our findings seem consistent with an alternate view that recognition performance declines as a function

of the amount of rotation required to reorient the stimulus (e.g. Shepard & Metzler, 1971; Valentine & Bruce, 1988). We suggest that when faces are learned upright and a number of angles of rotation are considered, the loss of configural information for whole faces is more gradual than previous research has suggested, while the features of a face appear to be encoded by a different means of processing, as isolated features are largely unaffected by rotation.

Chapter 4: Examining the Effect of Rotation on Configural Processing in a Face Matching Task

Overview

In this chapter we use a new paradigm, that of a same-different matching task, to look at the effects of rotation on a task which is known to disrupt configural information (the Thatcher illusion; Thompson, 1980). We look at the matching of different types of face pairs, and use same or different-identity pairs as a means of encouraging featural and configural processing respectively. Our findings are in contrast with those of previous research, and a number of accounts of these differences are considered, along with the implications of our findings for the dual-mode and holistic theories of face processing.

4.1. Experiment 4

The results from the experiments reported so far in this thesis have provided an important insight into the effects of rotation on face recognition memory, and suggest that these effects may differ according to the orientation in which a face is encoded. These experiments have considered the effects of rotation on a face recognition task thought to demonstrate holistic processing, and have provided some evidence to suggest that faces may not be processed in a purely holistic manner as previous authors have suggested (e.g. Tanaka & Farah, 1993). In this chapter, we turn to consider the effects of rotation on a face *perception* task which does not involve a recognition memory component to look at the effects of rotation on a task which is thought to disrupt configural information. The Thatcher illusion (Thompson, 1980) is widely thought to provide us with such a stimulus. This is an illusion where the eyes

and mouth are inverted within an upright face, resulting in a grotesque image. If the whole image is then inverted, the changes made to the face are disrupted, and the face no longer appears abnormal or gruesome. This illusion is thought to result from the disruption of configural information produced by inverting the face parts relative to their surrounding context. That is, upright faces are more sensitive to (changes in) configural processing (e.g. Bartlett & Searcy, 1993; Rhodes et al, 1993), whereas with inverted faces, the process of inversion has already disturbed this configural information, making configural processing less efficient.

Boutsen & Humphreys (2003) have recently considered the processing of Thatcherised faces within a same/different matching paradigm. Participants decided whether the features of two faces were the same (both normal or both Thatcherised) or not (one normal, one Thatcherised), in both upright and inverted orientations, for same-identity (Experiment 3) and different-identity (Experiment 4) face pairs. Same-identity pairs, where the two images are identical, can be based on a comparison of individual features, and thus, while still available, configural information becomes less useful in this task, making the matching of individual features more likely. With different-identity pairs, however, matching cannot be based on individual features, as the two images are always of different people, and so matching will therefore require information about the orientation of face parts relative to their context in order to complete the task.. For 'same' decisions, an inversion effect on reaction times was found for normal but not for Thatcherised faces. This was the case for both same-identity and different-identity pairs. For 'different' decisions, responses to inverted faces were significantly slower than those to upright faces (see also Lewis & Johnston, 1997).

Boutsen & Humphreys (2003) suggest that the presence of the face context in Experiment 3 (with normal faces) enabled the face parts to be encoded as part of a configuration, even when a feature-based strategy was encouraged. This is able to explain why an inversion effect was obtained with whole, normal faces but not with Thatcherised faces. In contrast, in Experiment 4 an inversion effect was found for normal whole faces, suggesting that the additional face context does not encourage a configural encoding strategy when matching cannot be image-based.

The absence of inversion effects for Thatcherised face pairs, on the other hand, may point to the presence of the face context in Experiments 3 and 4 having disturbed configural processing, particularly with upright faces, as this renders the context and the Thatcherised parts incompatible. Boutsen & Humphreys (2003) argue that a distinction between two types of configural processing is able to account for these findings. Local configural processing is thought to involve the encoding of the critical parts necessary for identification and their spatial relations, but this encoding is independent of the face context. Global configural processing, refers to the encoding of these parts and their spatial relations within the face context.

With regard to the findings for same-person matching, Boutsen & Humphreys (2003) argued that global configural processing was important for the matching of normal faces on this task. The presence of the face context in their Experiment 3 (with normal faces) enabled the face parts to be encoded as part of a configuration, even when a feature-based strategy was encouraged. These authors argued that this

configural information was global, and that this information is affected by inversion for normal faces, and by Thatcherisation, but only for upright faces.

The lack of an inversion effect with Thatcherised faces was also explained by these authors in terms of the role of global configural processing. Boutsen & Humphreys (2003) argued that inverting the eyes and mouth of a normal face disrupts configural information to the same extent as face inversion. The matching of upright Thatcherised faces is therefore as impaired as the matching of inverted (normal or Thatcherised) faces, so no inversion effect is observed. The face context is argued to have disturbed configural processing for Thatcherised faces, making the context and face pairs incompatible, so global configural processing becomes necessary for the matching of Thatcherised faces.

Concerning the linear or non-linear nature of the effect of rotation, in recent years a number of authors have looked at the effects of rotation on perception of the Thatcher illusion. Murray, Yong & Rhodes (2000) found that there was a discontinuity in bizarreness ratings of Thatcherised faces between 90° and 120° rotations, suggesting that there may be a decrease in the utility of configural information as the face is rotated from the upright view. Lewis (2001) considered reaction times to simply decide that rotated faces were (or were not) Thatcherised, and found a nonlinear increase in RTs as the face was rotated away from upright.

In this experiment we attempt to combine the methodologies of Boutsen & Humphreys (2003) and Schwaninger & Mast (2005) to look more closely at the effect of rotation on configural/featural or holistic information on a same/different matching

task. We look at the effect of rotation on performance with Thatcherised and normal faces with same-identity pairs, which encourages a featural processing strategy, to look at why configural information is difficult to retrieve from inverted faces, with particular reference to the predictions of the mental rotation hypothesis, and also concerning the linear or non-linear nature of the relationship between rotation and face recognition.

Previous research suggests that we should expect to find a greater detrimental effect of rotation for normal faces than for Thatcherised faces on this type of task, regardless of whether the images are same-person or different-person pairs (Boutsen & Humphreys, 2003). This would be consistent with the view outlined earlier that the matching of normal and Thatcherised faces involves global configural processing, and that this information is disrupted by inversion for normal but not Thatcherised faces. However, if same-person matching encourages a featural processing strategy, we might expect to see no effects of inversion (rotation) for normal or Thatcherised faces; featural information is thought to be relatively unaffected by inversion. Although there are a number of methodological differences between our study and that of Schwaninger & Mast (2005), if Thatcherised faces can be compared to their configurally distorted faces, we may expect to find similar effects of rotation to these authors for Thatcherised faces. In their study, when the two faces were configurally distorted, there was a significant effect of rotation on RTs, and this effect was approximately linear in nature, although no statistical analysis was conducted to verify this.

We also look at the effect of rotation on different trials in this simultaneous matching task, which has not previously been considered. We would expect a clear effect of inversion on this task because, as noted earlier, an inverted Thatcherised face and an inverted normal face look very similar to each other. By considering the intermediate angles of rotation, however, we aim to investigate whether the two faces become progressively more similar over rotation, as may be indicated by a linear increase in reaction times, or whether there is a definite point at which processing of the two faces becomes similar, in which case a sharp increase in RT would be expected. Schwaninger & Mast (2005) found evidence consistent with an abrupt change in processing in the detection of configural changes in a sequential matching task. Participants made the greatest number of errors for faces rotated 90° and 120° from upright, with this number dropping for 150° rotated and inverted faces, and a similar pattern was also found for RTs. Thus on different trials we may expect to find a definite switch in accuracy rates and reaction times as the faces are rotated away from the upright.

Thus a strict interpretation of the dual-mode view would argue that it is configural information which is affected by face rotation and inversion, and therefore if 'same' decisions are made on the basis of featural information (which is unaffected by inversion), then there should be no detrimental effect of inversion on face matching performance with same-person pairs, for normal or Thatcherised faces. If, as Boutsen & Humphreys (2003) suggest, global configural information is being used to complete the match, however, we may expect to find inversion effects for normal but not Thatcherised faces.

The predictions of the holistic hypothesis, on the other hand, are not quite as clear. A strict interpretation would appear to suggest that Thatcherisation should have no effect on face matching, as faces are encoded as wholes, rather than by an analysis of their composite parts. There should, therefore, be no difference between normal and Thatcherised faces, whether same or different-person pairs are seen. If 'holistic' is defined as whether or not the Gestalt of the overall image is coherent, however, this suggests a different prediction. With inverted Thatcherised faces, while the eyes and mouth are seen in the correct orientation, the Gestalt of the overall image is not coherent. Normal inverted faces, in contrast, have a coherent Gestalt, but the eyes and mouth are inverted, which may hinder featural processing of these faces.

Carbon & Leder (2005) looked at the role of featural information in the early stages of face recognition, but the same arguments may be able to be applied here. They found faster reaction times for recognition of inverted normal faces than inverted Thatcher faces over longer presentation times, while in a second experiment, reaction times to correctly recognise upright Thatcher faces were significantly longer than those for normal faces. While the effects of inversion were not directly tested in this study, Carbon & Leder (2005) concluded that more integrated, holistic processing is involved in the recognition of briefly presented upright faces, whereas the early encoding of inverted faces is of a more featural nature.

Thus, if featural information is important in perceptual matching tasks, inverted Thatcher faces should be processed faster than inverted normal faces, because the key featural information (the eyes and mouth) is already correctly oriented. If face processing is holistic in nature, however, one might expect inverted

normal faces to be processed faster, because they have a more coherent holistic Gestalt than inverted Thatcher faces. This version of the holistic processing hypothesis may therefore predict that there will be significant differences between upright normal and Thatcherised faces, and between inverted normal and Thatcherised faces, but it remains unclear what the overall effect of inversion would be for each type of face.

4.1.1. Method

4.1.1.1. Participants

Twenty undergraduates from Cardiff University with normal or corrected-to-normal vision received course credits for their participation in this experiment.

4.1.1.2. Materials

Six full frontal, colour faces (4 male, 2 female) obtained from the Psychological Image Collection at Stirling (PICS; <http://pics.psych.stir.ac.uk>) were used in this experiment. Each face was of approximately similar size, and each image was resized to measure 283 x 346 pixels on the screen. The images were centred vertically on the screen, and were 500 pixels apart horizontally. The eyes and mouth were selected and inverted separately for each face. The normal and Thatcherised upright faces were then rotated through the remaining four angles (45°, 90°, 135° and 180°). The faces were used to create a set of 120 face pairs, 60 for each type of response (same/different), 12 for each angle of rotation. Each face pair was repeated 3 times, making 360 trials in the experiment. On 'different' trials (when a normal and a Thatcherised face were paired together), the position of the normal face was counterbalanced between trials, such that it appeared on the left of the screen (and the

Thatcherised face on the right) in half of the trials, and on the right of the screen (with the Thatcherised face on the left) in the other half. All images were presented on a computer with a screen resolution of 800 x 600 pixels using Superlab, and subtended a visual angle of approximately 5°. Examples of the stimuli used in the two experiments are shown in Figure 4.1.




Figure 4.1. Examples of the stimuli used in the two experiments. The top row shows same-person normal and Thatcherised face pairs (upright and inverted). The bottom row shows a 'different' pair (one normal, one Thatcherised) and a different-person pair (both of the same type).

4.1.1.3. Procedure

In each trial, participants saw two faces on the screen. These faces were either both normal, both Thatcherised, or one normal and one Thatcherised. Participants were asked to indicate whether they thought the two images were the same or different, by pressing one of two keys on a keyboard. In each trial, the face pair was presented until the participant made a response, followed by a 1 second blank interval

before the next trial. Participants were asked to respond as quickly and accurately as possible.

4.1.1.4. Design

The two within-participants independent variables were the orientation of the faces (upright, 45°, 90°, 135° and 180°) and whether the two faces were the same or different (both normal, both Thatcherised, or one normal, one Thatcherised). The dependent variables were the accuracy and latency of responses to the same/different decision for same face pairs.

4.1.2. Results

Separate analyses were performed for the same and different pair trials. Only RTs for correct responses and for those not exceeding three standard deviations from the mean RT for each participant were analysed. The mean percentage of RTs which were excluded for being more than 3 standard deviations from the mean was 9.33% for normal faces and 11.5% for Thatcherised face pairs. The mean correct RTs for responses to normal and Thatcherised faces and the percentage of errors for each type of face were calculated for each angle of rotation. On same pair trials, an error was a 'different' response to a pair of normal faces or a pair of Thatcherised faces, while on different pair trials, an error was a 'same' response to a pair consisting of a normal face and a Thatcherised face. The mean RTs for same pairs are shown in Figure 4.2 below.




Figure 4.2. Mean reaction times for same pairs in Experiment 4 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA was conducted on the mean RTs shown in Figure 4.2. There was a significant effect of orientation on performance ($F(4, 76) = 11.667; p < 0.05$), but no significant effect of face type ($F(1, 19) = 2.452; p > 0.05$) and no significant interaction between orientation and face type ($F(4, 76) = .699; p > 0.05$). There was a slight trend for a decrease in RTs for normal faces between 135° and 180° while Thatcherised faces showed a continuous increase over the same distance, but this decrease for normal faces was not significant ($t(19) = 1.626; p > 0.05$), and the decrease for normal faces did not produce a significant difference between normal and Thatcherised faces in the inverted condition ($t(19) = 1.299; p > 0.05$). These results suggest that rotation had a detrimental effect on same-pair face-matching performance, and this effect was approximately the same for normal and Thatcherised faces.




Figure 4.3. Percentage of errors for same pairs in Experiment 4 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA was conducted on the error data shown in Figure 4.3. There was a significant effect of both rotation ($F(4, 76) = 2.955; p < 0.05$) and face type ($F(1, 19) = 23.699; p < 0.05$), with Thatcherised face pairs producing more errors than normal face pairs, but the interaction between rotation and face type ($F(4, 76) = 1.420, p > 0.05$) was not significant. The difference between normal and Thatcherised face pairs was only significant in the upright ($t(19) = 3.387; p < 0.05$) and inverted ($t(19) = 3.907; p < 0.05$) orientations.

The mean RTs and percentage of errors for different pairs are shown in Figures 4.4 and 4.5 respectively. The mean percentage of RTs which were excluded for being more than 3 standard deviations from the mean was 8.22%.




Figure 4.4. Mean reaction time for different pairs in Experiment 4 as a function of angle of rotation. Error bars show ± 1 standard error.




Figure 4.5. Percentage of errors for different pairs in Experiment 4 as a function of angle of rotation. Error bars show ± 1 standard error.

A repeated measures ANOVA showed a significant effect of rotation on RT performance ($F(4, 76) = 17.673; p < 0.05$). A series of paired-sample t-tests revealed significant differences between 0° and all other angles of rotated faces, and between each angle of rotation (i.e. between 0° and 45° , 45° and 90° , 90° and 135° and 135° and 180° ; $p < 0.05$ in all cases). Figure 4.4 shows a steady increase in RTs between 0°

and 90°, with a trend for a sharp increase in RTs between 90° and 135°, but a paired-samples *t*-test revealed that this increase in RTs was not significantly larger than that between 0° and 90° ($t(19) = 1.031; p > 0.05$), or larger than that between 45° and 90° ($t(19) = 1.765; p > 0.05$), suggesting that for RTs there is a progressive effect of rotation on this task.

A repeated measures ANOVA of the error data (Figure 4.5) also revealed a significant effect of rotation ($F(4, 76) = 28.839; p < 0.05$), but showed an abrupt increase in errors for faces rotated beyond 90° ($t(19) = 5.792; p < 0.05$). The increase in errors between 135° and 180° was also significant ($t(19) = 2.666; p < 0.05$). These analyses suggest that reaction times to detect that one face in an otherwise identical pair is Thatcherised become progressively longer with increasing rotation away from the upright, with a significant increase in RTs between each angle of rotation, but in terms of accuracy there is a clear point of rotation beyond which participants find it difficult to detect that the two faces are different.

4.1.3. Discussion

The results of this experiment show a number of differences with previous research. Boutsen and Humphreys (2003) found that the correct matching of two Thatcherised faces yielded significantly longer RTs than did the matching of two normal faces. Although this was also true of our findings, the difference between normal and Thatcherised faces was not significant in our study. Moreover, a significant effect of rotation was found for both normal and Thatcherised faces, suggesting that rotation may affect normal and Thatcherised faces in similar ways. This is in contrast to the findings of Boutsen and Humphreys (2003) with same-face

pairs which showed that although RTs were slower for (upright) Thatcherised pairs, they were not significantly reduced by inversion.

