Processing of emotional expression in subliminal and low-visibility images

Submitted by Hannah Lucy Filmer to the University of Exeter
as a thesis for the degree of
Doctor of Philosophy in Psychology
May 2012.

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other university.

Signature:……………………………………….

Abstract

This thesis investigated the processing of emotional stimuli by the visual system, and how the processing of emotions interacts with visual awareness. Emotions have been given ‘special’ status by some previous research, with evidence that the processing of emotions may be relatively independent of striate cortex, and less affected by disruption to awareness than processing of emotionally neutral images. Yet the extent to which emotions are ‘special’ remains questionable. This thesis focused on the processing of emotional stimuli when activity in V1 was disrupted using transcranial magnetic stimulation (TMS), and whether emotional properties of stimuli can be reliably discriminated, or affect subsequent responses, when visibility is low.

Two of the experiments reported in this thesis disrupted activity in V1 using TMS, Experiment 1 with single pulses in an online design, and Experiment 2 with theta burst stimulation in an offline design. Experiment 1 found that a single pulse of TMS 70-130 ms following a presentation of a body posture image disrupted processing of neutral but not emotional postures in an area of the visual field that corresponded to the disruption. Experiment 2 did not find any convincing evidence of disruption to processing of neutral or emotional faces. From Experiment 1 it would appear that emotional body posture images were relatively unaffected by TMS, and appeared to be robust to disruption to V1. Experiment 2 did not add to this as there was no evidence of disruption in any condition.

Experiments 3 and 4 used visual masking to disrupt awareness of emotional and neutral faces. Both experiments used a varying interval between the face and the mask stimuli to systematically vary the visibility of the faces. Overall, the shortest SOA produced the lowest level of visibility, and this level of visibility was arguably outside awareness. In Experiment 3, participants’ ability to discriminate properties of emotional faces under low visibility conditions was greater than their ability to discriminate the orientation of the face. This was despite the orientation discrimination being much easier at higher levels of visibility. Experiment 4 used a gender discrimination task, with emotion providing a redundant cue to the decision (present half of the time). Despite showing a strong linear masking function for the neutral faces, there was no evidence of any emotion advantage. Overall, Experiment 3 gave some evidence of an emotion advantage

under low visibility conditions, but this effect was fairly small and not replicated in Experiment 4.

Finally, Experiments 5-8 used low visibility emotional faces to prime responses to subsequent emotional faces (Experiments 5 and 6) or words (Experiments 7 and 8). In Experiments 5, 7 and 8 there was some evidence of emotional priming effects, although these effects varied considerably across the different designs used. There was evidence for meaningful processing of the emotional prime faces, but this processing only led to small and variable effects on subsequent responses.

In summary, this thesis found some evidence that the processing of emotional stimuli was relatively robust to disruption in V1 with TMS. Attempts to find evidence for robust processing of emotional stimuli when disrupted with backwards masking was less successful, with at best mixed results from discrimination tasks and priming experiments. Whether emotional stimuli are processed by a separate route(s) in the brain is still very much open to debate, but the findings of this thesis offers small and inconsistent evidence for a brain network for processing emotions that is relatively independent of V1 and visual awareness. The network and nature of brain structures involved in the processing of subliminal and low visibility processing of emotions remains somewhat elusive.

Acknowledgments

During the last 4 years I have received support – academically and personally – from a number of different people. This support has been crucial, and without the help of my supervisors, peers and family this thesis would have been impossible.

Firstly, my PhD would not have happened without financial support I was given by the ESRC.

I'd like to thank my supervisors. Prof. Stephen Monsell has given expert and patient guidance for which I am very grateful. I also received valuable comments and feedback from Dr. Aureliu Lavric at several points in the last few years. Early on in my PhD, Prof. Vince Walsh provided advice on TMS experimental design that was a huge help at a time when there was no TMS expert at Exeter.

Many members of the psychology department also provided valuable feedback on my research throughout my time at Exeter.

The last few years have been made much easier, and considerably more fun, by a number of MSc and PhD students, past and present, in the department. My officemate Felice provided many lighter moments, along with Charlotte, Tobias, Nik, Dale, Bibiana, Heike, Jonathan, and David. My parents were also a fantastic source of support, and provided many hours proof reading assistance. Finally, my partner, Sam, has provided love, support, and many delicious meals that made the last few years particularly enjoyable.

Table of Contents

ABSTRACT **3**

ACKNOWLEDGMENTS **5**

TABLE OF CONTENTS **7**

LIST OF FIGURES **11**

LIST OF TABLES **17**

DECLARATION **19**

CHAPTER 1: LITERATURE REVIEW **21**

1.1. INTRODUCTION 21
1.2. VISUAL AWARENESS 21
1.3. THEORIES OF VISUAL AWARENESS AND ITS DISRUPTION 26
1.4. SUBLIMINAL RESPONSE PRIMING AND CURRENT CONTROVERSIES 34
1.5. EMOTIONAL PROCESSING 42
1.6. SUBLIMINAL EMOTIONAL STIMULI AND THEIR IMPACT ON BEHAVIOUR 45
1.7. LOW VISIBILITY EMOTIONAL PRIMING: ALTERNATIVES TO VALENCE CONGRUENCE EFFECTS .. 50
1.8. PROBLEM OF MEASURING AWARENESS 52
1.9. THE PROCESSING OF EMOTIONAL STIMULI IN THE BRAIN 55
1.10. A DUAL ROUTE TO VISUAL PROCESSING? 58
1.11. A BRIEF WORD ABOUT SUBLIMINAL PROCESSING OF NON-EMOTIONAL STIMULI 66
1.12. A BRIEF WORD ABOUT VISUAL ATTENTION 67
1.13. SUMMARY 70

CHAPTER 2: TRANSCRANIAL MAGNETIC STIMULATION: BASIC PRINCIPLES, POSSIBILITIES, AND APPLICATIONS TO RESEARCH IN VISUAL PROCESSING **73**

2.1. INTRODUCTION 73
2.2. THE NEUROSCIENCE 'NICHE' OF TMS 73
2.3. HOW TMS WORKS 74
2.4. SPATIAL RESOLUTION OF TMS 75
2.5. STIMULATION PROTOCOLS AND TEMPORAL RESOLUTION 76
2.6. MAIN ADVANTAGES AND DISADVANTAGES OF TMS 79
2.7. IMPLICATIONS FOR THE PRESENT STUDIES 81

CHAPTER 3: TMS TO V1 SPARES DISCRIMINATION OF EMOTIVE RELATIVE TO NEUTRAL BODY POSTURES 83
3.1. ABSTRACT 83
3.2. INTRODUCTION 83
3.3. METHOD 89
  3.3.1. Participants 89
  3.3.2. Stimuli 89
  3.3.3. Materials 91
  3.3.4. Procedure 91
3.4. RESULTS 94
  3.4.1. Fixation check and stimulus eccentricity 94
  3.4.2. Classification performance 94
3.5. DISCUSSION 98
CHAPTER 4: THE EFFECT OF THETA BURST STIMULATION ON V1 103
4.1. INTRODUCTION 103
4.2. METHOD 105
  4.2.1. Participants 105
  4.2.2. Materials 105
  4.2.3. Stimulation sites and cTBS 106
  4.2.4. Procedure 107
4.3. RESULTS 109
  4.3.1. PT and MT changes 109
  4.3.2. Number task 110
  4.3.3. Face discrimination – overall results 110
  4.3.5. Emotion specific effects 113
  4.3.6. Face discrimination: impact of hemisphere stimulated 115
4.4. DISCUSSION 117
CHAPTER 5: EMOTIONS ARE SPECIAL: USING VISUAL MASKING TO DISSOCAITE PROCESSING OF NEUTRAL AND EMOTIONAL STIMULI 121
5.1. INTRODUCTION 121
5.2. EXPERIMENT 3 124
  5.2.1. Method 124
  5.2.2. Results 127
  5.2.3. Discussion 132
5.3. EXPERIMENT 4 133