Boutsen and Humphreys (2003) explained this lack of an inversion effect with Thatcherised faces by arguing that inverting the eyes and mouth of a normal face disrupts configural information to the same extent as face inversion. The matching of upright Thatcherised faces is therefore as impaired as the matching of inverted (normal or Thatcherised) faces, so no inversion effect is observed. However, studies such as that of Schwaninger and Mast (2005) have shown almost equivalent RTs and accuracy scores for upright configurally distorted as for featurally distorted faces on 'same' trials (when the two faces were identical); and Carbon, Schweinberger, Kaufmann & Leder (2005) and Carbon and Leder (2005) have also found effects of inversion for Thatcherised faces in recognition tasks. Carbon et al (2005) found an inversion effect for Thatcherised faces for the N170 ERP amplitude for familiar-face recognition, whereas Carbon and Leder (2005) found a similar effect for RTs in a recognition task over long (but not short) presentation times. Therefore, it may be that in a simultaneous matching task with same-person pairs, a featural strategy is encouraged, and this is the case for both normal and Thatcherised face pairs. Moreover, it would appear that featural information is disrupted by inversion on this task, a finding which is not predicted by the dual-mode hypothesis. It is important to note, however, that other methodological differences may also account for this difference between the two studies. The difference may be due, for example, to the differential salience of the feature inversion (used to produce the Thatcherised faces) between the two studies. It is possible that the Thatcherised faces were harder to

detect or less discriminable from normal faces in the Boutsen and Humphreys (2003) study, which may have given rise to the longer RTs for upright Thatcherised faces.

Thus, in terms of the predictions of the different accounts of the face-inversion effect, these findings are not supported by the dual-mode hypothesis, or by the idea that global configural information is important for this task but is differentially affected by inversion for normal and Thatcherised faces. The holistic hypothesis is able to account these findings, however. This view suggests that Thatcherisation should have no effect on face matching, as faces are encoded as wholes, rather than by an analysis of their composite parts. Inversion (rotation) should effect face matching, but there should be no difference between normal and Thatcherised faces, whether same or different-person pairs are seen.

Nevertheless, as they stand, our findings are consistent with the predictions of the mental-rotation hypothesis. When there is no change to be detected (same trials), the hypothesis predicts that there will be no difference between normal and Thatcherised face pairs, because trials always contain identical stimuli, and participants would mentally rotate both types of stimuli to determine that they were the same. On different trials, however, rotation should have a substantial effect on both accuracy scores and RTs, and again this pattern of results was observed. The findings are also largely consistent with the view that a same-person matching task, when the two faces are identical, encourages a feature-based matching strategy, and that this is the case for both normal and Thatcherised faces.

4.2. Experiment 5

Boutsen & Humphreys (2003) also looked at the matching of normal and Thatcherised faces in different-identity face pairs, where the two faces depicted different individuals, as a means of encouraging a more configural processing strategy to be employed. The task was again to decide as quickly and accurately as possible whether the two faces were of the same type (both normal or both Thatcherised) or different (one normal and one Thatcherised). There were no main effects of face type or of orientation for correct 'same' decisions, but normal faces again showed an inversion effect, while Thatcherised faces were responded to faster in the inverted condition than in the upright condition, yielding a significant interaction. A similar pattern of results was found with different-person face parts (Experiment 2), suggesting that the additional face context (i.e. global configural information) did not facilitate encoding when participants could not rely on an image-based matching strategy. For 'different' decisions, responses to inverted faces were significantly slower than those to upright faces. Again, Boutsen & Humphreys (2003) suggested that this inversion effect for normal but not for Thatcherised faces highlights a qualitative difference in the encoding of these two types of faces.

The absence of inversion effects for Thatcherised face pairs in this task is again consistent with the view that for upright Thatcherised faces, the face context disturbs configural processing, leaving the context and the Thatcherised parts incompatible. This produces longer RT's than for normal faces. Indeed, the effect of face context was particularly apparent in this experiment, as there was a trend for faster reaction times to inverted than to *upright* Thatcherised face pairs. Boutsen & Humphreys (2003) argued that this task encouraged local configural processing and

this information is disrupted by inversion for normal faces, and by Thatcherisation, but for upright faces only. These findings suggest that if there is a qualitative difference in the encoding of normal and Thatcherised faces on a different-person face matching task, we might expect similar results to those of Boutsen & Humphreys (2003); an interaction between face type and orientation, with a positive effect of inversion for Thatcherised faces, and an inversion effect for normal face matching.

Carbon et al (2005), however, have shown effects of inversion for Thatcherised faces, with significant differences N170 amplitudes between inverted and normal Thatcherised faces. While there are a number of differences between our study and those of Carbon et al (2005), in light of the findings from Experiment 4, it is possible that we may again find different effects of Thatcherisation and rotation to those of Bousten & Humphreys (2003). Indeed, previous perceptual matching studies have found significant inversion effects with spatially distorted (e.g. Searcy & Bartlett, 1996) and Thatcherised faces (e.g. Schwaninger & Mast, 2000, 2005), and these effects have been either larger than or in the complete absence of inversion effects for featurally distorted faces. This would appear to suggest that the effects of inversion and Thatcherisation might be additive on this task, with inversion (rotation) having a greater effect on Thatcherised than on normal face pairs.

4.2.1. Method

4.2.1.1. Participants

Twenty undergraduates from Cardiff University with normal or corrected-to-normal vision received course credits for their participation in this experiment.

4.2.1.2. Materials

The same stimuli detailed in Experiment 4 were used in this experiment. The six faces were used to create a set of face pairs, with each face paired with each other face once, making 15 pairs for each angle of rotation, 150 pairs for each type of response (same/different). Each face pair was repeated 3 times, making 900 trials in the experiment. On 'different' trials (when a normal and a Thatcherised face were paired together), the position of the normal face was counterbalanced between trials, such that it appeared on the left of the screen (and the Thatcherised face on the right) in half of the trials, and on the right of the screen (with the Thatcherised face on the left) in the other half, such that there were the same number of 'same' and 'different' trials. All images were presented on a computer using Superlab.

4.2.1.3. Procedure

In each trial, participants saw two faces on the screen. These faces were either both normal, both Thatcherised, or one normal and one Thatcherised. Participants were asked to indicate whether they thought the two images were *of the same type* (both normal or both Thatcherised) or different (one normal, one Thatcherised), by pressing one of two keys on a keyboard. In each trial, the face pair was presented until the participant made a response, followed by a 1 second blank interval before the next trial. Participants were asked to respond as quickly and accurately as possible.

4.2.1.4. Design

The two within-participants independent variables were the orientation of the faces (upright, 45°, 90°, 135° and 180°) and whether the two faces were the same or different (both normal, both Thatcherised, or one normal, one Thatcherised). The

dependent variables were the accuracy and latency of responses to the same/different decision for same face pairs.

4.2.2. Results

The results were calculated in the same way as in Experiment 4. The mean RTs and percentage of errors for same pairs (both normal or both Thatcherised) are shown in Figures 4.6 and 4.7 respectively. The mean percentage of RTs which were excluded for being more than 3 standard deviations from the mean was 1.2% for normal faces and 2.5% for Thatcherised face pairs.




Figure 4.6. Mean reaction times in Experiment 5 for same pairs as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA was conducted on the mean RTs shown in Figure 4.6. There was a significant effect of orientation on performance ($F(4, 76) = 15.703; p < 0.05$), a significant effect of face type ($F(1, 19) = 44.046; p < 0.05$) and a significant interaction between orientation and face type ($F(4, 76) = 7.696; p < 0.05$). There was no significant difference between upright normal and upright Thatcherised

faces ($t(19) = .982$; $p > 0.05$), but significant differences emerged between normal and Thatcherised faces at all other angles of rotation ($p < 0.05$ in all cases). For normal faces, there was no significant difference between upright and 45° or upright and 90° rotated faces ($t(19) = 0.733$; $p > 0.05$; $t(19) = 1.801$; $p > 0.05$ respectively), but significant differences emerged between upright and 135° and upright and inverted faces ($t(19) = 2.670$; $p < 0.05$; $t(19) = 3.647$; $p < 0.05$ respectively). There were significant differences between 45° and 90°, 90° and 135° and 135° and 180° rotated faces ($p < 0.05$ in all cases). For Thatcherised faces, there was no significant difference between 0° and 45° rotated faces ($t(19) = 1.190$; $p > 0.05$), but significant differences were found between upright and all other angles of rotation ($p < 0.05$ in all cases). Significant differences were also found between 45° and 90°, and 90° and 135° rotated faces ($t(19) = 5.247$; $p < 0.05$; $t(19) = 2.911$; $p < 0.05$), but not between 135° and 180° ($t(19) = 1.370$; $p > 0.05$).




Figure 4.7. Percentage of errors in Experiment 5 for same pairs as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA conducted on the error data in Figure 4.7 also showed significant effects of rotation ($F(4, 76) = 24.637; p < 0.05$), face type ($F(1, 19) = 114.553; p < 0.05$) and an interaction between face type and rotation ($F(4, 76) = 18.939; p < 0.05$), again showing that the effect of rotation was much greater for Thatcherised than for normal face pairs. The difference between normal and Thatcherised faces was significant at all angles of rotation ($p < 0.05$ in all cases). For normal faces, only inverted faces differed significantly from upright faces ($t(19) = 3.336; p < 0.05$). Significant differences were also found between 90° and 135° rotated faces, and 135° and inverted faces ($t(19) = 2.203; p < 0.05; t(19) = 2.559; p < 0.05$ respectively). For Thatcherised faces, there was no significant difference between upright and 45° rotated faces ($t(19) = .376; p > 0.05$), but significant differences were found between upright and all other angles of rotation ($p < 0.05$ in all cases). In addition, significant differences were found between 45° and 90°, and 90° and 135° rotated faces, but the decline in errors for Thatcherised face pairs between 135° and 180° was not significant ($t(19) = 1.640; p > 0.05$).

The data for same pairs, therefore, shows consistent effects of rotation. This effect was much greater for Thatcherised than for normal face pairs, and Thatcherised face pairs were consistently affected by rotations beyond 45° from the upright, whereas normal face pairs were affected by 90° rotations.

We note from these analyses, however, that while the RT data seems to suggest a steady increase in response times as the faces are rotated towards inverted, the error data shows an improvement in performance between 135° and 180°. We therefore also performed a combined analysis of the RT and error data to look at the

effects of rotation on matching performance overall. This is a method suggested by Townsend & Ashby (1983), in which the mean RT for each participant is divided by the proportion of correct responses for each condition. The sum of the mean RTs divided by the proportion of correct responses for 'same' decisions is shown in Figure 4.8.




Figure 4.8. Mean reaction time divided by proportion correct for same pairs in Experiment 5 as a function of angle of rotation. Error bars show ± 1 standard error.

A repeated measures ANOVA on this data showed significant effects of orientation ($F(1, 19) = 26.035; p < 0.05$), face type ($F(1, 19) = 89.003; p < 0.05$) and a significant interaction between face type and orientation ($F(1, 19) = 18.150; p < 0.05$). Linear and cubic functions were also fitted to the data to compare the ability of these functions to account for the data. A cubic function was not found to provide a better fit to the data than a linear function for normal or Thatcherised faces ($F(2, 1 < 199.50)$). This suggests that the overall effect of rotation on performance in the

matching of 'different' face pairs is one of a roughly linear decline in performance as the face is rotated away from the upright.

The mean RTs and percentage of errors for different pairs (one normal, one Thatcherised) are shown in Figures 4.9 and 4.10 respectively. The mean percentage of RTs excluded for being more than 3 standard deviations from the mean was 1.6%.

A repeated measures ANOVA was conducted on the mean RTs shown in Figure 4.9. There was a significant effect of rotation on performance ($F(4, 76) = 16.983; p < 0.05$). There was a small increase in RTs between 0° and 45° , with a larger increase between 45° and 90° , and a paired-samples t -test revealed that this increase was significantly larger than that between 0° and 45° ($t(19) = 2.847; p < 0.05$). The increase between 90° and 135° was also significant ($t(19) = 3.420; p < 0.05$), but not that between 135° and 180° ($t(19) = 1.812; p > 0.05$). This appears to suggest that when one is required to detect that one face in a *non-identical* pair is Thatcherised, rotation of the faces to 90° makes this task significantly more time consuming, and this may indicate that a different processing strategy is being employed for faces rotated to this point.




Figure 4.9. Mean reaction time in Experiment 5 for different pairs as a function of angle of rotation. Error bars show ± 1 standard error.




Figure 4.10. Percentage of errors in Experiment 5 for different pairs as a function of angle of rotation. Error bars show ± 1 standard error.

A repeated measures ANOVA of the error data in Figure 4.10 also revealed a significant effect of rotation ($F(4, 76) = 37.323$; $p < 0.05$) but, as in Experiment 4, showed an abrupt increase in errors for faces rotated beyond 90° ($t(19) = 5.888$; $p < 0.05$). The increase in errors between 135° and 180° was also significant ($t(19) = 4.663$; $p < 0.05$). These data appear to suggest that there is a clear point beyond which

participants find it difficult to detect that the two faces are of a different type, providing further support for our findings from Experiment 4.

We note from these analyses, however, that the pattern of errors for 'different' pairs is very flat up to 90° before rising sharply, whereas the pattern for the RT data appears to be of a more gradual increase as the faces are rotated away from the upright. We therefore suggest that it would be helpful to perform a combined analysis of the RT and error data to look at the effects of rotation on matching performance overall. This is a method suggested by Townsend & Ashby (1983), in which the mean RT for each participant is divided by the proportion of correct responses for each condition. The sum of the mean RTs divided by the proportion of correct responses for 'different' decisions is shown in Figure 4.11.




Figure 4.11. Mean reaction time divided by proportion correct for different pairs in Experiment 5 as a function of angle of rotation. Error bars show ± 1 standard error.

A repeated measures ANOVA on this data showed a significant effect of rotation ($F(4, 76) = 12.423; p < 0.05$). More importantly here, linear and cubic functions were fitted to the data to compare the ability of these functions to account for the data. A cubic function was not found to provide a better fit to the data than a linear function ($F(2, 1) < 199.50$). This suggests that the overall effect of rotation on performance in the matching of 'different' face pairs is one of a roughly linear decline in performance as the face is rotated away from the upright.

Comparing Experiments 4 and 5

A between-experiments comparison was performed to see if there was any difference in the effect of rotation between same-person and different-person face pairs. Separate analyses were conducted for 'same' and 'different' responses. A mixed ANOVA was conducted on the 'same' decision data, and showed significant effects of rotation ($F(4, 152) = 25.040; p < 0.05$) and face type ($F(1, 38) = 35.091; p < 0.05$) and significant interactions between rotation and face type ($F(4, 152) = 5.811; p < 0.05$), face identity (experiment) and rotation ($F(4, 152) = 4.447; p < 0.05$) and face identity and face type ($F(1, 38) = 14.255; p < 0.05$). The three-way interaction between these variables was also significant ($F(4, 152) = 4.250; p < 0.05$). The main effect of identity between the two experiments, however, was not significant ($F(1, 38) = 0.371; p > 0.05$). Thus, the interactions revealed a greater effect of rotation for Thatcherised than for normal face pairs, and for different-identity than for same-identity face pairs. For same-person pairs, a series of paired-sample *t*-tests found significant differences between 0° and 45°, 45°, and 90°, 90° and 135° but not 135° and 180° rotated faces ($t(39) = 0.264; p > 0.05; p < 0.05$ in all other cases). For different-person pairs, only the difference between 0° and 45° rotated faces was not



significant ($t(59) = 0.117; p > 0.05$); significant differences were found between all other angles of rotation.

A mixed ANOVA of the 'different' decision data for same- and different-person pairs also showed a significant effect of rotation ($F(4, 152) = 34.255; p < 0.05$). The interaction between rotation and experiment was not significant ($F(4, 152) = 0.204; p > 0.05$). Here the main effect of identity between the two experiments was significant ($F(1, 38) = 4.196; p < 0.05$). A series of paired sample t -tests revealed that the difference between same-person and different-person pairs emerged for upright ($t(19) = 2.259; p < 0.05$) and 90 rotated faces ($t(19) = 2.388; p < 0.05; p > 0.05$ in all other cases). For same-person pairs, significant differences were found between all angles of rotation between 0° and 45°, 45° and 90°, 90° and 135° and 180°; $p < 0.05$ in all cases). For different-person pairs, only the differences between 0° and 45° ($t(19) = 1.291; p > 0.05$) and 135° and 180° ($t(19) = 1.812; p > 0.05$) rotated faces were not significant.