5.3.1. Methods *134*

5.3.2. Results *136*

5.3.3. Discussion *141*

5.4. GENERAL DISCUSSION 143

CHAPTER 6: INFLUENCING BEHAVIOUR WITH SUBLIMINAL EMOTIONAL FACES: FOUR PRIMING EXPERIMENTS 145

6.1. INTRODUCTION 145

6.2. EXPERIMENT 5 149

6.2.1. Methods *149*

6.2.2. Results *153*

6.2.3. Discussion *156*

6.3. EXPERIMENT 6 157

6.3.1. Methods *158*

6.3.2. Results *159*

6.3.3. Discussion *161*

6.4. EXPERIMENT 7 162

6.4.1. Methods *163*

6.4.2. Results *166*

6.4.3. Discussion *170*

6.5. EXPERIMENT 8 172

6.5.1. Method *174*

6.5.2. Results *176*

6.5.3. Discussion *182*

6.6. GENERAL DISCUSSION 184

CHAPTER 7: SUMMARY, METHODOLOGICAL REFLECTIONS, AND FUTURE DIRECTIONS 189

7.1. INTRODUCTION 189

7.2. THEME SUMMARIES 189

7.2.1. V1's role in visual awareness and emotion processing *189*

7.2.2. The effect of visual masking on decision making about emotional stimuli190

7.2.3. Subliminal emotional priming *191*

7.3. WHY WHERE THE RESULTS INCONSISTENT? 192

7.3.1. Methodological difficulties *192*

7.3.2. Failures to replicate *194*

7.4. PRIVILEGED EMOTION PROCESSING? 195

7.4.1. Alternative accounts *196*

7.4.2. Anatomical basis *197*

7.5. MEASURES OF AWARENESS: CHOICE OF MEASUREMENT IN THIS THESIS 198
7.6. FUTURE RESEARCH DIRECTIONS 200
    7.6.1. Replications 200
    7.6.2. Extensions 201
7.7. CONCLUSIONS 202

REFERENCES 203

APPENDIX 1: ITEMS ANALYSIS FOR EXPERIMENT 2 221

APPENDIX 2: ITEMS ANALYSIS FOR EXPERIMENTS 6-8 223

APPENDIX 3: MEDIAN ANALYSIS EXPERIMENTS 5-8 225

List of figures

Figure 1.1..30

Connectivity matrix summarising the interconnections within visual cortex and between visual cortex and other cortical areas. Reproduced from Felleman & Van Essen, 1991.

Figure 2.1..78

Paradigms of TBS and Their Effects on MEPs. Reproduced from Huang et al. (2005).

Figure 3.1..90

The body posture images making up the two categories: (a) the aggressive postures and two similar neutral postures; (b) the pleasant postures and the two similar neutral postures. (c) Sequence of events on a standard trial. (d) Coordinates (in degrees of visual angle from fixation) of each participant's phosphene location.

Figure 3.2..95

Group mean categorisation performance (d') for each stimulus type (emotive and neutral) at each stimulus location (phosphene or control location).

Figure 3.3..96

Group mean categorisation performance (d') as a function of stimulus location, stimulus type, and pulse onset time.

Figure 3.4..97

Relation between overall categorisation performance and the size of the interaction depicted in Figure 2: interaction contrast = (phosphene location emotive – phosphene neutral) – (control location emotive – control location neutral).

Figure 4.1..106

Examples of the stimuli used (a), the location of the phosphenes reported by participants (b), and a standard trial outline (c).

Figure 4.2..111

Group mean categorisation performance (d') for each time, stimulus type and stimulus location.

Figure 4.3..112

Group mean categorisation performance (d') for each stimulus location, stimulus type and time, for version 1 (left) and version 2 (right) of the experiment. Version 1 represents angry male and happy female faces, version 2 happy male and angry female faces.

Figure 4.4..113

Group mean categorisation performance (d') for each of the conditions (stimulus type, stimulus location and time) split across quarters of each time set.

Figure 4.5..114

Group mean categorisation performance (d') for each condition for participants who noticed the emotional expressions (left) and those who did not (right).

Figure 4.6..115

Group mean categorisation performance (proportion correct) for each condition for the angry (left) and happy (right) emotions and equivalent neutral images.

Figure 4.7..116

Mean categorisation performance (d') for the right (top) and left (bottom) hemisphere participants who completed versions 1 (left) or 2 (right) of the experiment. Each graph represents data from 4 participants.

Figure 5.1..124

Experiment 3: (a) examples of the target stimuli (top happy, bottom angry, left young, right old), (b) examples of mask stimuli used as backwards and forwards masks, (c) standard trial outline.

Figure 5.2..128

Group mean discrimination accuracy for Experiment 3, for all 5 tasks at all SOAs (a), the emotion and orientation tasks at 10ms SOA (b), and the emotion and orientation tasks broken down by emotion type (upright trials – (c), inverted trials – d)).

Figure 5.3..131

Group mean discrimination from Experiment 3 for the emotion and orientation tasks per session run, with (a) showing session one, (b) session two, and (c) session three.

Figure 5.4..135

Examples of the stimuli used for the target faces (a), showing neutral (left), happy (middle) and angry (right) expressions, and examples of the mask stimuli (b), used in Experiment 4.

Figure 5.5..137

Group mean categorisation performance (d') for Experiment 4, for the emotional and neutral faces, for upright (left) and inverted (right) face blocks.

Figure 5.6..139

Group mean categorisation performance (proportion correct) for Experiment 4, for the upright (a & c) and inverted (b & d) faces for the angry and equivalent neutral faces (a & b) and for the happy and equivalent neutral faces (b & d).

Figure 5.7..141

Group mean categorisation performance (proportion correct) for Experiment 4, for the emotional and neutral faces, for upright (a & c) and inverted (b & d) blocks, in session one (a & b) and session two (c & d).

Figure 6.1..150

Examples of the prime, mask and target stimuli (top) and a standard trial outline for the priming part of the experiment (bottom) in Experiment 5.

Figure 6.2...153

Mean discrimination performance (d') for the three visibility tasks in Experiment 5. Error bars represent SEM.

Figure 6.3...154

Congruency effects in RTs (top) and error rates (bottom) for the upright and inverted prime conditions in Experiment 5.

Figure 6.4...155

Group mean performance at the forced choice neutral trials for Experiment 5. Error bars represent SEM for performance at each prime orientation.

Figure 6.5...159

Mean group discrimination performance (d') for the three visibility tasks in Experiment 6.

Figure 6.6...160

Mean congruence effects in RTs (top) and error rates (bottom) for each prime-target stimulus type in Experiment 6.

Figure 6.7...164

Examples of the stimuli used for the prime faces (a), showing neutral (left), happy (middle) and angry (right) expressions, examples of the mask stimuli (b), examples of the target word stimuli (c), and trial outline (d) in Experiment 7.

Figure 6.8...167

Group mean discrimination performance (d') at the visibility task (left) and the group mean subjective ratings (scale of 1-7) of visibility (right) from Experiment 7.

Figure 6.9...168

Mean RTs (top) and error rates (bottom) for the negative (left) and positive (right) target words, shown for each prime type (angry, happy and neutral) from Experiment 7.

Figure 6.10...169

Mean congruence effect (incongruent trials – congruent trials) between prime and target at each of the SOAs for the RTs (left) and the error rates (right).

Figure 6.11...177

Group mean discrimination performance (d') of prime face orientation for the inner and outer target word location conditions from the visibility test in Experiment 8.

Figure 6.12...178

The interaction between target valence and prime emotion for the RTs (top) and the error rates (bottom) from Experiment 8.

Figure 6.13...179

The interaction between target valence and prime emotion for the inner (a & b) and outer (c & d) locations for the RTs (a & c) and error rates (b & d) from Experiment 8.

Figure 6.14...180

The interaction between prime emotion and SOA for the inner (a & b) and outer (c & d) locations for the RTs (a & c) and error rates (b & d) from Experiment 8.

Figure 6.15...181

Group mean discrimination performance (d') for the decision bias blocks from Experiment 8.

Figure 6.16...182

Group mean RTs for the inner (top) and outer (bottom) locations for the decision bias blocks from Experiment 8.

List of Tables

Table 3.1..**89**

Average ratings and average discrimination performance (% correct) for each of the eight stimuli used. A low rating (1) means very aggressive, a high rating (7) very pleasant, and a mid-rating (4) neutral.

Table 3.2..**98**

Average discrimination performance (d') for each stimulus type and stimulus location, for participants stimulated in the left hemisphere, and the right hemisphere.

Table 5.1..**126**

All 4 of the possible block orders in Experiment 3, varied between participants.

Table 5.2..**129**

F values for the ANOVA's run on each task in Experiment 3, with the factors of SOA (and inversion for emotion and age tasks). * indicates p<0.05, ** indicates p<0.001.