In terms of the regression analyses, the data from the two experiments was combined for each condition (e.g. normal, Thatcherised, different) and linear and cubic functions were fitted to the data. Comparisons of the ability of the linear and cubic functions to account for the data showed that for normal face pairs, a cubic function provided a better fit to the data than a linear function ($F(2, 1 = 1422)$), but this was not the case in any of the other conditions. Thus, the overall effect of rotation in these experiments has been one of a roughly linear decline in performance as the face is rotated away from the upright, although there is some evidence of non-linearities.

4.2.3. Discussion

The results of this experiment show a number of differences with our findings from Experiment 4 and with Boutsen & Humphreys (2003) work. Here we found that the matching of Thatcherised faces yielded significantly longer reaction times than did the matching of normal faces. We also found a significant effect of rotation, and an interaction between face type and rotation. Boutsen & Humphreys (2003) also reported such an interaction, but showed that inversion affected normal and Thatcherised faces in opposite directions; our results appear to suggest that the direction of the effect is the same, but rotation has a more detrimental effect on the matching of Thatcherised face pairs than normal face pairs. Thus it would appear that Thatcherisation disproportionately disrupts the matching of face pairs when the task cannot rely on image-based matching. Furthermore, Thatcherised face pairs were consistently affected by rotation beyond 45° from the upright, whereas normal face pairs were affected by 90° rotations. The data from different trials were largely consistent with those of Experiment 4, although different pairs were not significantly affected by 45° rotations in this experiment. Nevertheless, when one is required to detect that one face in a *non*-identical pair is Thatcherised, rotation of the faces to 90° makes this task significantly more time consuming, and this may indicate that a different processing strategy is being employed for faces rotated beyond this point. There was an overall difference between same-person and different-person matching, with the latter producing longer reaction times and a greater percentage of errors. The error data was also consistent with our previous findings, suggesting that there is a clear point (once the face is rotated beyond 90°) beyond which participants find it difficult to detect that the two faces are of a different type.

Comparing the two experiments, for 'same' responses there was an interaction between rotation and identity, with the effect of rotation being greater for different-identity pairs than for same-identity pairs. Meanwhile, for 'different' responses the effect of rotation was of approximately the same magnitude for same and different-person pairs, but different-person pairs produced consistently longer reaction times and more errors, suggesting that this task is particularly difficult when matching cannot be purely image-based. We will now discuss the findings from these experiments, their differences (and similarities) with previous research, and their implications for accounts of the face inversion effect.

4.3. General Discussion

In these experiments we have attempted to look at the effects of rotation on configural information and to examine why configural information is difficult to retrieve from rotated faces. On a same-person matching task, which could be achieved by matching image-based featural information, an effect of rotation was observed for both normal and Thatcherised face pairs, and this effect was approximately the same for both types of faces. Thatcherised face pairs produced more errors than normal faces, however. On different-person trials, RTs became progressively longer with increasing rotation away from the upright, whereas accuracy scores were significantly affected by rotations beyond 90°. When the two faces depicted different people (Experiment 5), an effect of rotation was again observed for normal and Thatcherised face pairs, but the effect was significantly greater for Thatcherised than for normal faces, and this was true both of RTs and of accuracy scores. On different trials, the data mirrored that of Experiment 4, with RTs increasing with rotation from the

upright, and accuracy scores significantly affected by rotations beyond 90°. Both of these effects were more pronounced for different-person face pairs, and the effect of rotation was also greater for different-identity pairs than for same-identity pairs for 'same' responses.

Perhaps the most surprising finding here, in light of previous research, is that of an inversion effect for Thatcherised face pairs, and this was apparent for both same-person and different-person face matching. We argue that the mental-rotation hypothesis is able to account, at least in part, for this finding. When there is no change to be detected (same trials), the hypothesis predicts that there should be no difference between normal and Thatcherised face pairs, because trials always contain identical stimuli, and participants would mentally rotate both types of stimuli to determine that they were the same. The data from different trials were also consistent with the hypothesis, with rotation conferring a substantial effect on both accuracy scores and RTs.

The pattern of findings for 'same' responses with different-person pairs, however, is more consistent with holistic processing being used to perform this task. When 'holistic' is defined as whether or not the Gestalt of the overall image is coherent, the holistic hypothesis predicts that inverted normal faces will be processed faster than inverted Thatcher faces, as they have a more coherent holistic Gestalt (Carbon & Leder, 2005). Here, the effect of rotation was significantly greater for Thatcherised than for normal faces, and this was true of both RTs and accuracy scores. Thus, when matching cannot be based on image-based featural information, participants may revert to a holistic processing strategy, which is able to account for

the greater effect of rotation for Thatcherised than for normal face pairs. There was, however, no significant difference between upright normal and Thatcherised faces in this experiment, suggesting that this version of the holistic hypothesis is limited in the extent to which it is able to account for the current findings.

The findings presented here do not appear to support the distinction drawn by Boutsen & Humphreys (2003) between local and global configural information. We found little evidence to suggest that global configural information is involved in the matching of normal and Thatcherised faces, and that this information is disrupted by inversion for normal but not Thatcherised faces. Rather, our findings suggest that when featural information can be used to complete the match (i.e. when the face pairs contain identical images of the same person), the effect of rotation is independent of face type. When different-identity or different-image pairs are shown and configural information is required, however, Thatcherised faces are affected more by rotation than normal faces; rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces.

We have also considered the question whether the relationship between rotation and face recognition is of a linear or nonlinear nature. Although we are unable to conclude that there is definitely a linear relationship between rotation and face-matching performance in these tasks, the majority of the data suggested that a linear function provided the best account of the data. These findings are in broad agreement with the third prediction of the mental-rotation hypothesis mentioned in the introduction: that recognition performance is dependent on the amount of normalisation required. Moreover, our findings provide an insight into the effect of

rotation on face-recognition tasks, where previous studies reporting a nonlinear effect have looked at bizarreness ratings of Thatcherised faces, or the simple detection of a Thatcherised face.

Though only our study and that of Boutsen and Humphreys (2003) have looked at 'same' decisions in these types of matching types with Thatcherised faces (ie matching on the basis of image type), our findings are nevertheless consistent with other research, which shows that inversion is disruptive to the processing of Thatcherised faces (e.g. Murray et al, 2000; Sturzel & Spillmann 2000; Lewis, 2001). We propose an alternative explanation of the effect of configural information on face matching: that Thatcherisation disproportionately disrupts the matching of face pairs in a task which encourages configural processing (matching different-identity face pairs), but when matching can be image-based, as in Experiment 4, the loss of configural information through rotation is no greater for Thatcherised face pairs than for normal pairs. Our findings appear to suggest that there may not be a qualitative difference between the encoding of Thatcherised and normal faces as such, but that different tasks involve different processing or encoding strategies, and that, when configural encoding is required, Thatcherised faces are affected more by rotation than normal faces.

Chapter 5: Face rotation in a matching task: exploring the effects of inversion, the similarity of the image and identity matching

Overview

This chapter considers three further questions suggested by the face matching paradigm developed by Bousten & Humphreys (2003). One possible explanation for the discrepancies between our findings and those of previous research is that seeing the face pairs at a number of angles of orientation between upright and inverted somehow encouraged a different processing strategy to be employed than when only upright and inverted faces were seen. This question is addressed in Experiment 6. We then consider the effects of rotation, and the processing strategy employed when different images of the *same* person are to be matched (Experiment 7), before finally looking at the effects of rotation on matching of facial identity rather than face type in Experiment 8.

5.1. Experiment 6

The findings reported in Chapter 4 have shown a number of differences with previous research on a face matching task with same- and different-identity face pairs. Bousten & Humphreys (2003) found that for 'same' (both normal or both Thatcherised) decisions, there was an inversion effect on reaction times for normal but not for Thatcherised face pairs, and this was the case for both same-identity and different-identity pairs. For 'different' decisions, responses to inverted faces were significantly slower than those to upright faces (see also Lewis & Johnston, 1997). In our experiments, in which we have also looked at the intermediate angles of rotation between upright and inverted, however, we found different effects of rotation for

same-identity and different-identity face pairs. When a same-identity matching task was employed, an effect of rotation was observed for both normal and Thatcherised face pairs, and this effect was approximately the same for both types of faces. On different trials, reaction times became progressively longer with increasing rotation away from the upright, while accuracy scores were significantly affected by rotations beyond 90°.

When the two faces depicted different people (Experiment 5), however, an effect of rotation was again observed for normal and Thatcherised face pairs, but the effect was significantly greater for Thatcherised than for normal faces, and this was true of both RTs and accuracy scores. On different trials, the data mirrored that of previous research, with reaction times increasing with rotation from the upright, and accuracy scores significantly affected by rotations beyond 90°. Both of these effects were more pronounced for different-person face pairs. A between-experiments comparison showed that the effect of rotation was greater for different-identity pairs than for same-identity pairs for 'same' responses, but was roughly the same for both types of pairs for correct 'different' decisions.

We have suggested that our findings indicate that Thatcherisation disproportionately disrupts the matching of face pairs when the task encourages configural processing (matching different-identity face pairs), but when matching can be image-based as in Experiment 4, the loss of local configural information through rotation is no greater for Thatcherised face pairs than for normal pairs. This appears to be at odds with the view offered by Boutsen & Humphreys (2003) that there is a qualitative difference between the encoding of Thatcherised and normal faces.

It is possible, however, that the differences between our findings and those of Boutsen & Humphreys (2003) with regard to the overall effect of inversion are due to the fact that we considered a number of different angles of rotation in our study, rather than just upright and inverted faces. That is, seeing the face pairs at a number of angles of orientation between upright and inverted may have encouraged a different processing strategy, and produced different effects of inversion to those observed in previous research.

One way of testing this possibility is to use a similar same-different matching task using only upright and inverted faces. If, when only upright and inverted faces are seen, the effects of inversion are similar to those of Boutsen & Humphreys (2003), we would expect to find a detrimental effect of inversion for normal but not for Thatcherised face pairs (indeed we may even expect to find an improvement in performance for inverted Thatcherised faces, although this effect was not significant in their study). This would suggest that our previous findings with different-person pairs (Experiment 5) were due to the fact that we considered a number of intermediate angles of rotation in our study. If comparable effects of inversion are observed in the two studies (Experiment 5 and 6), however, this would suggest that the differences in the effect of inversion between these studies and those of Boutsen & Humphreys (2003) could not be accounted for simply by the additional angles of rotation considered. Moreover, the observation of an interaction between face type and rotation, as found in Experiment 5 with rotated faces, would provide further support for our hypothesis that Thatcherisation disproportionately disrupts the matching of face pairs when the task encourages configural processing.

5.1.1. Method

5.1.1.1. Participants

Twenty undergraduates from Cardiff University with normal or corrected-to-normal vision received course credits for their participation in this experiment.

5.1.1.2. Materials

The upright and inverted stimuli detailed in Experiment 4 were used in this experiment. The six faces were used to create a set of face pairs, with each face paired with each other face once, making 15 pairs for each angle of rotation, 60 pairs for each type of response (same/different). Each face pair was repeated 3 times, making 360 trials in the experiment. On 'different' trials (when a normal and a Thatcherised face were paired together), the position of the normal face was counterbalanced between trials, such that it appeared on the left of the screen (and the Thatcherised face on the right) in half of the trials, and on the right of the screen (with the Thatcherised face on the left) in the other half. This ensured that there was the same number of 'same' and 'different' trials. All images were presented on a computer using Superlab.

5.1.1.3. Procedure

In each trial, participants saw two faces on the screen. These faces were either both normal, both Thatcherised, or one normal and one Thatcherised. Participants were asked to indicate whether they thought the two images were of the same type or not, by pressing one of two keys on a keyboard. In each trial, the face pair was presented until the participant made a response, followed by a 1 second blank interval

before the next trial. Participants were asked to respond as quickly and accurately as possible.

5.1.1.4. Design

The two within-participants independent variables were the orientation of the faces (upright and inverted) and whether the two faces were the same or different (both normal, both Thatcherised, or one normal, one Thatcherised). The dependent variables were the latency and accuracy of responses to the same/different decision for same face pairs.

5.1.2. Results

The results were calculated in the same way as in Experiment 4. The mean RTs and percentage of errors for same pairs (both normal or both Thatcherised) are shown in Figures 5.1 and 5.2 respectively. The mean percentage of RTs which were excluded for being more than 3 standard deviations from the mean was 2.3% for normal faces and 3.7% for Thatcherised face pairs.

A repeated measures ANOVA was conducted on the mean RTs to make correct 'same' decisions shown in Figure 5.1. There was a significant effect of orientation on performance ($F(1, 19) = 15.260; p < 0.05$) and a significant effect of face type ($F(1, 19) = 19.373; p < 0.05$). The interaction between orientation and face type just failed to reach significance, however ($F(1, 19) = 3.563, p = 0.074$). These results suggest that inversion had a detrimental effect on 'same' judgements to different-identity face pairs, but this effect was roughly the same for normal and Thatcherised faces.




Figure 5.1. Mean reaction times for same and different pairs in Experiment 6 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA conducted on the error data shown in Figure 5.2 shows a similar pattern, with a significant effect of orientation ($F(1, 19) = 34.912; p < 0.05$) and a significant effect of face type ($F(1, 19) = 118.196; p < 0.05$). The interaction between orientation and face type was also significant, however ($F(1, 19) = 19.588; p < 0.05$).




Figure 5.2. Percentage of errors for same and different pairs in Experiment 6 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

The mean RTs and percentage of errors for correct 'different' decisions are also shown in Figures 5.1 and 5.2 respectively. Repeated measures ANOVAs on these data showed significant effects of orientation on RTs ($F(1, 19) = 23.240; p < 0.05$), and errors ($F(1, 19) = 121.698; p < 0.05$).

We note from these analyses that the interaction between face type and orientation was highly significant in the error data ($p < 0.001$), but just failed to reach significance in the RT data. We therefore decided to perform a combined analysis of the RT and error data to investigate the overall effects of face type and orientation on face matching performance (Townsend & Ashby, 1983). The sum of the mean RTs divided by the proportion of correct responses for each condition for 'same' decisions is shown in Figure 5.3. A repeated measures ANOVA on this data showed significant effects of orientation ($F(1, 19) = 30.496; p < 0.05$), face type ($F(1, 19) = 56.184; p <$

0.05) and a significant interaction between face type and orientation ($F(1, 19) = 24.034; p < 0.05$).




Figure 5.3. Mean reaction time divided by proportion correct for same pairs in Experiment 6 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

The combined RT and error data for 'different' decisions is not shown in Figure 5.3 in order to show clearly the main interaction of interest here, but this data was also combined, and a repeated measures ANOVA revealed a significant effect of inversion for these responses ($F(1, 19) = 38.794; p < 0.05$).

Comparing Experiments 5 and 6

A between-experiments comparison with the upright and inverted data from Experiment 5 was conducted to test the hypothesis that the findings reported in Experiments 4 and 5 might be due to the intermediate angles of rotation considered in these experiments, and that this may modify the effects of inversion. Separate

analyses were again conducted for 'same' and 'different' pairs. The mean reaction times for 'same' decisions between the two experiments are shown in Figure 5.4.

A mixed ANOVA was conducted on the data from these two experiments. The within-participant independent variables were orientation (upright and inverted), face type (normal and Thatcherised), and the between-participants variable was whether the faces were rotated (Experiment 5) or inverted (Experiment 6).




Figure 5.4. Mean reaction time for 'same' decisions in Experiment 6 as a function of angle of rotation and face type for rotated (Experiment 5) and inverted faces (Experiment 6).

The analysis showed significant effects of face type ($F(1, 38) = 33.703$; $p < 0.05$) and orientation ($F(1, 38) = 32.153$; $p < 0.05$), and a significant interaction between face type and orientation ($F(1, 38) = 16.552$; $p < 0.05$). The main effect between the two experiments, however, was not significant ($F(1, 38) = .010$; $p > 0.05$), and neither was the interaction between the experiment and orientation ($F(1, 38) = .507$; $p > 0.05$), or any of the other interactions ($p > 0.05$ in all cases).

Indeed, a similar pattern of findings was observed in a combined analysis of the RT and error data from these two experiments, again showing a greater overall effect of inversion Thatcherised than normal face pairs; the effects of face type ($F(1, 38) = 99.942; p < 0.05$), orientation ($F(1, 38) = 55.198; p < 0.05$) and the interaction between face type and orientation ($F(1, 38) = 44.427; p < 0.05$) were significant, while none of the variables were significant ($p > 0.05$ in all cases).

Thus, as can be seen in Figure 5.4, the effect of inversion was the same for rotated and inverted faces in this same-matching task, while the significant interaction between face type and orientation suggests that for different-identity face pairs, the overall effect of inversion is greater for Thatcherised than for normal face pairs. This interaction can be seen in Figure 5.5 below.




Figure 5.5. Mean reaction time for 'same' decisions as a function of angle of rotation, collapsed across rotated (Experiment 5) and inverted faces (Experiment 6).




Figure 5.6. Mean reaction time for 'different' decisions as a function of angle of rotation for rotated (Experiment 5) and inverted faces (Experiment 6).

A between-experiments comparison with the upright and inverted data from Experiment 5 was also conducted for 'different' pairs. A mixed ANOVA on the data shown in Figure 5.6 revealed a significant effect of orientation ($F(1, 38) = 42.872$; $p < 0.05$), but neither the effect of experiment nor the interaction between experiment and orientation was not significant ($F(1, 38) = .230$; $p > 0.05$ in both cases), and this was also the case in the combined analysis of the RT and error data from these two experiments; the effect of rotation was significant ($F(1, 38) = 34.313$; $p < 0.05$), but the effect of experiment ($F(1, 38) = .929$; $p > 0.05$) and the interaction between experiment and rotation ($F(1, 38) = .885$; $p > 0.05$) were both non-significant. Thus the effect of inversion on RT's to make correct 'different' decisions was the same for both rotated and inverted face pairs.

5.1.3. Discussion

The main aim of this experiment was to examine the possibility that the differences in the overall effect of inversion between our findings from Experiment 5

(reported in Chapter 4) and those of Boutsen & Humphreys (2003) were due to the intermediate angles of rotation considered in our studies. We suggested that seeing the face pairs at a number of angles of orientation between upright and inverted may have encouraged a different processing strategy, and produced different effects of inversion, particularly for Thatcherised faces, to those observed in previous research. We hypothesised that if this was the case, we would expect to find effects of inversion similar to those of Boutsen & Humphreys (2003) when only upright and inverted faces were seen in a matching task; a detrimental effect of inversion for normal but not for Thatcherised face pairs. If comparable effects of inversion are observed in these two experiments (5 and 6), however, this would suggest that the differences in the effect of inversion between these studies and those of Boutsen & Humphreys (2003) could not be accounted for simply by the additional angles of rotation considered in our studies.

Our findings clearly show comparable effects of inversion for both rotated and inverted faces on RTs to make 'same' decisions in these studies. The results of Experiment 6 are very similar to those of Experiment 5; main effects of orientation and face type were again found, but the interaction between orientation and face type just failed to reach significance in this experiment. This interaction was significant in a combined analysis of the RT and error data, however, and also in a between-experiments comparison of the upright and inverted faces in these two studies, suggesting that for different-identity face pairs, the overall effect of inversion is greater for Thatcherised than for normal face pairs, and this is the case for both rotated and inverted faces. The data for 'different' decisions was also very similar to that found in Experiment 5 with rotated faces, with a between-experiments

comparison of the upright and inverted faces in these two studies showing that the effect of inversion on RTs to make correct 'different' decisions was the same for both rotated and inverted face pairs.

With regard to the different accounts of the face inversion effect, the findings of the between-experiments comparison provide further support for our hypothesis that Thatcherisation disproportionately disrupts the matching of face pairs when the task encourages configural processing. These findings can also be accounted for by a holistic processing account, however, where holistic refers to the coherence of the overall Gestalt of the face. Carbon & Leder (2005) found faster recognition times for inverted normal faces than inverted Thatcherised faces. Thus, when matching cannot be based on image-based featural information, participants may revert to a holistic processing strategy for inverted faces, which is able to account for the greater effect of inversion for Thatcherised than for normal face pairs.

5.2. Experiment 7

A further question suggested by this face matching paradigm is whether the findings reported with same-person pairs, where the two images were identical on 'same' trials, are also observed when different images of the same person are shown. That is, when the task is again to decide whether two faces of the same person are the same (both normal or both Thatcherised) or different (one normal and one Thatcherised), whether the effect of rotation is the same when two different images of the same person are to be compared. This allows us to look at the effects of rotation on a task which may encourage a more purely configural processing strategy than other studies using same-person pairs.

Young et al (1985) have considered the matching of different images of the same person compared with matches made from comparisons of the same photograph. Participants saw one face as a whole, intact version of that person, while the other face in the pair had either the internal or external features removed. When the two images were different, the internal or external features were derived from a different picture of that person. When matching could be image-based (same images), reaction times were faster, and there were no differences between familiar and unfamiliar faces for either type of feature, but the matching of whole faces to external features was faster than that to internal features of the same face. When the two images were different, however, Young et al (1985) found that matching based on internal features was faster for familiar than for unfamiliar faces, but there was no difference between familiar and unfamiliar faces for matches based on external features. Thus, while we do not look at differences between familiar and unfamiliar faces in this study, and no direct comparison was made between these experiments for unfamiliar faces (which we are interested in here), Young et al (1985) have nevertheless shown differential effects of same-image versus different-image face matching.

A number of other studies have also compared the processing of different images of the same person with that of same-image (or pictorial) processing, using a range of tasks. Schweinberger, Pickering, Jentsch, Burton & Kaufmann (2002) looked at the immediate repetition priming of faces using same and different images. Participants saw a briefly presented prime face, followed by a short delay (1300ms) and then a target face. They were asked to decide whether the second face in each trial was familiar (famous) or unfamiliar. Targets were primed either by the face of a

different person (i.e. unprimed), or by the same or a different picture of that person. Reaction times revealed a significant effect of priming, with both same and different pictures responded to faster than unprimed faces. Critically, responses to faces primed by the same image were also faster than those primed by a different image, suggesting a greater priming effect of having seen the same image of that person than a different image. Schweinberger et al (2002) also looked at the activity of the N250 component ('N250r'), which is thought to be an ERP modulation evoked by the immediate repetition of faces, on this task, and found some evidence of image specificity; the amplitude of this component was reduced for repetitions of different images compared to those of the same image.

In addition, Eger, Schweinberger, Dolan & Henson (2005) found a similar pattern of results to Schweinberger et al (2002) for participants making gender decision judgements to faces; same picture targets showed a stronger priming effect than different picture targets of that person. Schweinberger, Pickering, Burton & Kauffman (2002), however, found evidence to suggest that when repetition priming is considered over longer prime-test intervals (in this case 15 minutes), reaction times and ERP responses are both independent of whether the same or a different image of a person is seen. This appears to suggest that when the interval between prime and test is increased, priming no longer reflects changes in processing at the face recognition unit, or perceptual, level.

Thus it would appear that showing two different images of the person may encourage a different processing strategy, and produce different behavioural results in terms of RT's, at least for immediate repetitions or short prime-test intervals. In the

context of the face matching studies employed here, we suggest that showing two different images of the same person may encourage a more purely configural encoding strategy, whereas our previous experiment with same-person pairs was designed to encourage an image or feature-based strategy, but participants were still able to make use of configural information. Here, participants should not be able to rely on the simple matching of image-based features, but configural information should remain. If this is indeed the case, we might expect to observe similar findings to those of Experiment 5, in which a configural processing strategy was encouraged by using different-identity face pairs. We reported a greater effect of rotation for Thatcherised than for normal face pairs in that experiment, and argued that this indicated that rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces when the task requires this type of information.

5.2.1. Method

5.2.1.1. Participants

Eighteen undergraduates from Cardiff University with normal or corrected-to-normal vision received course credits for their participation in this experiment.

5.2.1.2. Materials

Six new male faces obtained from the Psychological Image Collection at Stirling (PICS; <http://pics.psych.stir.ac.uk>) were used in this experiment. All the faces were full frontal, colour images. For each person, three different images of that person were obtained. Each image showed a full frontal view of the face, but was a different picture of that person, such that image-based matching could not be used to complete

the task. As in the previous experiments, each face was of approximately similar size, and each image was resized to measure 283 pixels wide x 354 pixels in height on the screen. The images were centred vertically on the screen, and were 500 pixels apart horizontally. The eyes and mouth were selected and inverted separately for each face. The normal and Thatcherised upright faces were then rotated through the remaining four angles (45°, 90°, 135° and 180°).

These faces were used to create a set of face pairs. Each face was paired with itself and with the other two faces depicting that person, making 360 'same' trials (72 for each of the five angles of rotation), and a further 360 'different' (one normal, one Thatcherised) trials (72 for each angle of rotation). On 'different' trials (when a normal and a Thatcherised face were paired together), the position of the normal face was counterbalanced between trials, such that it appeared on the left of the screen (and the Thatcherised face on the right) in half of the trials, and on the right of the screen (with the Thatcherised face on the left) in the other half. All images were presented on a computer with a screen resolution of 800 x 600 pixels (horizontal x vertical) using Superlab, and subtended a visual angle of approximately 5°. Examples of the stimuli used in this experiment are shown in Figure 5.7.

5.2.1.3. Procedure

In each trial, participants saw two faces on the screen. These faces were either both normal, both Thatcherised, or one normal and one Thatcherised. Participants were asked to indicate whether they thought the two images were of the same type (both normal or both Thatcherised) or different (one normal, one Thatcherised) by pressing one of two keys on a keyboard. In each trial, the face pair was presented until

the participant made a response, followed by a 1 second blank interval before the next trial. Participants were asked to respond as quickly and accurately as possible.



Figure 5.8. Examples of the stimuli used in Experiment 7. Each face pair depicts two different images of the same person. The two left hand pairs (top and bottom) are of the same type, while the two right hand pairs (top and bottom) show 'different' type faces.

5.2.1.4. Design

The three within-participants independent variables were the orientation of the faces (upright, 45°, 90°, 135° and 180°), the image of the person (same or different), and whether the two faces were of the same type or different (both normal, both Thatcherised, or one normal, one Thatcherised). The dependent variables were the accuracy and latency of responses to the same/different decision for same face pairs.

5.2.2. Results

The results were calculated in the same way as in Experiment 4. The mean RTs and percentage of errors for same pairs (both normal or both Thatcherised) are shown in Figures 5.8 and 5.9 respectively. The mean percentage of RTs which were

excluded for being more than 3 standard deviations from the mean was 3.06% for same image and 1.94% for different image pairs.

A repeated measures ANOVA was conducted on the mean RTs to make correct 'same' decisions shown in Figure 5.8. There were significant main effects of orientation ($F(4, 68) = 28.418; p < 0.05$), face type ($F(1, 17) = 31.709; p < 0.05$) and image ($F(1, 17) = 56.476; p < 0.05$) on matching performance. The interactions between image and rotation ($F(4, 68) = 5.765; p < 0.05$) and face type and rotation ($F(4, 68) = 4.513; p < 0.05$) were also significant, but that between image and face type was not significant ($F(4, 68) = 1.996; p > 0.05$). The three way interaction between these variables was significant ($F(4, 68) = 3.389; p < 0.05$); for Thatcherised faces, the effect of rotation was greater for different-image than for same-image pairs, while for normal faces, different-image pairs yielded longer RTs than same-image pairs, but the effect of rotation was roughly the same for both types of face pairs.

These results show that whether the image of a person is the same or different (i.e. whether image-based matching is possible or not) can impact upon matching performance in this task. The effect of rotation was greater for different-image pairs than for same-image pairs, but whether the image of the person was the same or different did not differentially affect normal and Thatcherised faces. As has been found previously, the effect of rotation was much greater for Thatcherised than for normal face pairs.




Figure 5.8. Mean reaction times for same pairs in Experiment 7 as a function of angle of rotation, image and face type. Error bars show ± 1 standard error.

A series of paired-sample t tests looking at the intermediate angles of rotation found that only the 'different image, normal' condition showed a significant difference between upright and 45° rotated faces ($t(17) = 2.478$; $p < 0.05$). Significant differences were observed between 45° and 90° and 90° and 135° rotated faces in all conditions ($p < 0.05$ in all cases). The difference between 135° rotated and inverted faces was only significant in the 'different image, Thatcher' condition ($t(17) = 4.381$; $p < 0.05$).

A repeated measures ANOVA was conducted on the error data for 'same' decisions shown in Figure 5.9. Consistent with the RT data, there were significant main effects of orientation ($F(4, 68) = 6.436$; $p < 0.05$), image ($F(1, 17) = 33.609$; $p < 0.05$) and face type ($F(1, 17) = 25.722$; $p < 0.05$). The interactions between image

and rotation ($F(4, 68) = 6.588; p < 0.05$) and image and face type ($F(1, 17) = 7.948; p < 0.05$) were also significant, but that between face type and rotation was not significant in this analysis ($F(4, 68) = 2.164; p > 0.05$). The three way interaction between these variables was again significant ($F(4, 68) = 2.557; p < 0.05$); for same-image pairs there was no effect of rotation, and this was the case for normal and Thatcherised faces, while for different-image pairs, the effect of rotation was greater for Thatcherised than for normal faces.




Figure 5.9. Percentage of errors for same pairs in Experiment 7 as a function of angle of rotation, image and face type. Error bars show ± 1 standard error.

The mean RTs and percentage of errors for correct 'different' decisions are shown in Figures 5.10 and 5.11 respectively.




Figure 5.10. Mean reaction times for different pairs in Experiment 7 as a function of angle of rotation and image. Error bars show ± 1 standard error.

A repeated measures ANOVA was conducted on the mean RTs to make correct 'different' decisions shown in Figure 5.10. There was a significant effect of rotation on performance ($F(4, 68) = 35.022; p < 0.05$). The main effect of image, however, was not significant ($F(1, 17) = 1.152; p > 0.05$), as was the interaction between image and rotation ($F(4, 68) = 0.991; p > 0.05$). Thus, in contrast with the findings for 'same' decisions, while there was an effect of rotation on matching performance, the effect was the same for both same- and different-image pairs.

A series of paired-sample *t*-tests looking at the intermediate angles of rotation found that for 'different' decisions, the difference between upright and 45° rotated faces was significant for same image pairs ($t(17) = 3.288; p < 0.05$) but not when two different images of that person were shown ($t(17) = 1.898; p > 0.05$). Significant

differences were observed between 45° and 90° and 90° and 135° rotated faces in both conditions ($p < 0.05$ in all cases), but there was no significant difference between 135° rotated and inverted faces in either the same-image ($t(17) = 1.274; p > 0.05$) or different image ($t(17) = .741; p > 0.05$) conditions.

A repeated measures ANOVA was conducted on the error data for 'different' decisions shown in Figure 5.11. There were significant main effects of orientation ($F(4, 68) = 13.559; p < 0.05$) and image ($F(1, 17) = 10.292; p < 0.05$) but the interaction between image and rotation was not significant ($F(4, 68) = .290; p > 0.05$). Thus, although there was an overall effect of image, the effect of rotation was the same for same and different-image pairs on this task.




Figure 5.11. Percentage of errors for different pairs in Experiment 7 as a function of angle of rotation and image. Error bars show ± 1 standard error.

A regression analysis was also conducted on the data from this experiment, for each condition (normal, Thatcherised, different) and for each type of image (same or different), with linear and cubic functions fitted to the data. Comparisons of the ability of the linear and cubic functions to account for the data showed that for different-image normal face pairs, a cubic function provided a better fit to the data than a linear function ($F(2, 1) = 1093$), but this was not the case in any of the other conditions. Thus, the overall effect of rotation in these experiments has been one of a roughly linear decline in performance as the face is rotated away from the upright, although there is some evidence of non-linearities.

5.2.3. Discussion

The aim of this experiment was to look at the effect of rotation on a face matching task in which the two faces were either the same or a different-image of the person shown. We suggested that configural information should be required to complete this task, as matches can not be made solely by matching the features of the two images. If this was the case then we may expect to find similar effects of rotation to those reported in that experiment. The results for 'same' decisions were consistent with these predictions, with a greater effect of rotation for different-image pairs than for same-image pairs.

Indeed, we note that this effect is also consistent with the findings reported in the comparison of Experiments 4 and 5. This analysis showed a greater effect of rotation for different-identity than same-identity pairs. Thus, when the matching task cannot be completed by comparing image-based features, there is a greater effect of rotation than when this information is available and such a strategy can be employed.

Further, the effect of rotation was greater for different-image than for same-image pairs for Thatcherised faces, while for normal faces, different-image pairs yielded longer RTs than same-image pairs, but the effect of rotation was roughly the same for both types of face pairs. Thus the effect of rotation appears to be independent of face type only when the face pairs contain identical images of the same person, that is, when featural information can be used to complete the match. The data from Experiments 5 and 6 with different-identity pairs are also consistent with this suggestion, showing that Thatcherised faces are affected more by rotation than normal faces. These findings are therefore consistent with the view that rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces.