Table 6.1..**154**

Mean RT (ms) and error rate (%) for each condition in Experiment 5.

Table 6.2..**160**

Mean RT (ms) and error rates (%) for each trial type for Experiment 6.

Table A1.1..**221**

Mean accuracy (% correct) for each of the ninety-six stimuli used in Experiment 2. * Stimuli were outliers in that performance was less than 2SD below the overall mean.

Table A1.2..**222**

Group mean categorisation performance (% correct) for each condition of Experiment 2 before items exclusions.

Table A1.3..**222**

Group mean categorisation performance (% correct) for each condition of Experiment 2 following items exclusions.

Table A2.1..**223**

Mean, standard deviation, minimum and maximum RT (ms) and error rate (% incorrect) for the stimuli used in Experiments 6-8.

Table A2.2..**224**

Summary of the mean RTs (ms) with and without the exclusion of outlier stimuli for Experiment 7.

Table A2.3..**224**

Summary of the mean RTs (ms) with and without the exclusion of outlier stimuli for Experiment 8.

Table A3.1..**225**

Median RT (ms) for each prime orientation, prime congruency, and target emotion from Experiment 5.

Table A3.2..**225**

Median RT (ms) for each prime orientation, prime congruency, and target emotion from Experiment 6.

Table A3.3..**226**

Median RT (ms) for each Target valence, prime emotion, and prime – target SOA from Experiment 7.

Table A3.4..**226**

Median RT (ms) for each target valence, prime emotion, and prime – target SOA from Experiment 8. The data in this table have been collapsed across location.

Declaration

The research reported in this thesis was carried out at the University of Exeter between October 2008 and December 2011. During this time I was supervised by Prof. Stephen Monsell and Dr. Aureliu Lavric.

This thesis has not been submitted, in whole or in part, for any other degree, or qualification, at any other University. Chapter 3 is an article that will be submitted to a journal for publication in the next few months. I conducted all the experiments reported in this thesis, ran the analyses, wrote the first drafts and prepared all the figures and tables (unless otherwise stated). Prof. Stephen Monsell is a co-author on Chapter 3, and provided editing and feedback on the other chapters.

Hannah Filmer
May 2012

CHAPTER 1

Literature review

1.1. Introduction

This thesis investigates processing of emotional stimuli by the visual system and how this interacts with visual awareness. Each chapter that follows outlines an experiment (or set of experiments) and introduces the particular issues and background relevant at that point. This chapter outlines previous research that is relevant to the whole thesis, including work on: visual processing, visual awareness & processing emotions. It covers research from a variety of methodologies, including: behavioural, TMS, EEG, MEG, fMRI, patient studies and animal research.

The first part of the review concentrates on visual processing and awareness, in particular focusing on V1's role in awareness, and how this can be disrupted. The rest of the review covers theories of visual awareness, residual processing ability and behavioural effects when awareness is disrupted, the impact of emotional stimuli on behaviour, the problem of measuring awareness, the processing of emotions in the brain, the possibility of a dual route to the processing of emotion beyond sensory systems, and the separation between awareness and attention. Chapter 2 will consider TMS as a methodology in more detail, as TMS research has given important insights into visual processing and awareness, and it is used in experiments reported in subsequent chapters (Chapters 3-4).

1.2. Visual awareness

In 1974, Humphrey published a case study based on a rhesus monkey called Helen. Helen had the whole of the striate cortex removed, and was studied carefully for 8 years. Following the removal, it was initially clear that Helen's vision was severely impaired. Indeed, she appeared to be, unsurprisingly perhaps, blind. However, over the 8 years she was tested, it became apparent that this was not the case - Helen showed increasing evidence of having retained or reacquired some visual capacity. Whilst her vision did not return to normal, it did reach the point where she could accurately move around her environment, and find/reach for small objects. This ability suggested that although striate cortex plays an important role in visual processing, some degree of visual processing could still be achieved if striate cortex is damaged/removed. This is also supported by work by Schneider (1969) who reported that hamsters with the striate cortex removed can orientate to their surroundings but are poor at pattern recognition. Whilst not identical to the findings of Humphrey (1974), this also shows the duality of deficit and remaining ability following the ablation of striate cortex.

In the same year, Weiskrantz, Warrington, Sanders & Marshall (1974), published a paper on a patient who had part of the striate cortex removed following persistent, severe headaches and the detection of an abnormality within part of the striate cortex. The removal resulted in an area of the visual field in which the patient was blind – he could not report the presence of stimuli. However, when presented with a choice between alternative stimuli (for example, vertical versus horizontal stripes), the patient could guess significantly above chance what he had been shown in the blind part of the visual field. This counter-intuitive result is commonly referred to as ‘blindsight’, and has provided substantial evidence that V1 plays a crucial role in visual awareness.

Research with blindsight patients has now been conducted with a whole host of different stimuli, including those displaying emotions. When images are presented in the blind part of the visual field, patients can discriminate between emotions in body posture images (de Gelder & Hadjikhani, 2006), and emotions in faces (de Gelder, Vroomen, Pourtois & Weiskrantz, 1999), and can even be fear conditioned to simple line drawings paired with a small shock (Hamm, Weike, Schupp et al., 2003).

From this research on blindsight, it appears that V1 is necessary for visual awareness but is not crucial for unconscious processing of at least some properties of a visual stimulus, including emotional expression and salience. In addition to studying patients who have damage to V1 it is possible to use Transcranial Magnetic Stimulation (TMS) to disrupt processing in healthy participants. This carries the advantage of allowing the study of the timing of effects in the visual system, as disruption with TMS can be highly temporally specific. Amassian, Cracco, Maccabee et al. (1989) were the first to find that at a specific point in time after a stimuli has been presented (80-100 ms after stimulus onset) a single pulse of TMS could prevent a participant from being able to name a trigram, and participants reported seeing nothing or only a blur. Hence V1 is important for visual awareness, and particularly so 80-100 ms after a stimulus is presented.

Later studies used a wider range of pulse onset times, and implied that TMS can disrupt processing of letters within two stimulus onset asynchrony (SOA) windows

– an early one, estimated to be somewhere between 10-60 ms SOA, which tends to be less robust (present in fewer participants), and a later window, estimated around 100-140 ms SOA (Couthout, Uttl, Walsh, Hallett & Cowey, 1999; Couthout, Uttl, Ziemann, Cowey & Hallett, 1999; Kammer, 2007). It has been shown that these suppression effects are not caused by eye blinks from the TMS pulse (Couthout, Uttl, Juan, Hallett & Cowey, 2000). The fact that V1 is involved both very early, and in a later time window, strongly suggests that V1 plays an important and dynamic role in visual awareness.

Research has begun to examine possible residual processing of emotional stimuli during disruption of awareness in normal participants like those found in blindsight patients. Jolij & Lamme (2005) applied TMS to V1, and compared ability to locate a schematic face (a "happy" or "sad" emoticon) among four otherwise neutral faces arranged in a square to the ability to discriminate between the valence of the discrepant emotion (happy or sad). They found that when TMS was applied to V1 at around 110 ms after stimulus onset, participants were no better than chance when trying to locate the emotional face to the left or right of a central fixation point, but could discriminate the emotion significantly above chance (the difference between location and emotion performance was also significant). As no control condition (with ineffectual TMS) was used, it is not possible to say whether this result was due to the impact of TMS to V1 at 110 ms, or simply the generalised effect of a pulse (and not specifically disruption to V1). For example, the effect could be due to the specific timing of the pulse driven by the sound and tactile sensation (or a general cortical stimulation effect). This seems relatively unlikely, but control stimulation conditions (or in the very least, sham stimulation) are experimentally good practice to rule out arguments of this nature. The study by Jolij & Lamme (2005), as well as the previous noted TMS experiments, used a round coil for stimulation that gives quite a large area of stimulation that may not be restricted to V1. However, the results do suggest that similar sparing of emotion processing to that seen in blindsight patients can be generated in healthy participants with TMS, and V1's critical role in processing of non-emotional stimuli appears to be at around 110 ms. I follow this experiment up in the experiments reported in Chapters 3 and 4.

Another way to disrupt processing in healthy participants is with visual masking. A large amount of research into subliminal processing has been conducted using masking, in particular backwards masking. This is a method whereby a mask image follows a target very quickly (within tens of milliseconds), leading to interference with processing and/or visibility of the preceding target (hence "backwards"). Backwards masking has long been available to anyone who can control the timing of successively displays precisely, so there is a much larger literature on its effects than on effects of TMS.

A number of different types of mask image can be used in a backward masking paradigm. These include: metacontrast, pattern masks, noise masks and light flashes (Breitmeyer & Öğmen, 2006). Metaconstrast masking involves presenting a spatially adjacent image following the target, usually contrasting the target in colour/lightness. Noise masking uses random noise as the mask image, whereby the constituents of the mask do not relate to the target. Pattern masks are composed of components of the target structure/shape – i.e. similar contours, blobs, corners. Light flash techniques involve a bright flash of light following the target image. These varying mask types can also be used with a forward mask, where the mask precedes the target. Both can be used in conjunction to maximise interference, which is known as “sandwich” masking.

The level of disruption a mask can have will vary on a number of factors. One of these is the time separating the target and the mask(s). A common technique is to vary the interval between stimulus and mask onset and plot the psychometric masking function relating target discrimination performance to the interval. Backwards masking with a SOA of variable length can produce one of three main types of masking function. The first is type A, and involves maximal disruption at the shortest SOA, with increasing visibility as SOA increases. The next is type B, which represents a U shape, with little disruption at the shortest SOA, increasing disruption as SOA increases, then decreasing again (typical maximal disruption around 40 ms SOA). The third is multimodal, with several points of disruption (Breitmeyer & Öğmen, 2006). The type of masking function achieved is closely linked to the interaction of mask and target energy (i.e. strength), along with the type of mask image used (Breitmeyer & Öğmen, 2006).

Masking functions appear to be related to the speed of processing, and how this is affected by stimuli properties and masking. As an example, Breitmeyer, Kafaligönül, Öğmen et al. (2006) found different masking functions for different tasks based on the same stimuli. The target stimuli (circles) were metacontrast masked and participants had to discriminate (2AFC) the contour (complete circle or a slight flattening at the top) or the contrast (lightness). Forced choice contour discrimination showed type A masking, with maximal suppression around 10 ms, whereas subjective contrast discrimination showed type B masking with maximal suppression around 40 ms. They concluded that the differences here reflect different speeds (and possibly pathways) involved in contour and surface-brightness processing. The findings demonstrate that the masking function achieved can be useful as a way of assessing visual processing, and that differences in disruption can be evident for different tasks with the same stimuli.

Marcel (1983) used backwards masking, with target stimuli of words and a pattern mask. In some experiments properties of the target stimuli were discriminated by choosing between two probe stimuli – which was more similar in terms of orthographic, or semantic, properties. In other experiments a subsequent target stimulus was presented and responded to, and Marcel measured the effect of the initial (masked) target stimulus on responses to the subsequent (unmasked) target stimulus. If the masked stimulus influences (e.g. speeds up or slows down) responses to the second stimulus, the effect is generally referred to as 'priming.' Priming experiments will be covered in more detail in subsequent sections of this chapter. For now, the experiments of Marcel (1983) and the critique that followed will be considered in more detail.

Through a series of experiments with the SOA between target and mask varied, Marcel established a range of SOAs at which a backward mask disrupted awareness of the target leaving participants unable to report whether a stimulus was present. However, judgments based on other properties of the words (e.g. basic visual discrimination, semantics) at these SOAs were relatively intact when participants were forced to choose. This suggests that processing of the target was mostly disrupted in terms of access and awareness, with comparatively little disruption to basic visual and semantic processing. Findings such as these are in line with the results found by Weiskrantz and others with blindsight patients (Weiskrantz et al., 1974; de Gelder, Vrooman, Pourtois & Weiskrantz, 1999; de Gelder & Hadjikhani, 2006; Hamm et al., 2003).

The pioneering experiments of Marcel (1983) and others were subjected to a barrage of methodological critiques, especially by Holender (1986). In their recent review of the history of masking studies of subliminal processing, Kouider & Dehaene (2007) summarise the main objections to the Marcel (1983) study as follows. The visibility test occurred before the priming or discrimination measurement, leaving open the possibility that visibility increased with experience of the stimuli. The visibility estimate was also based on a small amount of data (no more than 20 trials), leading to questions about the measure's reliability. The conditions and general visual experience of the visibility test and priming/discrimination were different (slowly decreasing SOA for the visibility blocks, more consistency in the priming blocks) which could lead to differences in phenomenal experience and dark adaption making the two difficult to compare. None of the target stimuli used in the priming/discrimination part were present in the visibility test, leading to the possibility that the target could affect the visibility of the prime (e.g. through a semantic relatedness effect). Since the work of Marcel (1983) there have been numerous experiments using masked stimuli as primes that have (largely) tried to avoid these problems. Measuring discrimination performance for masked stimuli is a less used paradigm. They, along with general issues surrounding the measurement of awareness, will be covered in more detail in subsequent sections (subliminal response priming and current controversies, the problem of measuring awareness).

The above-summarised research shows comparable effects of TMS to V1, visual masking, and damage to V1 on awareness of visual stimuli.. This leads to the important question of how these methods work to disrupt awareness, and whether there is any evidence that all three disrupt the same underlying neural network/process. The next section will summarise the main theories of visual awareness, and how it is that damage to V1, TMS to V1 at particular times, and visual masking all lead to disrupted awareness and some residual discrimination ability.

1.3. Theories of visual awareness and its disruption

There have been three main (neurologically based) theories of visual awareness. The first highlights the different *cortical routes* from visual processing (Milner & Goodale, 1992), the second focuses on the role of *recurrent processing* in the visual system (Lamme, 2001), and the third suggests the *amount* of processing is important (Kouider & Dehaene, 2007). All of these theories will now be discussed in light of recent research. They are not necessarily mutually exclusive, but it is the recurrent processing model that appears to have the greatest support from patient, EEG, TMS studies and behavioural effects such as masking.

Grounded in patient’s data, Milner & Goodale (1992) put forward a dual route theory of visual processing highlighting the distinction between the dorsal and ventral streams. The idea of there being two streams to visual processing was first put forward by Ungerleider & Mishkin (1982), who proposed a ‘where’ dorsal pathway, and a ‘what’ ventral pathway. Milner & Goodale's (1992) theory was based on patients suffering from visual agnosia (in particular, D.F., who incurred brain damage following carbon dioxide processing) and showing a profound deficit for recognising and discriminating between even simple shapes. Despite this deficit, when asked to grasp an object in front of her, the patient D.F. could post a card through a small slot and adjust her grasp depending on the size of the object. This led to the theory that the ventral stream, the ‘what’ processing stream (starting at V1 and progressing through to the temporal lobe), was damaged and important for visual awareness and basic discriminations. However, the dorsal stream (also starting at V1 and progressing through MT/V5 and into the parietal lobe) could process visual information for action relatively independent of the ventral stream, and in patients such as D.F. could be sufficient for action outside of awareness. The location of the damage was later confirmed as being in the ventral stream of visual processing (James, Culham, Humphrey, Milner & Goodale, 2003).

Milner & Goodale’s (1992) theory certainly accounts for visual agnosia patients like DF, and builds on the pre-existing evidence of two streams of visual processing in the cortex. However, disorders of awareness are possible if part of the dorsal stream is damaged. Neglect patients who have suffered damage to parietal cortex are evidence to this, as they show reduced awareness linked to areas of space (Vuilleumier, Armony, Clarke et al., 2002). A counter argument here would be that neglect is a disorder of attention, which is a separate (although related) process (see a later section, *A brief word about visual attention*, for some discussion of the differences). However, the findings of Jolij & Lamme (2005, discussed above) show disruption to ability to locate a face, with relative intact discriminations of the properties of faces – essentially the opposite to what would be predicted if the ventral/dorsal, aware/unaware distinction holds true. The idea of awareness being based in the ventral stream (or indeed, any other specific area of the brain) also does not give an account of how visual masking disrupts awareness. One can postulate that a mask may mostly affect processing in the ventral stream, but this does not account for how exactly masking works, i.e. how masking effects the processing of a stimulus.