For 'different' decisions, our findings with regard to the overall effects of inversion are largely consistent with those of Experiment 4 and previous research when the task is to decide that a normal and a Thatcherised face are different (Boutsen & Humphreys, 2003; Lewis & Johnston, 1997). Same and different-image pairs appear to be affected differently by 45° rotations of significant effects of rotation in this experiment, however, with an effect for same-image but not different-image pairs. Indeed, this is consistent with the findings from Experiments 4 and 5. When same-identity pairs were to be matched in Experiment 4, there was a significant difference between upright and 45° rotated faces, but when different-identity pairs were matched in Experiment 5, this effect was not significant. Thus, when one is required to detect that one face in a *non*-identical pair is Thatcherised, rotation may have an earlier

effect on the matching of same-image pairs, but ultimately the effect of rotation is the same for both same- and different-image pairs.

5.3. Experiment 8

A final question we explore in this chapter can be seen by simply incorporating Experiments 4 and 5 and turning the questions they address around. Experiment 5 looks at normal and Thatcherised face pairs of different identities and asks participants to judge whether the type of image is the same for each face. Here we consider the effect of rotation on the reverse question, whether the *identity* of the two faces is the same in same (both normal or both Thatcherised) and different (one normal, one Thatcherised) face pairs. To address this, we incorporate the same identity trials from Experiment 4 with the different identity trials from Experiment 5, such that participants will now ignore the type of image shown (normal or Thatcherised) and instead judge whether the identity of the two faces is the same or different.

Judgements of the identity of face pairs differing in image type has been considered briefly in previous research by Bartlett & Searcy (1993). In their study, one face was always a normal, smiling view of a person, and the other face was a Thatcherised, a configurally distorted (in which spatial relations between parts are changed, such as the distance between the eyes), a neutral-expression or a grotesque-expression face, either of the same person or of a different person. Each pair was seen both upright and inverted. The task, designed to encourage featural processing, was to judge whether each pair depicted the same person or two different people. Bartlett & Searcy (1993) hypothesised that if same/different face judgements are based on

featural processing, inversion should not reduce the accuracy or latency of correct 'same' judgements to same face pairs. There was no reliable effect of inversion on accuracy for Thatcherised or spatially distorted same face pairs, and no consistent effects of inversion on response latencies, although there was a significant advantage for upright faces for correct different-face judgements. Bartlett & Searcy (1993) suggest that their data support the idea that the Thatcher illusion results from a disruption of configural processing in upright faces.

White (2002) also looked at the effect of configural changes on identity and expression matching tasks, considering faces in which either one or both eyes in a face were moved up into the forehead region. These changes are thought to primarily affect categorical and configural changes of the spatial relations of face parts. Categorical changes refer to changes to the general directional layout of the face, such as one eye being to the side of the other. Configural changes refer to the violations of the distances between these features. Thus, moving one eye up primarily alters categorical assumptions, whereas moving both eyes up (but maintaining the level and distance between them) primarily affects configural information.

Participants saw a pair of faces on the screen in each trial, presented for only 2 seconds, and simply had to decide whether the two faces were of the same or different women, whilst ignoring any difference in the eyes or in facial expressions between the two images. Incompatible responses (pairs showing the same women with different expressions, or different women with the same expressions) were slower than compatible responses (same women with same expressions, or different women with different expressions). Identity matches were found to be slower in pairs where both

eyes were moved than one-eye-moved pairs, while the opposite pattern was true of expression matches. These findings appear to suggest that changes to configural information are important for face identity matching, and that this appears to affect the perceptual encoding of identities, rather than retention or retrieval processes.

Here we adopt a similar task to Bartlett & Searcy (1993), although we aim to look only at normal and Thatcherised faces. Also, in line with our previous experiments, participants are not given a time limit in which to make a same/different response to each face pair, as they were in Bartlett & Searcy's (1993) experiments. If this task encourages featural processing, as Bartlett & Searcy (1993) suggest, we would expect to find no effect of rotation or inversion for 'same' judgements, but an inversion effect for 'different' identity decisions. Although a perceptual matching task is still employed, this experiment addresses more closely the question of facial recognition and the impact of identity processing on the Thatcher illusion and the role of configural information in recognition, an area which has not attracted much attention, despite being arguably one of the most important questions for face processing research.

5.3.1. Method

5.3.1.1. Participants

Eighteen undergraduates from Cardiff University with normal or corrected-to-normal vision received course credits for their participation in this experiment.

5.3.1.2. Materials

The same stimuli detailed in Experiment 4 were used in this experiment. The six faces were used to create a set of face pairs, with each face paired with itself once. In addition, 6 further 'different identity' face pairs were chosen at random from the 15 possible face pairings. This was to ensure that there was the same number of 'same' and 'different' identity trials (12) for each angle of rotation. Each face pair was repeated 3 times, making 36 trials for each angle of rotation, 720 trials in total. There were 360 trials for each type of response (same identity vs. different identity), 180 of which consisted of faces of the same type (both normal or both Thatcherised) and 180 of a different type (one normal, one Thatcherised). On 'different' trials (when a normal and a Thatcherised face were paired together), the position of the normal face was counterbalanced between trials, such that it appeared on the left of the screen (and the Thatcherised face on the right) in half of the trials, and on the right of the screen (with the Thatcherised face on the left) in the other half, such that there were the same number of 'same' and 'different' trials. All images were presented on a computer using Superlab.

5.3.1.3. Procedure

In each trial, participants saw two faces on the screen. These faces were either both normal, both Thatcherised, or one normal and one Thatcherised. Participants were asked to indicate whether they thought the two images were of the same *identity* or not, by pressing one of two keys on a keyboard. In each trial, the face pair was presented until the participant made a response, followed by a 1 second blank interval before the next trial. Participants were asked to respond as quickly and accurately as possible.

5.3.1.4. Design

The three within-participants independent variables were the orientation of the faces (upright, 45°, 90°, 135° and 180°), the identity of the two faces (same or different) and whether the two faces were the same or different (both normal, both Thatcherised, or one normal, one Thatcherised). The dependent variables were the latency and accuracy of responses to the same/different decision for same face pairs.

5.3.2. Results

The results were calculated in the same way as in Experiment 4. The mean percentage of RTs which were excluded for being more than 3 standard deviations from the mean was 1.63% for same identity pairs faces and 1.79% for different identity pairs. A repeated measures ANOVA was conducted on the three independent variables; identity (same or different), face type (normal, Thatcherised or different) and rotation (upright, 45°, 90°, 135°, inverted). There was a significant effect of rotation on performance ($F(4, 68) = 15.395; p < 0.05$), but no significant effect of face identity ($F(1, 17) = 1.697; p < 0.05$) or of face type ($F(2, 34) = .966; p > 0.05$). There was a significant interaction between face type and face identity ($F(2, 34) = 6.911; p < 0.05$). None of the other interactions were significant, however.




Figure 5.12. Mean reaction times in Experiment 8 as a function of identity, face type and angle of rotation. Error bars show ± 1 standard error.

To investigate the interaction between face type and face identity, we conducted a series of ANOVAs, this time looking at ‘same’ and ‘different’ identity judgements separately. Responses and errors were also analysed separately for same type (both normal or both Thatcherised) and different type (one normal, one Thatcherised) pairs for each level of the ‘identity’ variable, making four separate ANOVAs in total. A repeated measures ANOVA was conducted on the mean RTs for ‘same identity, same type’ pairs shown in Figure 5.12. There were significant effects of rotation ($F(4, 68) = 5.101; p < 0.05$) and face type ($F(1, 17) = 7.042; p < 0.05$) on performance, and also a significant interaction between face type and rotation ($F(4, 68) = 2.642; p < 0.05$). This shows that the effect of rotation was greater for Thatcherised than for normal face pairs.

Examining the effects of rotation in more detail, for normal faces, a series of paired-sample t-tests showed a significant difference between upright and 45° rotated faces ($t(17) = 3.271; p < 0.05$), but the difference between 135° and inverted faces

was not significant ($t(17) = .693; p > 0.05$), as were the differences between all other intermediate angles of rotation. For Thatcherised faces, only the decrease in RTs between 135° and inverted faces was significant ($t(17) = 4.224; p < 0.05$; $p > 0.05$ in all other cases).

A separate ANOVA was conducted to look at the effect of rotation on 'same identity, different type' responses shown in Figure 5.12. There was no significant overall effect of rotation on performance ($F(4, 68) = 1.673; p > 0.05$), but the difference between upright and inverted face pairs was significant in a paired-samples t-test ($t(17) = 2.255; p < 0.05$), suggesting that there was an inversion effect on this task.




Figure 5.13. Mean reaction times for different identity pairs in Experiment 8 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA was also conducted on the mean RTs for 'different identity, same type' pairs shown in Figure 5.13. There was a significant

effect of rotation ($F(4, 68) = 7.569; p < 0.05$), but the effect of face type was not significant, ($F(1, 17) = 2.826; p > 0.05$), as was the interaction between face type and rotation ($F(1, 17) = .132; p > 0.05$). Thus, as can be seen in Figure 5.13, the effect of rotation was very similar for both types of faces with different-identity pairs. Indeed, a similar effect of rotation was also found for different identity decisions to faces differing in face type, as revealed in a repeated measures ANOVA on this data ($F(4, 68) = 8.301; p < 0.05$).




Figure 5.14. Percentage of errors for same identity pairs in Experiment 8 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A repeated measures ANOVA of the error data was conducted on the three independent variables; identity (same or different), face type (normal, Thatcherised or different) and rotation (upright, 45°, 90°, 135°, inverted). This data is shown in Figures 5.14 for 'same identity' and in 5.15 for 'different identity' pairs. There was a significant effect of rotation ($F(4, 68) = 2.636; p < 0.05$), but not of face type ($F(2,$

34) = .702; $p > 0.05$) or of identity ($F(1, 17) = .483$; $p > 0.05$), and none of the two or three-way interactions were significant ($p > 0.05$ in all cases).




Figure 5.15. Percentage of errors for different identity pairs in Experiment 8 as a function of angle of rotation and face type. Error bars show ± 1 standard error.

A regression analysis was also conducted on the data from this experiment, for each condition (normal, Thatcherised, different) and for each type of identity pair (same or different), with linear and cubic functions fitted to the data. Comparisons of the ability of the linear and cubic functions to account for the data showed that cubic functions did not provide a better fit to the data than linear functions in any of the six conditions. Thus, the overall effect of rotation in these experiments has been one of a roughly linear decline in performance as the face is rotated away from the upright.

5.3.3. Discussion

The aim of this experiment was to look at the effect of rotation on the matching of facial identity with same and different face type pairs. We suggested that if featural information is used to complete this task, one might expect to find

significant effects of rotation for normal and Thatcherised faces, but no difference between these types of face pairs. This prediction was supported in the overall analysis, with no effect of face type or interaction between face type and rotation found. When 'same' and 'different' identity judgements were analysed separately, however, there was an effect of face type on 'same' identity judgements for same type pairs. Further, the interaction was also significant, showing that the effect of rotation was greater for Thatcherised than for normal face pairs.

For 'different' identity judgements, however, there was no effect of face type, or interaction between face type and rotation. It can therefore be concluded that the effect of rotation was similar for both types of faces with different identity pairs. Thus it would appear that when judging the identity of two simultaneously presented faces, participants use a different strategy to when the type of stimuli is to be compared. That is, they use featural information to distinguish between different identity pairs, and this information is not differentially affected by face type. Configural information, on the other hand, is used to judge that the two faces are the same and this is affected by rotation and Thatcherisation.

5.4. General Discussion

In this chapter we have built on the findings from Chapter 4 by examining three further questions suggested by this face matching paradigm. First, we looked at the matching of different-identity face pairs in the upright and inverted orientations only, and compared the findings to those reported with rotated different-identity pairs. We found no evidence to suggest that the intermediate angles of rotation had affected

the overall effects of inversion in Experiment 5, or were responsible for the differences observed between our studies and those of Boutsen & Humphreys (2003).

Second we examined the matching of different images of the same person, and the effect of rotation on this task compared with that of same-image matching. This encouraged a more 'pure' configural processing strategy than has previously been possible with same-person pairs. We found that for 'same' decisions, there was a greater effect of rotation for different-image than for same-image pairs. Further, the effect of rotation was greater for different-image than for same-image pairs for Thatcherised faces, while for normal faces the effect of rotation was roughly the same for both types of face pairs. We suggest that, taken together, these findings indicate that the effect of rotation is independent of face type only when the face pairs contain identical images of the same person. The data from Experiments 5 and 6 with different-identity pairs are also consistent with this suggestion, showing that Thatcherised faces are affected more by rotation than normal faces. These findings are therefore consistent with the view that rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces.

Finally we looked at the effect of rotation on the matching of facial identity with same and different face type pairs. We found that when 'same' and 'different' identity judgements were analysed separately, there was an effect of face type on 'same' identity judgements for same type pairs. Further, the interaction was also significant, showing that the effect of rotation was greater for Thatcherised than for normal face pairs. For 'different' identity judgements, however, there was no effect of face type, or interaction between face type and rotation; the effect of rotation was

similar for both types of faces with different identity pairs. We have suggested that these findings indicate that judging the identity of two simultaneously presented faces requires a different processing strategy, with featural information used to distinguish between different identity pairs, and this information is not differentially affected by face type. Configural information, on the other hand, is used to judge that the two faces are the same and this is affected by rotation and Thatcherisation.

With regard to the predictions of the holistic and dual-mode accounts of face processing outlined in Chapter 4, we suggest that the holistic processing theory is able to account for some, but not all, of the findings reported here. A strict interpretation of this account would suggest that Thatcherisation should have no effect on face matching, as faces are encoded as wholes, rather than by an analysis of their composite parts. There should, therefore, be no difference between normal and Thatcherised faces, whether same or different-person pairs are seen. These predictions are inconsistent with the findings of the experiments reported in this chapter.

If 'holistic' is defined as whether or not the Gestalt of the overall image is coherent, however, the holistic view is able to account for some of our findings. In all of the experiments considered here, we found faster recognition times for inverted normal faces than inverted Thatcherised faces. Thus, when matching cannot be based on image-based featural information, participants may revert to a holistic processing strategy for inverted faces, which is able to account for the greater effect of inversion for Thatcherised than for normal face pairs. We note, however, that previous tests of this version of the holistic account have not compared upright and inverted faces

directly, and as such, this theory is currently unable to offer strict predictions about the effects of inversion and rotation on face matching tasks.

In this chapter we have presented further evidence against the distinction between local and global configural information, suggesting that this does not provide a useful framework for understanding the effects of rotation on face matching tasks. Furthermore, we have investigated the roles of these types of configural information in a variety of tasks, and shown that when featural information can be used to complete the match (i.e. when the face pairs contain identical images of the same person), the effect of rotation is independent of face type. When different-identity or different-image pairs are shown and configural information is required, however, Thatcherised faces are affected more by rotation than normal faces; rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces.

We have also considered the question of whether the relationship between rotation and face recognition is of a linear or non-linear nature. While we are unable to conclude that there is definitely a linear relationship between rotation and face matching performance in these tasks, the majority of the data suggested that a linear function provided the best account of the data. These findings are in broad agreement with the prediction of the mental rotation hypothesis that recognition performance is dependent on the amount of normalisation required. We found little evidence of an improvement in performance between 135° rotated and inverted faces in these experiments; this difference was only significant in the 'different image, Thatcher' condition of Experiment 7. Thus, in terms of the means by which face rotation tasks

are accomplished, our overall pattern of findings appears to suggest that performance is dependent on the amount of normalisation required.

Chapter 6: Conclusions and Directions for Future Research

While many studies have looked at the effects of inversion in face perception over the last thirty or forty years, it is only recently that researchers have considered the intermediate angles of orientation in an attempt to develop a broader understanding of this face inversion effect. Specifically, these studies have looked at the way in which different types of information are affected by rotations of the face in the picture plane, and have used a variety of methodologies, including perceptual matching tasks, face familiarity and gender classification judgements and the Thatcher illusion to address this question. Moreover, whether the relationship between the degree of rotation of the face from upright and reaction times or accuracy data is of a linear or non-linear nature has been taken in many studies as an indicator of whether there is a switch in processing between upright and inverted faces, and if so, where in the rotation process this might occur. As highlighted in Chapter 1, however, researchers have presented evidence for both types of relationship, and research which has used the Thatcher illusion to look at this question has also produced largely inconsistent findings with regard to the face processing-rotation relationship.

6.1. Task Demands and Methodological Issues

Before we examine the main conclusions of this thesis and consider the implications of our findings for configural and holistic theories of face processing, it is important to look at the main methodologies employed in this thesis. In this section we provide an overview of the studies which have employed these methods, and where appropriate, consider some of the flaws of these methodologies, both in terms

of their underlying assumptions as paradigms, and in the way in which they have been implemented in previous research.

6.1.1. The Part-Whole Paradigm

A number of studies have made use of the part-whole paradigm, developed by Tanaka & Farah (1993), on which the experiments in Chapter 3 are based. To recap, Tanaka & Farah (1993) showed participants a series of faces with an associated name. At test, participants were required to identify these faces from their isolated features or from their whole face, in a two-alternative forced-choice task. For the facial features, two different facial parts were shown (e.g. two different mouths), while for the whole faces, the distractor differed from the target by just one feature. Participants were asked to decide as quickly but as accurately as possible which feature belonged to the named person, or which face showed that person. Features were recognised with greater accuracy when seen in the context of the whole face than when seen in isolation, even though the two faces differed by only one feature. However, for inverted and scrambled faces, and schematic houses, the advantage for wholes over parts disappeared; facial features (and house features such as doors and windows) were recognised with the same accuracy, whether they were seen in isolation or in the whole (face) context. Tanaka & Farah (1993) argued that participants formed explicit representations of the face parts for inverted and scrambled faces (because configural information was not accessible), but not when learning upright faces.

Donnelly & Davidoff (1999) replicated Tanaka & Farah's (1993) study, comparing the recognition of parts and wholes of faces and houses. The only difference between the two studies was that Donnelly & Davidoff (1999) used two

further learning conditions, with participants given either 30, 60 or 90 trials. Their findings with faces replicated those of previous studies; there was an advantage for whole faces over isolated features. However, this pattern of recognition performance was also true for houses, a finding which is in direct contrast to that of Tanaka & Farah (1993). Donnelly & Davidoff (1999) suggested that the advantage for whole objects could be due to a potential flaw of the part-whole paradigm. That is, whole faces at test were comprised of a unique combination of features, and only that combination of features could be that person. Isolated features, however, were seen in two different targets, so a nose may belong to, for example, David, but would also be the correct nose for a different individual (e.g. Michael). Thus, the whole-face advantage may reflect the recognition of a unique combination of features, rather than the importance of the information between those features itself.

Donnelly & Davidoff (1999) tested this prediction in their second experiment by looking at scrambled faces and houses. Tanaka & Farah (1993) have previously found no whole over part advantage for scrambled faces, but both the learning and test stimuli were scrambled, so this could be due to scrambled faces being difficult to learn. Here the stimuli were normal in the learning phase. Donnelly & Davidoff (1999) argued that the advantage for whole stimuli should persist if it is indeed based on the recognition of a unique combination of features, rather than a genuine advantage for the additional spatial information they contain. If the advantage for whole objects *is* due to the additional spatial information, however, scrambling these stimuli should disrupt this information and abolish the whole over part advantage.

The results showed that the recognition of scrambled faces removed the whole advantage for faces, but not for houses. It would therefore appear that holistic encoding is used in the recognition of faces, but that this strategy might be unique to faces on this type of task. Interestingly, Donnelly & Davidoff (1999) also found support for this conclusion from a two-alternative forced-choice version of this task.

These findings appear to support the conclusions drawn by Tanaka & Farah (1993) that face perception is holistic in nature, as evidenced by the whole-face advantage in part and whole recognition. Moreover, the finding that scrambling faces at test removed the advantage for whole faces addresses, and appears to counter, a possible methodological flaw of the part-whole paradigm, that the recognition of whole faces may be based on the recognition of a unique combination of features (see also Leder & Carbon, 2005). This appears to suggest that the whole-face advantage is a genuine effect of the additional spatial information which is seen, and these findings are interpreted as being consistent with the holistic view of face processing.

Leder & Carbon (2005) also used the part-whole paradigm to consider the possibility that this whole-face advantage would also occur when isolated features, as well as whole faces, were learned. Participants learned upright whole faces or isolated features, and were tested on their recognition of both, in both upright and inverted orientations. When whole faces were learned, a whole-face advantage was once again observed for upright faces; recognition of whole faces was consistently superior to that of isolated face parts. Again, this is consistent with the holistic processing hypothesis, suggesting that separating the face into parts is particularly difficult. When face parts were learned, however, recognition performance was better when

tested with face parts than with whole faces, suggesting that the context of the whole face interferes with the recognition of features when these features have been learned. Thus the whole-face advantage may be one-directional in nature (an advantage is only seen when whole faces have been learned previously), and holistic information may interfere with the processing of face parts. This was shown to be the case even when participants learned only one feature (the critical feature at test), suggesting that these findings are not due to the uncertainty as to which would be the critical features (Experiments 1 and 2).

It is interesting to note, however, that these effects were reduced when real faces were used as the stimuli, and participants learned only the critical isolated feature (Experiment 4). When full faces were learned, the evidence was consistent with a whole-face advantage for recognition, but when parts were learned, participants were only slightly better at recognising parts than whole faces; in previous experiments this advantage for face parts had been clear and significant. This appears to suggest that participants were able to ignore irrelevant information more easily with real faces. This conclusion was supported by a comparison of Experiments 2 (in which participants also only learned the critical isolated feature, with artificial faces) and 4; the advantage for part over whole recognition when parts had been learned was significantly smaller for real than for artificial faces. This suggests that participants were able to ignore irrelevant information from whole faces when real faces were shown, suggesting that there is less interference from additional information when this information is not identical over different trials, as in Experiments 1 and 2 and most other research which employs the part-whole paradigm (e.g. Tanaka & Farah, 1993).

Lewis & Glenister (2003) also noted the problem of the part-whole paradigm being based largely on the use of schematic images, which are of poor quality and lacking in high-level detail. In their study, Lewis & Glenister (2003) replicated Tanaka & Farah's (1993) original study, using greyscale faces created using face reconstruction software, to consider the effect of 90° rotations on performance in this task, to assess how configural encoding is affected by rotation. An advantage for whole faces over isolated face parts was found for upright and 90° rotated faces, but not inverted faces, consistent with previous findings. Tanaka & Farah (1993) found no inversion effect for isolated feature tests from upright faces; recognition performance was equivalent to that for feature tests for inverted faces. Lewis & Glenister (2003), however, found a significant effect of inversion for isolated feature tests, suggesting that configural encoding may occur for individual features as well as for the whole face, albeit only when high quality facial images are used (see also Palermo & Rhodes, 2002). This appears to suggest that when good quality facial images are seen, the inversion effect may not be simply due to the disruption of holistic processing.

A range of studies have also used the part-whole paradigm to look at face recognition in young children, arguing that it provides a direct means by which to test the encoding switch hypothesis. This is the notion that young children encode faces in a featural manner until roughly the age of 10, before processing begins to switch to a more holistic nature (Carey & Diamond, 1977). This hypothesis was formed on the basis of a study conducted with 6, 8, and 10 year old children. They were shown upright and inverted photographs of unfamiliar faces, and were then asked to identify the previously seen face from a completely new face. The results showed that 8 and 10 year olds recognised the 'old' faces better when they were seen upright than when

inverted, but for 6 year old children, inverted faces were recognised just as well as upright faces. Carey & Diamond (1977) concluded that young children process faces in a featural manner initially, before switching to a more holistic approach over time.

Tanaka, Kay, Grinnell, Stansfield & Szechter (1998) used the part-whole paradigm to provide a direct test of this encoding switch hypothesis. This view would predict that if young children do indeed encode faces featurally, the recognition of face parts should be the same, whether presented in isolation or in the whole face. If holistic processing is used, however, a whole-face advantage similar to that found with adults should be observed. The results supported the latter prediction; by 6 years of age, children recognised face parts more accurately in the whole face condition than in isolation.

A further methodological problem with the part-whole paradigm, however, is that all the faces are learned together, and as a consequence participants are required to remember the faces over a reasonably long period of time. Tanaka et al (1998) noted that long-term memory tasks may invoke, or require, different processing strategies to those of immediate memory tasks, for example by allowing the formation of a holistic representation over time, whereas immediate memory tasks may be based on more featural strategies. To test this, Tanaka et al (1998) altered the part-whole paradigm to test recognition immediately after the faces had been seen. Each face was seen for 5 seconds, after which children identified a face part or a whole face in the usual two-alternative forced-choice task. Children were again found to recognise face parts more accurately in the whole face condition than in isolation, and, further, the size of this whole-face advantage did not change across the four age groups (6, 7, 9

and 11 years). Thus it would appear that the original part-whole paradigm does not invoke a different processing strategy, and, moreover, younger and older children (and adults) do not rely on different face processing strategies.

The research reviewed here has raised a number of important issues about the part-whole paradigm, and appears to counter at least two methodological issues. Firstly, Donnelly & Davidoff (1999) have shown that the original whole face advantage reported by Tanaka & Farah (1993) was not simply due to the recognition of a unique combination of features; rather, this effect appears to reflect the importance of the additional spatial information whole faces contain. Second, research has also refuted the claim that the whole face advantage may be the product of the long-term memory task employed in the paradigm, and may reflect a strategy which is employed in long-term but not short-term memory tasks. Tanaka et al (1998) found no evidence to support this view. Our review has, however, noted the importance of image quality (of artificial faces; Lewis & Glenister, 2003), and the type of image seen (artificial vs. real; Leder & Carbon, 2005) for the whole face advantage. Nevertheless, we suggest that the part-whole paradigm provides an excellent tool for investigating the effects of rotation in face recognition memory.

6.1.2. Face Matching Tasks: Simultaneous or Sequential?

The experiments described in Chapters 4 and 5 of this thesis employed a number of variations on a face matching task. In these studies, participants were asked to decide whether two simultaneously presented faces were of the same type (both normal or both Thatcherised), or not (one normal, one Thatcherised). In Experiment 4, the two faces were always the same image of the same person, while in Experiments 5

and 6, the two faces always depicted different people. In a further set of experiments, we also considered the situation where the two images depicted the same person, but were different images of that person. This eliminates the possibility of the match being based on the image rather than the face itself. Finally, we also considered the matching of the identity of the two faces in same and different face pairs. That is, participants ignored the type of image (normal or Thatcherised) and instead judge whether the identity of the two faces was the same or different.

As we have seen in the research discussed so far, face matching studies have used both sequential and simultaneous paradigms to look at the effect of manipulations such as inversion, Thatcherisation, changes to spatial positioning and feature saliency on face processing (e.g. Boutsen & Humphreys, 2003; Schwaninger & Mast, 1999, 2005; Searcy & Bartlett, 1996; Valentine & Bruce, 1988; Yovel & Kanwisher, 2004). Whilst it is not possible to compare these studies directly, it is important to acknowledge that this distinction between perceptual and memory-based face matching may, in part at least, account for some of the discrepancies in the findings of these studies (see, for example, the differences between the findings of Searcy & Bartlett (1996) and Yovel & Kanwisher (2004), although other methodological differences may also be partly responsible for these discrepancies).

A number of studies have directly compared the use of simultaneous and sequential matching in face processing tasks. Leder (1999) considered why face recognition is so impaired when a face is transformed into a line drawing. One possibility is that it is particularly difficult to match a line drawing to a stored memory representation of that face. Indeed, Bruce, Hanna, Dench, Healey & Burton (1992)

argued that the lack of surface information in line drawings makes this task particularly difficult. The second possibility is that this loss of information makes it difficult to perceive the two faces as being the same. Leder (1999) also looked at the type of representation necessary for the processing of line drawings of faces. He argued that structural representations of an image are necessary for recognition to occur across differences in, for example, viewpoint and lighting direction. Following this assertion, Leder (1999) also compared the recognition of same and different-viewpoint faces in his study. This is a similar manipulation to that used in Experiment 7 reported here, although in order to look at the effects of Thatcherisation, we only considered whether the images of the person were the same or different. Leder (1999) predicted that if line drawings give weaker structural representations than real faces, then the matching of line drawings differing in viewpoint should be particularly difficult in this task.

Participants saw two blocks of trials, one consisting of sequentially presented face pairs, the other of simultaneously presented pairs. Within each block, half of the trials showed two faces from the same viewpoint (same or different identity), and half showed different viewpoints (same or different identity). The two images were either both line drawings, both photographs or one line drawing and one photograph, with the order of presentation counterbalanced in these last two conditions (the 'mode' condition). The task was simply to decide whether the two pictures in each trial were of the same person or of two different people.

The results showed that when one or both of the images was a line drawing, sequential matching was more difficult than simultaneous matching. That is, whether

a line drawing was matched to a photograph, a photograph was matched to a line drawing or two line drawings were matched, sequential matching made this task significantly more difficult than simultaneous matching. There was no difference of matching task type when photographs were matched, however, suggesting that the representation formed from a photograph enables participants to store that face in memory and recall it after a brief delay. In the case of line drawings, the face representation may only be strong enough to match simultaneously presented items, and is weak when the task requires a representation to be held in memory.

With regard to the effects of viewpoint, the matching of two images differing in viewpoint was more difficult than the matching of same viewpoint faces, and again this interacted with the mode condition. There was a smaller effect of a change in viewpoint when two photographs were shown, which, when removed, also removed the interaction. This effect of change in viewpoint for photographs was still significant, but overall it would appear that the matching of two images differing in viewpoint is particularly difficult for line drawings of faces. Leder (1999) suggests that in this case the lack of surface information in line drawings may make the creation of a three-dimensional representation difficult, and it is this information which is lacking with line drawings. Importantly, the interaction between viewpoint and the type of matching task was not significant in the analysis, however, indicating that the effects of viewpoint were the same for sequential and simultaneous matches. Leder's (1999) findings suggest, therefore, that photographs form sufficiently strong representations that they can be matched equally well regardless of whether the matching task is of a perceptual or memory-based nature. Furthermore, the matching

of two images differing in viewpoint was more difficult than the matching of same viewpoint faces, and this was particularly true of line drawings of faces.

Barton, Deepak & Malik (2003) have also compared the use of sequential and simultaneous matching techniques. These authors tested the proposal that differential effects of inversion for spatial and featural changes are due to inversion impairing the perception of spatial changes in less salient regions of the face. Barton et al (2003) used a change blindness task to address this question. In this methodology, the presentation of two brief images is interrupted by a very briefly presented blank screen. The term change blindness refers to the fact that large changes (or differences) can be introduced between the two images, and these are undetected when presented in this way. What *is* seen is dependent on the distribution of attention across the image, meaning that changes to more salient areas of the image are more likely to be detected than those in less salient regions (e.g. Rensink, O'Regan & Clark, 1997).

Barton et al (2003) looked at the discrimination of changes to eyes and mouths in upright and inverted faces, and changes were made to feature colour, by lightening the eye or mouth colour, and to the spatial location of the eyes or mouth, by reducing the interocular distance for eyes or elevating the mouth from its original position. These changes were made to two base faces, one male and one female, producing 32 target faces in total. Before the experiment began, participants were familiarised with the base faces, and the most extreme examples of each of the four manipulations. This meant that participants were partially cued, as they knew what type of changes to expect, if not where they would occur. Each test block consisted of 128 trials, 64 of the male and 64 of the female face, and there were equal numbers of upright and

inverted and 'same' and 'different' trials for each face. Participants saw two faces for 555ms each, separated by an interval (a white screen) of 120ms, and their task was to decide whether the two faces were the same or not.

Inversion had a small effect on the detection of changes in eye position, eye colour and mouth colour, but had a significant effect on the perception of changes in mouth position, to the point at which these changes were made almost invisible. This was the case even when participants were partially cued as to what these changes were. Barton et al (2003) suggest that these findings provide support for the notion that both spatial relational information and regional salience are important factors in accounting for the face inversion effect.

The research reviewed in this section, therefore, suggests that, at least in the type of matching experiments we consider in this thesis, the distinction between perceptual and memory-based matching may only be an important consideration if facial images are degraded in some way; sequential, memory-based matching was more difficult when line drawings were involved in the match (even when two line drawings were to be compared), while the matching of photographs of faces was the same in simultaneous and sequential matching tasks (Leder, 1999). Further, Barton et al (2003) showed very similar effects of inversion on changes to spatial relations and facial features in simultaneous and sequential matching tasks. Thus, differences in the type of matching task employed may also be responsible for differences between studies where one study uses degraded facial images, but the distinction between perceptual and memory-based matching does not appear to be important when comparable images are used. We have also shown that the matching of two faces

differing in viewpoint is more difficult than that of same-viewpoint pairs, and while this effect also appears to be more pronounced with degraded images, normal face matching is also affected by changes in viewing angle. Indeed, our findings from Experiment 7, in which we looked at the matching of same versus different images of a person, are consistent with this research; the effect of rotation was greater for different-image than same-image pairs, but whether the image of the person was the same or different did not differentially affect normal and Thatcherised faces. We consider the implications of these findings later in this chapter.