Further evidence against the idea that visual awareness is based in the ventral stream has recently come to light. Tapia & Breitmeyer (2011) ran a response priming experiment using masked and unmasked arrows of varying contrast. Response priming paradigms present two stimuli in rapid succession, with participants responding to the latter of the two. The stimuli can be congruent in the response they would require –i.e. both require a left button response – or incongruent – i.e. the first stimulus would require a left response, but the second requires a right response. The effect of the first stimulus on the response to the second can be measured in terms of mean reaction times for congruent and incongruent conditions. Subliminal response priming is the same, with the addition of a mask following the first stimulus, leaving the first stimulus failing to reach awareness. Tapia & Breitmeyer (2011) found priming increased rapidly then levelled off with increasing contrast when the arrows were not masked (and hence were visible), whereas there was a more steady increase in priming with increasing contrast for masked (subliminal) primes, and this increase in priming was spread over a greater range of contrasts. Properties of contrast-response functions for magnocellular (predominant input to the dorsal processing stream) and parvocellular (predominant input to the ventral processing stream) channels were compared to the results of the response priming experiment. These properties were known from previous research measuring the response of ganglion cells that projected to either magnocellular or parvocellular cells in the Macaque monkey (Croner & Kaplan, 1995). This comparison showed the priming effects found with the unmasked primes were best modelled with properties of the magnocellular channels and the effects for masked stimuli by properties of parvocellular channels. In other words, properties of the inputs to the dorsal stream can be more closely linked to visual awareness than the properties of the inputs to the ventral stream. These results indicate, if anything, a complete reversal of the predictions from Milner & Goodales's theory (1992). They are also in line with Jolij & Lamme's (2005) finding that TMS to V1 can disrupt ability to locate a face (the dorsal stream) but not to discriminate which emotion is present (the ventral stream).

An alternative theory has been put forward that phenomenal visual awareness is a result of a series of recurrent loops throughout the brain, that all feed back into early visual processing areas (Lamme, 2001). When backwards masking is used, processing of this feedback is disrupted, reducing/eliminating awareness for the initial target image. This theory, and this type of activity, is frequently referred to as feedback, recurrent or re-entrant processing. Breitmeyer (2007) also suggested that backwards masking is effective through disrupting feedback processes in V1 (as opposed to forwards masking, which more likely interferes with the initial processing).

To be able to understand and assess the theory of feedback processing in visual awareness, it is first necessary to explore the connectivity of the visual system in more detail. The existence of feedforward, feedback, and horizontal connection in the visual system has been established in animal studies, and was thoroughly reviewed by Felleman & Van Essen in 1991, and by Salin & Bullier in 1995. Both note that most areas of the visual cortex in the Macaque monkey that are connected have connections that are reciprocal – they have the potential of functioning in both directions. Another point, which is clear from Felleman & Van Essen's review, is that there really are a lot of connections between visual areas. Figure 1.1 shows a summary of the connectivity of visual cortex with other cortical areas. V1 alone appears to be connected to at least 8 other visual areas including many areas of the occipital cortex, and several in parietal cortex. (It is worth noting here that there is also evidence that V1 is directly connected to a number of subcortical structures, in particular the amygdala (Amaral & Freese, 2005)). In return, V1 appears to have projections from at least 9 different areas (again from the occipital and parietal cortex: Felleman & Van Essen, 1991). It would seem V1 receives more inputs from "above" than it sends outputs upwards. In addition, Salin & Bullier (1995) note that there are more neurons involved in feedback than feedforward processing between V1 and V4, and claim that feedback connections are more numerous than feedforward connections in ventral occipitotemporal pathways in general. Feedback, then, appears to be a well established and, in terms of numbers of connections, a high proportion of processing in the visual system.

Concerning feed-forward processing, Lamme & Roelfsema (2000) conducted a meta-analysis of 48 studies of the timing of visual activation in the Macaque monkey. The earliest activity was in V1 and MT (about 40 ms after stimulus onset), followed by V3, SMA, FEF and MST at about 50 ms, and by 65 ms after stimulus onset many areas of the brain show activity including V2, V4 and areas of temporal cortex. This is clearly demonstrating that initial cortical activity relating to visual processing in the monkey is fast, and spreads rapidly to numerous areas of the cortex. Feed forward processing is thought to be hierarchical, starting at low levels of processing (e.g. processing basic contours) and moving (rapidly) to higher levels of processing (e.g. for faces or objects) (Lamme & Roelfsema, 2000).

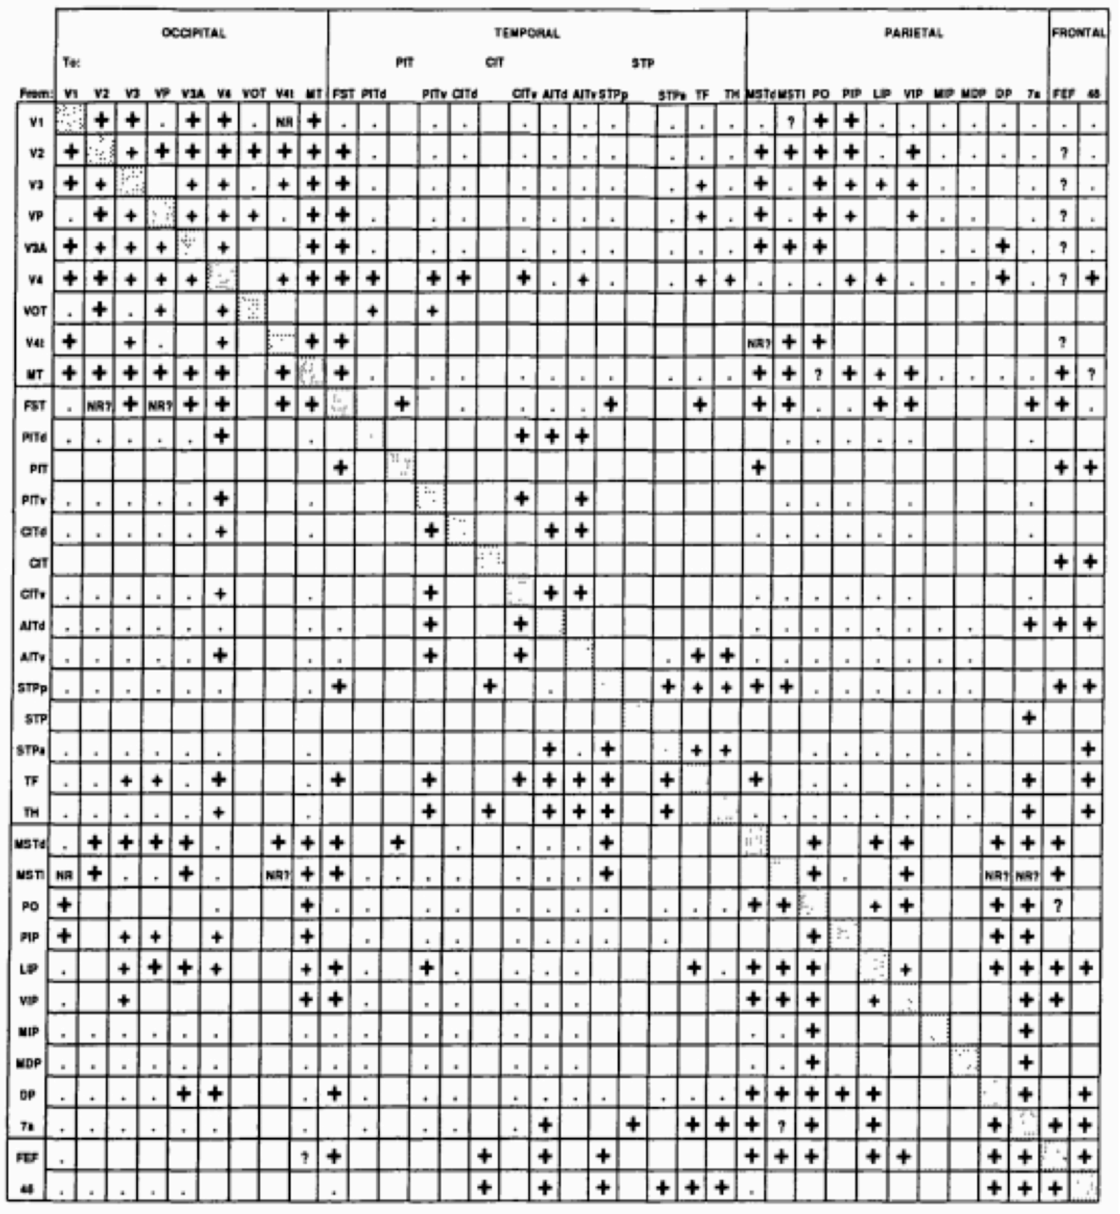

Figure 1.1: Connectivity matrix summarising the interconnections within visual cortex and between visual cortex and other cortical areas. Reproduced from Felleman & Van Essen, 1991.