6.1.3. Visual Search

The visual search paradigm, in which participants search for a target stimulus among an array of distractors, has also been used to inform the debate between the two major theories of face processing and the face inversion effect, again looking at the *type* of information being processed which differentiates upright from inverted faces (e.g. Kuehn & Jolicoeur, 1994; Murray, 2004). Early research in the visual search field conducted by Treisman & Gelade (1980) suggested that objects that can be identified by a unique feature from other items tend to pop out immediately from an array. For such items, the size of the array does not affect the time required to detect it. Complex objects that can only be detected by a specific conjunction of features, however, must be searched for in a serial fashion by examining every possible item for the correct conjunction. Detection of these items within a scene will slow as the size of the array increases. A visual-search task, therefore, can either be performed serially or in parallel, and the method being employed is usually determined from the change in reaction time as the array size increases (i.e. the slope of the target-present reaction-time graph).

More recent evidence from visual-search tasks, however, suggests that the distinction between parallel and serial encoding may not be as simple as first believed. Duncan & Humphreys (1989), for example, demonstrated that, in the plots of reaction time against array size, the angle of the target-present slope is influenced by the nature of the distractors. Search becomes more difficult the more similar the distractors are to the target, and also the more dissimilar distractors are *to each other*. Wolfe (1994) suggests that some items offer preattentive information that will guide more focused attention, and therefore search should be interpreted as being efficient or nonefficient, with efficient searches leading to parallel processing being observed in the reaction times. In a later paper, Wolfe (1998b) suggested that it is necessary to consider all the data from the target-present and target-absent slopes in order to determine whether search is efficient or inefficient.

This research has guided thinking in the visual search paradigm, and performance on visual search tasks is usually interpreted using one, or both, of two measures. First, the increase in reaction time to detect an item per additional distractor can be used (Wolfe, 1998a). Where this increase is greater than 10 ms per item, search is deemed to be serial, whereas where it is less than this it is parallel, and item can be said to 'pop out' of a visual array. This is admittedly an arbitrary distinction, as Wolfe (1998b) acknowledges, but nevertheless it is often applied as a cut-off point for determining the nature of the search. Second, following Wolfe (1998b), a self-terminating search hypothesis predicts that a serial search will have a 2:1 ratio between target absent and target present slopes (Treisman & Gelade, 1980). A ratio greater than 2:1 can be used to indicate pop-out.

These methods of distinguishing parallel and serial search are useful in describing the efficiency of visual search, but as Wolfe (1998b) suggests, there are problems with using either method on its own. In the experiment reported in this thesis, we will report slopes and slope ratios, but we will be more interested in comparing the efficiency of the search employed between conditions (for example between configural and featural changes, and upright and inverted faces). This will be achieved by comparing the relative slope sizes and slope ratios in different conditions.

The question of whether faces pop out of a visual search array has been addressed in a number of studies. In an early investigation into face detection – the task of detecting the presence of a face in a scene – Nothdurft (1993) investigated the time required to search for a schematic face in an array of jumbled schematic faces. Using array sizes of up to 48 items, he found that the reaction time increased by 113 ms for each additional distractor in the array, suggesting that search was serial rather than parallel in nature.

Kuehn & Jolicoeur (1994) conducted a similar study with schematic faces, looking at the effects of orientation and the similarity of the distractors on visual search performance. Over four experiments, which used inverted schematic faces, inverted digitized faces and scrambled ('configurally distorted') faces as distractors, they failed to find a pop out effect when searching for an upright, intact face in an array of between 1 and 12 faces. The findings with scrambled faces as distractors are of particular note here. Search was slowest amongst scrambled faces which preserved the top-down order and symmetry of a normal face, while scrambled faces which

violated both of these dimensions produced the fastest search rates, suggesting that the more face-like a distractor item is, the more it is likely to slow ones search for an intact face.

Inverted faces, which preserve the spatial relations of faces compared to scrambled stimuli, however, produced faster search times than scrambled faces that violated only one of the above dimensions, but only when trials were grouped by distractor type; when trials were mixed, search among inverted faces was slower than among all other types of distractor. Grouping by distractor type also facilitated searches for the target among 'violated top-down order' distractors compared with mixed trials, but had no effect on searches among 'symmetry violated' distractors. Kuehn & Jolicoeur (1994) suggest that all of these factors (the means by which faces are scrambled, orientation of the face, predictability of the distractor type and similarity of the distractor to the target) may all have an interactive effect on the search strategy used.

Brown, Huey & Findlay (1997) also found no evidence to suggest that faces pop out in a visual search task which used real rather than schematic faces. Participants looked for an upright face among inverted faces, or vice versa, or among jumbled faces, and Brown et al (1997) monitored their eye movements, tracking them to where they first looked after presentation of the array. Participants practised locating either upright faces in inverted distractors or inverted faces in upright distractors. At test, those who had practised locating upright faces were faster and more accurate at locating upright target faces than inverted, while those who practised locating inverted faces showed no difference between upright and inverted targets. A

similar pattern of findings was also found when scrambled faces were used as distractors. This appears to suggest that faces do not pop out, even when real faces are used, and upright faces appear to confer an advantage over inverted faces in peripheral vision.

Lewis & Edmonds (2005) conducted a series of visual search experiments looking at how we detect faces in scrambled natural images, arguably a more ecologically valid method of assessing real face detection. The faces to be detected were arranged within arrays containing distractor cells taken from the same set of images but not containing faces. While such arrays cannot be considered equivalent to natural scenes, the content of the image is largely similar. Experiment 1 explored face detection in a basic visual-search paradigm in arrays of either 4, 9 or 16 items and found that an upright, colour face does appear to pop out of a scrambled scene. A second experiment found that this parallel, preattentive search was eliminated when participants searched for upright faces among inverted faces, suggesting that the increased similarity of the distractors may prevent pop out in this case (e.g. Duncan & Humphreys, 1989; Kuehn & Jolicoeur, 1994).

Two further experiments investigated what aspects of a face may cause it to pop out of a scrambled scene in a visual search experiment. By transforming the images, it was hoped to identify aspects that were important in the preattentive processing. These experiments led to a rejection of colour, fine detail, and orientation as being preattentive guides. The only transformation that significantly reduced the parallel nature of the face detection task was reversal of luminance. It is possible that it is the general shape of the head that is being detected and that luminance reversal

interferes with the extraction of shape-from-shading information. An alternative account, however, is that when detecting a face we look for areas of dark surrounded by areas of light. This strategy has been effectively employed in automatic face-detection systems (Han, Liao, Yu & Chen, 2000).

In the final experiment, the effect of familiarity of the faces and the differences between the detection and recognition of faces was explored. Familiarity was not found to have any discernable effect on face detection. The faces still popped out of the arrays unless they were luminance reversed. The nature of the task required (either detection or recognition) affected the search slope. This suggested a degree of interaction between the processes of detection and recognition with recognition processing beginning before detection was completed. Orientation and luminance-reversal effect were found for familiar faces for detection but larger effects were found for recognition suggesting independent effects of inversion and luminance reversal at two separate stages of processing.

Lewis & Edmonds (2005) suggest that these findings indicate that there are at least three stages involved in the decision that a face is present in an array. The first of these appears to be a preattentive or parallel processing of the array in order to guide attention to a region that may contain a face. The size of the array does affect this stage but only at a level of between 1 and 5ms per item (clearly within the range that is considered to be parallel search; Wolfe, 1998a). This stage of processing is also affected by luminance reversal, suggesting that it is either the distinctive luminance pattern (possibly eyes and mouth) that produces pop-out or the extraction of three-dimensional information that provides the preattentive information. The remaining

two stages are influenced by inversion alone and by colour and clarity together, respectively. The former of these two may involve some rotation of the attended face region. This would explain why it is not influenced by other factors. The remaining stage may be a decision process where a face template is compared with the rotated attended region. The more information present in the image, the faster a decision can be made that the image is in fact a face. If face recognition is required rather than face detection then the third stage is replaced by a recognition process.

The research reviewed in this section has considered the process of detecting a face in a visual array, comprised either of other faces (inverted or scrambled) or of a natural scene, and has looked at the factors, or attributes of faces, which are important for this face detection process. In particular, we (Lewis & Edmonds, 2005) have suggested that the luminance of faces appears to be important in enabling faces to pop out of a scene, while a variety of other factors, including familiarity and orientation, did not effect this search process. Research has also shown the importance of the distractors in determining the search strategy employed, with search being more difficult the more similar (face-like) the distractors were to the target, while the predictability of the distractor type (grouped by distractor type or random) also influenced visual search, interacting with the type of scrambling used. We now turn our attention to a general discussion of our findings, starting with our own visual search experiment, in which we looked at the effects of configural and featural changes to faces in a visual search paradigm, together with the effects of inversion on these processes.

6.2. General Discussion and Conclusions

We began this thesis by looking at the effect of inversion on the detection of configural and featural changes to faces in a visual search task (Chapter 2). We found some evidence of a greater detrimental effect of inversion for configural than featural changes in this experiment, suggesting that there may be a role of configural information in the processing of featural information. We noted in Chapter 1 that a number of discrepant findings have been reported in relation to this issue. While our research appears to support that of Leder & Carbon (2006), further research is clearly needed to investigate whether this role for configural information in processing featural information is task dependent, or even dependent on the measure of performance (see also Yovel & Kanwisher, 2004). These findings, although only observed in the error data in this experiment, nevertheless appear to provide some support for the dual-mode hypothesis of face processing.

Having established a face inversion effect, we then attempted to address the issues outlined above by looking at the effects of face rotation on two main types of task, which are thought to encourage, and have been taken as indicators of, holistic and configural processing respectively; recognition memory of faces in a part-whole paradigm, and the matching of Thatcherised faces. We have used these tasks as a means of addressing three inter-related questions about face processing. First, we have considered the effects of face rotation in terms of the two major theories of face processing (the dual-mode hypothesis and the holistic processing hypothesis), and looked at the extent to which each of these theories is able to account for the effects of rotation on performance on these tasks. A second question is whether the relationship between the degree of rotation and face recognition performance is of a linear or non-

linear nature in these tasks, while a third, related question concerns the means by which face rotation tasks are accomplished. Here we considered the possibility that a mental rotation strategy may be used to complete these tasks by looking at the extent to which the mental rotation hypothesis (e.g. Rock, 1973; Schwaninger & Mast, 2005) was able to account for our findings.

In Chapter 3, we looked at the effects of rotation on a recognition memory task thought to encourage holistic processing, in which participants learned a series of faces (together with their associated names), and were then asked to identify the person, from their whole face and from a target feature, in a two-alternative forced-choice task. The test stimuli were seen in the same orientation as they had been shown in the learning phase. In contrast with our predictions from Chapter 1 we did not find a whole-face advantage for upright faces, however, and we also did not observe any overall effects of inversion on this task, for either whole faces or isolated features. While these findings were surprising, we noted in Chapter 1 that recognition memory studies have found that featural information is explicitly represented for upright faces. Schwaninger et al (2002) found that whole, upright faces could be recognised even when they were scrambled at test, while Bauml et al (1997) also found no effect of face inversion with equally rotated faces in a recognition memory task. Furthermore, Tanaka & Farah (1993) and Rhodes et al (1993) both found that isolated feature tests were not affected by inversion, and more recent studies (Leder & Carbon, 2005) have also found no effects of inversion for feature tests with high quality digital photographs using this part-whole paradigm. This would appear to rule out the possibility that poor image quality was able to account for the lack of an inversion effect for isolated features in previous studies (Lewis & Glenister, 2003).

In conclusion, we suggest that both of these findings pose problems for the holistic view of face processing, as they appear to indicate that faces are not encoded as wholes with no explicit representation of featural information. We suggest that a mental rotation strategy may be employed for rotated faces, with faces rotated to upright (due to our expertise with viewing faces at this orientation), while inverted whole faces may be flipped for recognition, a proposal which is able to account for the higher recognition rates for inverted faces than for those presented at intermediate angles of orientation.

In Chapter 3 we also tested a further prediction of this mental rotation hypothesis suggested by this paradigm - that participants may mentally rotate faces to a stored memory representation prior to recognition – by looking at recognition performance when faces were always seen in the upright orientation in the learning phase, and seen at different angles of rotation at test. Here, consistent with Tanaka & Farah (1993), a whole-face advantage was found for upright but not inverted faces, and there was a clear effect of inversion for whole faces on this task. There was also no effect of inversion for isolated features, and no advantage for upright over inverted isolated features, consistent with our findings from Experiment 2. Taken together, these findings suggest that when faces are learned as upright images and rotated at test, the recognition of whole faces declines as an approximately linear function of the difference in orientation between the learning and test faces, suggesting that participants may be mentally rotating faces prior to recognition (see also Schwaninger & Mast, 2005; Valentine & Bruce, 1988), while we did not find any evidence to suggest that inverted faces may be flipped for recognition on this task. Moreover, the

whole-face advantage for upright faces and the absence of an inversion effect for isolated features can be accounted for by both the dual-mode and holistic hypotheses, as both predict a detrimental effect of rotation for whole faces but not isolated features.

Thus, in these two chapters we have looked at the effects of rotation on a task thought to encourage holistic face processing, and found differential effects for equally compared to differently rotated faces. Indeed, in contrast with previous research (e.g. Tanaka & Farah, 1993; Lewis & Glenister, 2003) we found no evidence of holistic processing when the faces were always seen in the same orientation, with no inversion effects observed for whole faces, and no differences between whole faces and isolated features. In contrast with equally rotated faces, for differently rotated faces the recognition of whole faces declined as an approximately linear function of the difference in orientation between the learning and test faces, suggesting that participants may be mentally rotating faces prior to recognition.

It is important to note, however, that in the experiments reported here, participants only saw each of the faces in the learning phase twice, whereas other studies have used five presentations (Tanaka & Farah, 1993; Lewis & Glenister, 2003; Leder & Carbon, 2005). While this may account for the lower recognition rates observed in the current studies, particularly in Experiment 2, a comparison of Experiments 2 and 3 suggests that these low accuracy rates can at least in part be seen as the result of learning the faces at different angles of orientation in Experiment 2. Nevertheless, while there are many other differences between the two sets of studies, it is possible that a lack of power in these experiments means that the findings cannot

easily be compared with those of the tasks using the Thatcher illusion. A clear goal for future research, therefore, is to replicate these experiments with a greater number of presentations of the faces in the learning phase, in order to achieve the maximum possible recognition rates for each angle of orientation. Meanwhile, a further challenge for future research is to find a face perception task which encourages holistic processing, and displays a pattern of performance consistent with holistic processing when equally rotated faces are oriented away from the upright.

In Chapter 4, we turned our attention to look at the effects of rotation on a face processing task known to disrupt configural information, the Thatcher illusion. A same-different matching task, in which participants decided whether two simultaneously presented faces were both normal, both Thatcherised or one normal and one Thatcherised was used to address this question, and facial identity (same or different) was manipulated in an attempt to encourage featural and configural processing strategies respectively.

Our findings were in contrast with previous research (Boutsen & Humphreys, 2003), which failed to find an effect of inversion for Thatcherised face pairs when only upright and inverted faces were considered. For same identity matches, approximately equivalent effects of rotation were observed for normal and Thatcherised face pairs, suggesting that the matching of these face pairs involved a similar featural processing strategy. The matching of images depicting two different people, however, showed effects of rotation for both types of face pairs, but here the effect was significantly greater for Thatcherised than for normal faces. In both experiments, reaction times to make correct 'different' type responses increased as a

function of rotation from the upright. Thus, our findings suggest that when matching can be image-based as in Experiment 4, the loss of featural information through rotation is no greater for Thatcherised face pairs than for normal pairs, but Thatcherisation disproportionately disrupts the matching of face pairs in a task which encourages configural processing (matching different-identity face pairs).

While the results obtained here can be explained by the distinction between configural and featural processing, these findings pose problems for a strict interpretation of the holistic view of face processing, while the view of holistic as the extent to which the image forms a coherent Gestalt is also limited in its ability to account for these findings. Meanwhile, the majority of the data are consistent with a linear relationship between rotation and face matching performance, a finding which is consistent with the view that participants may be using a mental rotation strategy for rotated faces, with faces rotated to upright due to our expertise with viewing faces at this orientation.

One possible explanation for the differences between these findings and those of previous research is that seeing the face pairs at a number of angles of orientation between upright and inverted somehow encouraged a different processing strategy to be employed than when only upright and inverted faces were seen. This question was addressed in Experiment 6 by looking at the matching of different-identity face pairs in the upright and inverted orientations only, and comparing the findings to those reported with rotated (different-identity) pairs. There was no evidence to suggest that the intermediate angles of rotation had affected the overall effects of inversion in

Experiment 5, or were responsible for the differences observed between our studies and those of Boutsen & Humphreys (2003).

Experiment 7 looked at the matching of same or different images of a person, in an attempt to encourage a more configural processing strategy than has previously been possible with same-person pairs. There was a greater effect of rotation for different-image than for same-image pairs for 'same' decisions, while the effect of rotation was greater for different-image than for same-image pairs for Thatcherised faces. For normal faces, however, the effect of rotation was roughly the same for both types of face pairs. These findings appear to suggest that the effect of rotation is independent of face type only when the face pairs contain identical images of the same person. The finding from Experiments 5 and 6 with different-identity pairs that Thatcherised faces are affected more by rotation than normal faces are also consistent with this suggestion. Thus, the evidence presented here appears to be largely consistent with the view that rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces.

In the final experiment of this chapter, we looked at the effect of rotation on the matching of facial identity with same and different face type pairs. We found that when 'same' and 'different' identity judgements were analysed separately, the effect of rotation was greater for Thatcherised than for normal 'same' identity judgements. For 'different' identity judgements, however, the effect of rotation was similar for both types of faces. These findings appear to indicate that judging the identity of two simultaneously presented faces requires featural information to distinguish between different identity pairs, and this information is not differentially affected by face type.

Configural information, on the other hand, is used to judge that the two faces are the same and this is affected by rotation and Thatcherisation.

One possible confound of these findings, however, is that the faces used in this experiment were not matched for visual similarity in terms of age, gender, hair colour etc. It is therefore possible that participants were making image-based 'different' identity decisions, rather than using featural information as described above. A further goal for future research, therefore, is to re-address this question with faces which can be matched for visual similarity, to look at the effects of rotation on the matching of different-identity face pairs.

With regard to the implications of our findings for the holistic and dual-mode accounts of face processing, we suggest that only when holistic is defined in terms of whether the Gestalt of the overall image is coherent is this theory able to account for any of the findings reported in Chapter 5. We found faster recognition times for inverted normal faces than inverted Thatcherised faces in all of the experiments reported here, suggesting that when matching cannot be based on image-based featural information, participants may revert to a holistic processing strategy for inverted faces, which is able to account for the greater effect of inversion for Thatcherised than for normal face pairs. We note, however, that previous tests of this version of the holistic account have not compared upright and inverted faces directly, and as such, this theory is currently unable to offer strict predictions about the effects of inversion and rotation on face matching tasks.

In both of these chapters we have presented evidence to suggest that the distinction between local and global configural information does not provide a useful framework for understanding the effects of rotation on face matching tasks, as Boutsen & Humphreys (2003) suggest (see Chapter 1). Rather, the distinction between configural and featural information on the other hand is readily able to account for the findings observed in these experiments. The evidence presented here shows that when featural information can be used to complete the match (i.e. when the face pairs contain identical images of the same person), the effect of rotation is independent of face type. When different-identity or different-image pairs are shown and configural information is required, however, Thatcherised faces are affected more by rotation than normal faces. Rotation of Thatcherised faces disrupts configural information and does so to a much greater extent than for normal faces.

Meanwhile, as was the case in Chapter 4, the majority of the data suggested that a linear function provided the best account of the data. These findings are in broad agreement with the prediction of the mental rotation hypothesis that recognition performance is dependent on the amount of normalisation required, and that participants may be using a mental rotation strategy to discriminate between rotated faces. Again, we found little evidence to suggest that an alternative, flipping strategy was being used for inverted faces, with a significant improvement in performance between 135° rotated and inverted faces only being observed for different image Thatcher faces in Experiment 7. Thus, in terms of the means by which face rotation tasks are accomplished, our overall pattern of findings appears to suggest that performance is dependent on the amount of normalisation required. These findings appear to be at odds with those from other face rotation studies cited in the

Introduction (e.g. Schwaninger & Mast, 2005; McKone, 2004), although Schwaninger & Mast (2005) did not consider the nature of this relationship directly, whilst McKone (2004) observed curvilinear effects in a set of experiments with very different task demands. We suggest that future studies in this area may seek to address the relationship between face processing and rotation in greater detail to assess the extent to which the effects of rotation are task dependent, and under which circumstances each pattern of performance may be observed.

Thus in Chapters 4 and 5 we have found that when a face matching task requires the matching of featural information, the effect of rotation is approximately equal for normal and Thatcherised faces, but Thatcherisation disproportionately disrupts matching performance when different-identity pairs are to be matched, requiring configural information. Moreover, the effect of rotation is independent of face type only when the face pairs contain identical images of the same person, while the matching of facial identity requires featural information to distinguish between different identity pairs (although this issue clearly requires further investigation), and this information is not differentially affected by face type. Configural information, on the other hand, is used to judge that the two faces are the same and this is affected by rotation and Thatcherisation.

6.3. Future Research

In addition to the improvements to the methodologies employed in this research suggested earlier, we suggest that a sequential matching paradigm could be employed to investigate whether participants may use a mental rotation strategy to match normal and Thatcherised faces, as used in Chapters 4 and 5. Instead of a

simultaneous matching task, here participants would see one face (normal or Thatcherised) in the upright orientation, followed by another face, of either the same or a different type, at either the same or a different angle of rotation, from upright through to inverted. This would allow us to look in detail at the strategies taking place for the matching of same and different face types, and also to look at how configural and featural information is affected on this type of task.

This sequential matching strategy would also enable us to look at and track the effects of face rotation over time, an area which has been considered with object and picture matching tasks, but not in the face processing literature. Murray (1999) found evidence to suggest that in a sequential same-different matching task with drawings of objects, the effects of orientation were reduced following experience with the misoriented items. Murray (1999) suggests that over time, participants may develop object-invariant representations, replacing object-specific representations which are initially formed upon the single presentation of a rotated object. Schwaninger & Mast (2005) found no evidence of a reduced effect of orientation for faces in their study. We suggest, however, that an experiment which is designed to address this question directly is necessary to further our understanding of the processes of face rotation. Such an experiment would allow us to look more directly at the processing strategies used by participants, and also at the types of information which may be important for these object-invariant representations, if indeed they are formed with faces as they are with objects.

We also aim to investigate further the impact of rotation on visual search tasks. The research presented here appears to suggest that the detection of configural

and featural changes in a visual search array is of a serial, rather than parallel search. Here the question we wish to address is that of the effect of rotation on a task which appears to involve holistic rather than configural information, specifically considering whether the relationship is linear or non-linear, and what processing strategies might be being used to complete this task.

7.2. General Conclusions

In sum, the findings reported in this thesis have furthered our understanding of the effects of rotation in face processing, and have considered (and found some evidence to support) the mental rotation hypothesis as a means of understanding how face rotation tasks are accomplished. The evidence presented has shown differential effects of rotation on whole and part-based information in a recognition memory task, while different processing strategies appear to be involved in the processing of different and equally rotated faces. By looking at the matching of normal and Thatcherised faces, we have also provided a detailed account of the effects of rotation on configural and featural information in face processing, with Thatcherisation disproportionately disrupting matching performance when configural information is required. Moreover, these effects are dependent on the nature of the image seen and, for identity matches, the nature of the decision to be made, although further research is clearly needed in this regard.

References

- Bartlett, J. C., & Searcy, J. (1993). Inversion and configuration of faces. *Cognitive Psychology*, 25, 281-316.
- Barton, J. J. S., Deepak, S., & Malik, N. (2003). Attending to faces: change detection, familiarization, and inversion effects. *Perception*, 32, 15-28.
- Barton, J. J. S., Keenan, J. P., & Bass, T. (2001). Discrimination of spatial relations and features in faces: effects of inversion and viewing duration. *British Journal of Psychology*, 92, 527-549.
- Barton, J. J. S., Press, D. Z., Keenan, J. P., & O'Connor, M. (2002). Lesions of the fusiform face area impair perception of facial configuration in prosopagnosia. *Neurology*, 58, 71-78.
- Bauml, K. H., Schnelzer, M., & Zimmer, A. (1997). Discrimination and association processes for faces and non-faces: The effect of rotation. *Acta Psychologica*, 95, 107-118.
- Boutsen, L. & Humphreys, G. W. (2003). The effect of inversion on the encoding of normal and "thatcherized" faces. *Quarterly Journal of Experimental Psychology A*, 56 (6), 955-975.

- Brown, V., Huey, D., & Findlay, J. M. (1997). Face detection in peripheral vision: do faces pop out? *Perception*, *26*, 1555-1570.
- Bruce, V., Hanna, E., Dench, N., Healey P., & Burton, A. M. (1992). The importance of mass in line drawings of faces. *Applied Cognitive Psychology*, *6*, 619-628.
- Bruce, V., & Young A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*, 305-327.
- Burton, A. M., Bruce, V. & Hancock, P. J. B. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, *23*, 1-31.
- Burton, A. M., Bruce, V. & Johnston, R. A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology*, *81*, 361-380.
- Cabeza, R., & Kato, T. (2000). Features are also important: Contributions of featural and configural processing to face recognition. *Psychological Science*, *11*, 429-433.
- Carbon, C. C., & Leder, H. (2005). When feature information comes first! Early processing of inverted faces. *Perception*, *34*, 1117-1134.
- Carbon, C. C., Schweinberger, S. R., Kaufmann, J. M. & Leder, H. (2005). The Thatcher illusion as seen by the brain: An event-related brain potentials study. *Cognitive Brain Research*, *24*, 544-555.

Carey, S., & Diamond, R. (1977). From piecemeal to configurational representation of faces. *Science*, *195*, 312-314.

Carey, S., & Diamond, R. (1994). Are faces perceived as configurations more by adults than by children? *Visual Cognition*, *1*, 253-274.

Collishaw, S. M., & Hole, G. J. (2000). Featural and configurational processes in the recognition of faces of different familiarity. *Perception*, *20*, 893-909.

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

Corballis, M. C., Zbrodoff, N. J., Shetzer, L. I., & Butler, P. B. (1978). Decisions about identity and orientation of rotated letters and digits. *Memory and Cognition*, *6*, 98-107.

Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107-117.

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

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433-458.

Eger, E., Schweinberger, S. R., Dolan, R. J., & Henson, R. N. (2005). Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *Neuroimage*, *26*, 1128-1139.

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*, 628-634.

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

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

Gauthier, I., & Tarr, M. J. (2002). Unravelling mechanisms for expert object recognition: Bridging brain activity and behaviour. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 431-446.

Han, C. C., Liao, H. Y. M., Yu, G. J., & Chen, L. H. (2000). Fast face detection via morphology-based pre-processing. *Pattern Recognition*, *33*, 1701-1712.

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

Kuehn, S. M., & Jolicoeur, P. (1994). Impact of quality of the image, orientation, and similarity of the stimuli on visual search for faces. *Perception, 23*, 95-122.

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

Leder, H. (1999). Matching person identity from facial line drawings. *Perception, 28*, 1171-1175.

Leder, H., & Bruce. V. (2000). When inverted faces are recognised: the role of configural information in face recognition. *Quarterly Journal of Experimental Psychology A, 32*, 513-536.

Leder, H., & Carbon, C. C. (2005). When context hinders! Learn-test compatibility in face recognition. *Quarterly Journal of Experimental Psychology A, 58*, 235-250.

Leder, H., & Carbon, C. C. (2006). Face-specific configural processing of relational information. *British Journal of Psychology, 97*, 19-29.

Lewis, M. B. (2001). The Lady's not for turning. Rotation of the Thatcher illusion. *Perception, 30*, 769-774.

Lewis, M. B., & Edmonds, A. J. (2005). Searching for faces in scrambled scenes. *Visual Cognition, 12*, 1309-1336.

Lewis, M. B., & Glenister, T. E. (2003). A sideways look at configural encoding: Two different effects of face rotation. *Perception, 32*, 7-14.

Lewis, M. B., & Johnston, R. A. (1997). The Thatcher illusion as a test of configural disruption. *Perception, 26*, 225-227.

Malcolm, G. L., Leung, C., & Barton, J. J. S. (2005). Regional variation in the inversion effect for faces: differential effects for feature shape, feature configuration, and external colour. *Perception, 34*, 1221-1231.

Martini, P., McKone, E., & Nakayama, K. (2006). Orientation tuning of human face processing estimated by contrast matching in transparency displays. *Vision Research, 46*, 2102-2109.

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

McKone, E. (2004). Isolating the special component of face recognition: Peripheral identification and a Mooney face. *Journal of Experimental Psychology: Learning, Memory & Cognition, 30* (1), 181-197.

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

Mondloch, C. J., Le Grand, R., & Maurer, D. (2002). Configural processing develops more slowly than feature face processing. *Perception, 31*, 553-566.

Mooney, C. M. (1957). Age in the development of closure ability in children. *Canadian Journal of Psychology, 11*, 219-226.

Murray, J. E. (1999). Orientation-specific effects in picture matching and naming. *Memory and Cognition, 27*, 878-889.

Murray, J. E., Yong, E., & Rhodes, G. (2000). Revisiting the perception of upside-down faces. *Psychological Science, 11*, 498-502.

Nothdurft, H. C. (1993). Faces and facial expressions do not pop out. *Perception, 22*, 1287-1298.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science, 8*, 368-373.

Riesenhuber, M., Jarudi, I., Gilad, S., & Sinha, P. (2004). Face processing in humans is compatible with a simple shape-based model of vision. *Proceedings of the Royal Society of London: Series B-Biological Sciences, 271*, S448-450.

Rhodes, G. (1988). Looking at faces – first and second order features as determinants of facial appearance. *Perception, 17*, 43-63.

Rhodes, G., Brake, S., & Atkinson, A. P. (1993). What's lost in inverted faces? *Cognition*, *47*, 25–57.

Robbins, R., & McKone, E. (2003). Can holistic processing be learned for inverted faces? *Cognition*, *88*, 79–107.

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

Rock, I. (1988). On Thompson's inverted-face phenomenon. *Perception*, *17*, 815-817.

Schwaninger, A., & Mast, F. (1999). Why is face recognition so orientation-sensitive? Psychophysical evidence for an integrative model. *Perception*, *28* (Suppl.), 116.

Schwaninger, A., & Mast, F. (2005). The face-inversion effect can be explained by the capacity limitations of an orientation normalization mechanism. *Japanese Psychological Research*, *47*(3), 216-222.

Schwaninger, A., Lobmaier, J., & Collishaw, S. M. (2002). Role of featural and configural information in familiar and unfamiliar face recognition. *Lecture Notes in Computer Science*, *2525*, 643-650.

Schweinberger, S. R., Pickering, E. C., Burton, A. M., & Kauffman, J. M. (2002). Human brain potential correlates of repetition priming in face and name recognition. *Neuropsychologia*, *40*(12), 2057-2073.

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

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

Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology, 81*, 275-280.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171*, 701-703.

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

Tanaka, J. W., & Farah, M. J. (1991). Second-order relational properties and the inversion effect: Testing a theory of face perception. *Perception and Psychophysics, 50*, 367-372.

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

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

Tanaka, J. W., Kay, J. B., Grinnell, E., Stansfield, B., & Szechter, L. (1998). Face recognition in young children: when the whole is greater than the sum of its parts. *Visual Cognition*, 5(4), 479-496.

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

Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. New York: Cambridge University Press.

Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.

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

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

Valentine, T., & Bruce, V. (1985). What's up? The Margaret Thatcher illusion revisited. *Perception*, 14, 515-516.

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

White, R. (2002). Different spatial-relational information is used to recognise faces and emotional expressions. *Perception*, *31*, 675-682.

Wolfe, J. M. (1994). Guided search 2.0 - a revised model of visual-search. *Psychonomic Bulletin & Review*, *1*, 202-238.

Wolfe, J. M. (1998a). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13-74). Hove, UK: Psychology Press.

Wolfe, J. M. (1998b). What can 1 million trials tell us about visual search? *Psychological Science*, *9*, 33-39.

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

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

Young, A. W., Hellowell, D., & Hay, D. C. (1987). Configural information in face perception. *Perception*, *16*, 747–759.

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

Yovel, G., & Kanwisher, N. (2004). Face perception: domain specific, not process specific. *Neuron*, 44, 889-898.