Salin & Bullier (1995) note that while feedforward processing has well-defined receptive fields – each cell processes a small area of the visual field — feedback appears more diffuse, without clear receptive fields: individual cells respond to stimulation in a range of different locations. The possibility that feedforward may be linked more directly to receptive fields than feedback has also been noted by Lamme, Super & Spekreijse (1998) who cite research from the figure-ground literature that has measured single cell responses to various stimuli – those displaying a texture with a ‘figure’ (e.g. a texture defined square), and those displaying a texture with no figure. In the initial feedforward sweep, it appears that processing relates to a receptive field and is relatively unaffected by activity outside of this area. However, feedback seems to behave differently. For example, the feedback relating to activity based on the centre of a figure is affected, in the feedback sweep, by whether the ‘figure’ is actually a figure or a constant pattern across the screen. That is, feedback processing implements contextual modulation – information relating to activity outside of a specific receptive field.

Lamme, Super & Spekreijse (1998) & Lamme (2001) point to evidence that a cell’s response is moderated by stimuli outside its perceptual field, but only after around 100 ms from stimulus onset. That is, a cell with a receptive field that lies in an area of textured stripes pointing in one direction responds differently from 100 ms after stimulus onset depending on what is happening outside this area i.e. whether the texture is in the same direction outside the receptive field. This could be at least part of the purpose of feedback – initially processing is limited to the receptive field of the cell(s) being recorded, but after information begins to return in the form of integrated feedback the activity is now moderated by events outside the receptive field.

To try to understand feedback processing in more detail, it will be helpful to briefly consider EEG/MEG studies and reviews conducted since those of Felleman & Van Essen (1991), and Salin & Bullier (1995). The huge benefit of using these methods is their temporal specificity, which allows researchers to study activity occurring milliseconds after a subliminal image is presented. One question that can be looked at with EEG is how a backwards mask disrupts processing of a target.

Fahrenfort, Scholte & Lamme (2007) investigated visual processing using stimuli consisting of texture-defined squares and EEG. They reduced awareness of a target using backwards masking, and found associated disruption of signals in the ERP. The signal thought to be associated with the figure was isolated by computing a difference wave between masked and unmasked figures, showed differences emerged between masked and unmasked conditions at 109-141 ms after target onset. This suggests that masking has very little impact in the earliest stages of processing (thought to represent feedforward processing) but does so later on, in a time-window appropriate for feedback. If these assumptions were correct, then it would appear that a mask interrupts the later, feedback processing whilst leaving the feedforward processing relatively intact.

In a subsequent study, Fahrenfort, Scholte & Lamme (2008) correlated behavioural measures with the ERP difference signal associated with figure versus no figure followed by a mask. Only the 'later' activity thought to reflect 'recurrent' processing (activity only present when a stimulus was present), occurring at around 160 ms post target onset, correlated with behavioural performance (detection of figure), with earlier signals relating to the presence of a figure (peaking around 121 ms post target onset) not correlating with performance. This study suggested a slightly later time window than the previous study (Fahrenfort, Scholte & Lamme, 2007). But the finding of two stages of processing, with only one associated with detection of the figure, again suggests two levels/stages of processing consistent with the hypothesis of recurrent processing leading to awareness.

VanRullen & Thorpe (2001) used a natural scene to study the timing of processing related to behaviour/perception in the absence of masking. They asked participants to complete a go/no-go task to natural scenes showing animals or tractors presented for 20 ms. Participants were asked to respond to only one of the two types of images, and alternated between the two types (between blocks). They found that task-independent (differences in the ERPs between the animal and tractor images, regardless of which task was active) processing occurred at 75-80 ms after stimuli onset, with task-relevant (related to the task in question, regardless of specific stimuli type, and correlated with behaviour) processing occurring at 150 ms. Whilst it is not perhaps surprising that stages of processing would begin with basic processing that does not correlate with behaviour, and move to more meaningful processing, it is possible from the timings that the initial processing was feed-forward, and processing only reflected the behavioural response following recurrent feedback through the brain. Indeed, this point was put forward by VanRullen & Koch (2003), who suggested feed-forward was enough to influence behaviour, but feedback 'reverberates' and impacts behaviour to a greater extent. These results indicate that the 'feedback' processing correlating to behavioural performance for masked stimuli found by Fahrenforte, Scholte & Lamme (2008) can be echoed using natural scenes without masking. Again, though, it must be noted that there is no specific evidence from this that the 'later' processing was recurrent activity, and it is somewhat of leap from the data.

If V1 is feeding information to, and receiving information back from, other visual areas then we should be able to find evidence of information passing between two areas, at least once in either direction. This would provide more substantial evidence of recurrent activity leading to visual awareness than the above-summarised EEG findings. Juan & Walsh (2003) studied V1's potential interactive relationship with V5. They applied double pulse TMS to V1 and V5 whilst the participants responded to a display of dots, indicating whether they were moving or stationary. When TMS was applied to V1 at 60 and 80 ms, there was disruption to performance. There was another window of disruption with TMS at 100 and 120 ms. In addition, stimulation at 80 and 100 ms (the in-between timing) reduced performance when TMS was applied to V5. So disruption starts off selective to V1, then V5, and then V1 again. This is strong evidence for a feedback model of visual processing, as it not only highlights two time periods in which V1 is involved, but also gives evidence of specific involvement of another brain area in between.

A third theory of visual awareness, mentioned at the start of this section, must also be noted. This was described by Kouider & Dehaene (2007), and suggests that for a stimulus to reach awareness, it requires sufficient activity to activate 'global reverberation.' These reverberations are defined as bi-directional, self sustained loops across the brain (i.e. including sensory areas, and sensory non-specific areas such as the frontal cortex). This theory shares considerable ground with the feedback theory of visual awareness where the integration and communication between brain areas is required for awareness, but describes a broader network of areas that 'reverberate,' and the trigger for this reverberation as being a sufficient level of processing. Sufficient processing can be prevented by visual masking, and prevent awareness of the stimulus being reached. This is, perhaps, a subtle difference to the feedback theory, whereby methods to disrupt awareness (such as TMS and backwards masking) do so not by preventing the initiation of this 'reverberation,' but by directly interfering with the feedback itself. The two possibilities are not, of course, mutually exclusive – processing could be disrupted before or after the initiation of a feedback stage of processing, and it might be that the trigger for reverberation requires a degree of initial feedback processing in the visual system (in particular, feedback to V1). The difference, in essence, is more about the processing level(s) that need to be reached for awareness to occur: does it start from the early feedback loops within visual cortex, or does it require a larger network of the brain? Perhaps worth a note here is that treating conscious awareness, as measured in priming and masking experiments, as an all or nothing concept may not be accurate. There may well be *levels* of awareness of a stimulus, and these levels could correspond to the *amount* of feedback, or reverberation, present.

The global reverberation theory can be used to explain behavioural effects of subliminal stimuli. When awareness of the stimulus is not reached, i.e. the reverberation was prevented/minimised by a backwards mask, processing of the stimulus may still occur in many areas. This processing, whilst not enough to trigger the reverberation, may be sufficient for perceptual and semantic effects on behaviour. The reverberation theory is also mostly concerned with cortical areas, leaving open the possibility that disruption to this reverberation may have little or no impact on relevant processing in subcortical structures.

Gaillard, Dehaene, Adam et al. (2009) measured activity in the visual system using intracranial recording in 10 patients, with recordings taken from multiple sites (during surgery for medical reasons) in the occipital, temporal and frontal lobes. Activity in the measured populations of cells was compared when patients were shown masked or unmasked words to study the differences between the neural correlates of conscious and unconscious visual processing. They discovered that unconscious conditions were associated with a fast, dissipating wave of activity (the feed-forward processing) which was both brief and largely in the gamma band of activity. When, in contrast, stimuli were consciously perceived, there were a convergence of different markers such as higher spectral power in the gamma band of activity, and voltage changes in the prefrontal cortex and increases in beta range phase synchronies. This could be interpreted as evidence of Kouider & Dehaene’s (2007) reverberation.

The fact that conscious awareness appears to be linked to many processing changes suggests that awareness is borne out of complex interactions in the brain. However, the theory of recurrent processing now has considerable weight with mounting support from EEG and TMS studies. Whilst the neural correlate of visual awareness is not yet pinned down, it seems sensible to conclude that awareness results from complex interactions between brain regions and most likely involves feedback processes.

1.4. Subliminal response priming and current controversies

I have briefly mentioned above research that found subliminal (or very low visibility) stimuli can influence behaviour: Weiskrantz, 1974; de Gelder, Vrooman, Pourtois &

Weiskrantz, 1999; de Gelder & Hadjikhani, 2006; Hamm et al., 2003; Kouider & Dehaene, 2007; Marcel, 1983. In this section of the review, I will now discuss these residual behavioural effects in more detail, in particular those found in response priming paradigms. Later in this chapter I will discuss another form of priming – the priming of arousal with emotional stimuli.

In the 1950s a report emerged by James Vicary that briefly presented images (presented for 1 frame) in a cinema encouraging people to buy more popcorn, or drink more Coca-Cola led to an increased in consumption of those products. This finding sparked a considerable amount of anxiety about the possible effects of subliminal images on behaviour. A flurry of research ensued, and was reviewed by Holender in 1986, who concluded that either the studies had failed to show sufficient evidence of any semantic processing of a subliminal image, or the images in question were not convincingly subliminal. The criticism was focused on Marcel (1983) and the many attempts that were made to replicate the findings of Marcel (1983). These attempts were either unsuccessful, or suffered from their own limitations. Holender’s (1986) critique of the Marcel (1983) experiments was discussed earlier in this chapter. The sceptical view taken by Holender (1986) was perhaps appropriate in the 1980s, but since then methodologically superior research has found that subliminal images can be processed semantically, and can influence some behaviours (Kouider and Dehaene, 2007). This literature, and the current points of debate and controversy is the focus of this section.

The term ‘priming’ refers to the ability of a stimulus to influence a response(s) to a subsequent stimulus. In particular, subliminal response priming is where a ‘prime’ stimulus is (subliminally) presented very shortly before a ‘target’ stimulus, to which participants respond. If the prime image specifies the same (congruent) or opposite (incongruent) response as the target then this may speed up or slow down RTs. For example, many priming studies have used stimuli consisting of arrows pointing either left or right to test whether a masked arrow (the prime) can influence the RT for a left or right key press to a subsequent unmasked arrow (the target). Vorberg, Mattler, Heinecke, Schmidt & Schwarzback (2003) used stimuli of arrows and found that an arrow prime made undetectable by masking (i.e. d' of zero when its direction had to be discriminated) could nevertheless facilitate RT to a subsequent target with the same response. That is, the detection of a prime, and the influence of a prime seem to have a degree of independence.

Schlaghecken & Eimer (2004) also found RT priming effects with arrow stimuli when the prime was subliminally presented. They also included 'free choice trials', where there was no right or wrong response and participants were instructed to make a random choice. On these free choice trials, participants were significantly more likely to choose a response that was congruent to the prime. A similar effect was found by Kiesel, Wagener, Kunde, et al. (2006) using numbers, where participants had to classify whether a number was higher or lower than 5. In this experiment, RT priming effects were found along with an impact on free choice trials. The numbers used as primes and targets were semantically the same, but in different format (i.e. "7" versus "seven") indicating a degree of subliminal semantic priming as the prime and target stimuli share only semantic (and not basic visual) properties.

Kiesel, Kunde, Pohl, Berner & Hoffmann (2009) manipulated the familiarity of the prime to study its impact on behaviour. They found that long term chess players showed priming and choice effects following a complex, familiar chess move. If the move was unfamiliar, no priming effects were present. This is introducing the prospect of personal saliency/familiarity being important for unconscious processing (and measurable behavioural effects), at least in terms of relatively complex stimuli. However, this study reports no measure of awareness being used, so the actual level of awareness (and thus whether these results are subliminal or supraliminal) is questionable.

Debate still exists in the literature as to the nature and cause of subliminal response priming. The main points of debate will be outlined here. In a subsequent section I will review the problem of measuring awareness to determine that stimuli were presented subliminally.

An important debate in the recent priming literature concerns whether a prime can influence behaviour if it has not been (consciously) seen before the present trial. Many studies use the same images repeatedly as both primes and targets (e.g. Kiesel et al., 2009; Schlaghecken & Eimer, 2004), with the result that the prime has been seen consciously at some point. It has been argued that subliminal priming effects are actually due to partial awareness of prime stimuli where familiar primes can be reconstructed from small amounts of (aware) information of the prime (Kouider & Dupoux, 2004). For example, detection of a slight curve around the mouth area can be interpreted, if the curve is upwards, of representing a happy expression. This can, in theory, be achieved with no (or very little) awareness of the rest of the stimulus –

hence the term ‘partial awareness.’ There is one main intuitive argument against this: the time pressure present in response priming paradigms does not seem to allow for some form of elaborate reconstruction of the prime image (an image that participants are frequently unaware is even present). Another issue raised by the use of the same stimuli as targets and primes is whether the priming effects arise from the subliminal prime being processed semantically or instead affecting behaviour due to learned stimulus –response (S-R) mappings. If the prime was previously presented as a target, participants may have developed simple response associations (S-R mappings). When the same stimulus is then presented as a prime, processing is sufficient to activate the S-R association but not necessarily semantic processing.

Both the partial awareness hypothesis and the issue of whether primes are processed semantically or as S-R responses can be refuted if priming can be found in the absence of any prior knowledge of the prime stimulus or use of it as a target stimulus. Many studies use the same stimuli as prime and target images leading to a considerable degree of familiarity of the stimulus, but if different items are used as targets and subliminal primes and priming is found, it would argue against partial awareness and suggest semantic processing. Abrams & Grinspan (2006) found subliminal effects with ‘novel’ primes (not seen consciously at any point). They used a pleasant/unpleasant discrimination task, with (different) words for the primes and targets. They do not report the usual RT effect of congruence but its effect on d' (estimated from error rates for the two possible prime types). The priming effect reported only occurred following practice at the task, and as pointed out by Kouider & Dupoux (2007) this and the use of a difficult discrimination allows for an interpretation in terms of reconstruction and partial awareness. Fully automatic subliminal priming intuitively should not require practice if it is hard wired in the brain. Another possibility is that people start off responding slower to the targets as the task, buttons and words are ‘unfamiliar’ and this slow and cautious response strategy initially masks a priming effect. This is not possible to rule out, as there was not a condition where the target task was practiced without any primes, and overall response speeds were not compared.

Whether a subliminal, target independent prime can influence rapid responses is still unclear. Many studies do not actually achieve full suppression of awareness, although effects do not correlate with visibility measures and so indicate independence from awareness. A slightly different approach to the question of semantic processing and prime awareness is to look more holistically at the problem – there are undoubtedly many possible sources of priming effects, and these may interact within each priming experiment depending on the design. A meta-analysis was run by Van den Bussche, Van den Noortgate & Reynvoet (2009) of 24 studies investigating subliminal, semantic categorisation, priming. Overall, the analysis suggested that non-semantic explanations (such as S-R mappings) could not explain all of the priming effects reported. However, non-semantic factors did *enhance* the priming effect. It is possible, then, that true subliminal semantic priming is possible, but the effects are 'boosted' by availability of other relevant information – e.g. learnt S-R mappings. This leaves open the question of the strength of the semantic priming effect in the absence of other factors – partial awareness, learnt S-R mappings – but certainly highlights that subliminal, semantic, processing is at least one possible contributor to priming effects.

A recently published set of studies (Eckstein & Henson, 2012) investigated the conditions under which S-R mappings may be generated to prime stimuli. They manipulated the primes (faces) so they occurred only as primes, or as both primes and targets. Under prime-only prime face conditions, the authors consistently found (for fame and gender judgments) congruence effects were absent. If the prime also occurred as a target, congruency effects were found, but only if the trials were mixed – if the prime stimuli occurred as targets throughout the experiment. If the prime stimuli occurred as targets in a previous session, no congruence effects were found. These findings hint that congruence effects with primes that do not occur as targets may not be measurable, and that the development of S-R mappings requires frequent exposure to the prime stimuli as targets.

If there are subliminal, unconscious effects then it may be intuitive (and certainly appears from the research currently reviewed) that this impact is outside of our 'control', as well as awareness. Kiesel, Kunde & Hoffmann (2007) used a task switching paradigm where participants switched between discriminating target letters as consonants or vowels, or numbers as higher or lower than 5. The primes could be congruent or incongruent to the response option, and relevant to the active task or not. They found evidence that irrelevant primes (e.g. primes of the inactive task) can influence behaviour, suggesting that processing of the prime can be automatic. This provides an example of a case where primes affect responses to the subsequent stimuli on the basis of S-R mappings, not semantic processing, as there is no semantic overlap between the irrelevant primes and the targets.

In contrast to the above research, Eckstein & Perrig (2007) asked participants to practice two different tasks (animate/inanimate, or positive/negative decision on words) then complete only one of the two tasks. They found that the task given influenced the type of priming when the targets and the primes were otherwise the same. That is, despite practicing responding to another task, participants were not then influenced by these primes – there was no evidence of automatic response priming. This is evidence that there is some form of top down control that can influence subliminal effects. Eckstein & Perrig (2007) tested the visibility of their primes, and found them to be above chance level. However, the priming effect was not correlated with performance so they argued that priming effects and visibility are not directly related.

The nature of processing of prime images was further explored by Tapia, Brietmeyer & Shooner (2010) who used shapes that varied in form and colour, and were metacontrast-masked. Properties of the prime stimuli could be fully congruent (same colour and form as the target), incongruent (colour and form different to the target), or partially incongruent (one dimension congruent, the other incongruent). Participants responded to both dimensions at the same time, or just one of the two dimensions. When the stimuli were presented unmasked, a stimulus behaved as if incongruent for the dual-dimension task if either (or both) feature was incongruent. For the masked primes and the dual-dimension task, stimuli that were partially incongruent were ‘neutral’ in their priming effect. If only one dimension was responded to, an incongruent feature irrelevant to the task did not impact performance for masked or unmasked primes. This result implied participants were able to focus on one dimension of the prime at a time, and were able to control which dimension was attended when needed. If no specific dimension was useful, the dimension attended appeared to be relatively random. As a result, for the dual-dimension task, on some trials performance would reflect incongruent responses, and on some trials congruent – averaging out to neutral. Processing, by this interpretation, was featural with attention focused on only one feature of the prime.

Whether processing of the prime, and the priming effect, are truly automatic remains unclear, however the main studies quoted on this topic used different structures: Kiesel, Kunde & Hoffmann (2007) used a regularly changing task, whereas in Eckstein & Perrig (2007) two tasks were initially practiced but only one was completed. Tapia, Brietmeyer & Shooner (2010) asked participants to complete all of the ‘tasks’ (attend both features, or one of the two only), but this changed relatively infrequently. It might be sensible to conclude here that irrelevant features can be ignored for some tasks, but can interfere in others, with interference more likely if there is a greater amount of switching between tasks (and regular ‘activation’ of the irrelevant S-R mappings).

There is some evidence that even very low visibility primes are processed semantically rather than merely activating a learned response category. Kiesel, Kunde, Pohl & Hoffmann (2006) asked participants to classify words (nouns) as large or small objects, with the target words preceded by a congruent or incongruent prime word. Primes could either be from the target set (and seen consciously) or presented only as a subliminal prime. They also manipulated the size of the target set, which consisted of either 40 or 4 words. They found priming effects with familiar (consciously seen) primes for both large and small target sets, but novel primes only produced priming for the large target set. The experiment found higher than ideal d' values for visibility of the primes (above chance), however visibility did not correlate with the priming effects. The main results imply that the number of target images determined the presence of a priming effect. This may be due to a different relation being established between the primes and targets, i.e. S-R response mappings when the target set is small, and a semantic approach when the target set is large. This study provides further evidence that priming can occur when the primes are target independent, although this may only be under certain conditions and it would appear that properties of the target stimuli set are important.

The notion of there being two different types of priming is supported by recent work by Finkbeiner & Friedman (2011) who asked participants to classify photos as animals or people, using masked primes that were congruent or incongruent, and target-independent or presented as targets and primes through the experiment. Participants responded via reaching to the left or the right depending on the response to be made (e.g. reach left for person). They found two different patterns of reaching; reaching for repeated incongruent primes gave a larger and earlier path offset (i.e. deviation from the optimal reaching path) than target-independent incongruent primes, although effects of congruency were present for both. They suggested that the repeated primes showed a peripheral (perceptual) priming effect, and novel primes a central (semantic) effect. The data did, in the very least, suggest different patterns of priming for stimuli with different levels of experience (i.e. previously seen consciously or not).

There are still unanswered questions in the priming literature as to what specifically drives priming effects. The evidence for response priming, for top down influences (Eckstein & Perrig, 2007; Tapia, Breitmeyer & Shooner, 2010), influence of processing of the target (Kiesel, Kunde, Pohl & Hoffmann, 2006) and the suggestion that practice may be required when the primes are target-independent (Abrams & Grinspan, 2006) indicate that there are many potential factors. These are points which need to be considered in designing priming experiments, both to control potential confounds as far as possible, and to try and disentangle the source(s) of priming effects.

Another puzzle in the priming literature is that of Negative Compatibility Effects (NCEs). Most of the discussion of priming so far has assumed that a prime that is congruent to a target speeds up reaction times in comparison to incongruent prime-target pairs (a Positive Compatibility Effect – PCE). However, this is not always the case. A NCE refers to an observation of slower responses on congruent than incongruent trials. This has been found in a variety of contexts, but namely seem to be related to the prime visibility and the temporal relationship between the prime and target. When NCEs occur (in more detail) and the impact of these effects on understanding priming, will now be discussed.

Sumner, Tsai and Nachev (2006) manipulated the visibility of a prime arrow where participants responded to a second (always visible) prime arrow. They found NCEs for the ‘invisible’ primes and PCEs for the ‘visible’ primes. This suggests that the direction of priming effect, positive or negative, may be influenced (to some extent) by prime visibility. However, the suggestion that awareness is linked to the priming effect calls into question the validity of using a lack of correlation between awareness and priming as proof that the two are unrelated (e.g. Eckstein & Perrig, 2007; Kiesel et al., 2006). It appears that the issue of independence between priming and awareness may not be a simple one.

Other research has argued that NCEs can be driven by inhibition of motor pathways (Sumner, 2008; Boy & Sumner, 2009; Reith & Huber, 2010; Schlaghecken & Eimer, 2004). This theory suggests that the direction of priming effect is dependent upon the timing of the target following the prime. Initially, the prime may activate motor pathways and, if a target is presented in close proximity, this would lead to a PCE. However, after the initial activation following the prime, inhibition occurs, as a response is not immediately required to the prime. If the target is then presented, the inhibition will actually slow responses relative to an alternative choice. This theory essentially posits that at a certain temporal separation the prime and target are not longer treated as one event, but as separate events, so that any response activated by the (irrelevant) first event is suppressed and becomes more difficult to activate as the response to the target.

There is considerable evidence for this theory from research by Sumner (2008) and Boy & Sumner (2009) showing that a time variation between target and mask (whilst the time between prime and mask is constant, and awareness comparable) can lead to PCEs at short time lags (40 ms), and NCEs at longer intervals (150 ms). With this in mind, it appears that time interval can explain a large portion of the NCE, under some conditions at least. This holds when the response was not a highly compatible left or right response, as in the earlier studies. Boy & Sumner (2009) used an association that was learnt with practice during the experiment, with stimuli of horizontal and vertical line pairs. The prime and the targets did not have obvious responses associated with them prior to the start of the experiment. The trajectory of learning and of priming was compared for a long and short SOA between the prime and the target. Both the NCE at the longer SOA, and the PCE at the shorter SOA developed at a similar rate. Moreover, when the association was reversed halfway through the experiment, both showed the same reversal of priming direction, followed by recovery at a similar rate. Boy & Sumner (2009) concluded that the NCE and PCE are closely linked to behave in a similar fashion.

The data considered so far has focused on stimuli that are mostly (with the exception of Kiesel et al.'s 2009 study using chess) relatively arbitrary and of little social salience. But what happens when you subliminally present an image that may carry more social importance?

1.5. Emotional processing

The priming data reviewed above have provided evidence that RTs to target images can be impacted by subliminal (or, at least, low visibility) prime images. Most of the evidence reviewed used stimuli of neutral emotional valence, and low social saliency (e.g. arrows, numbers, and trigrams). Emotionally valenced stimuli, and social stimuli in general, may generate a greater subliminal impact, and could conceivably result in a greater influence on our behaviour. The focus of this thesis is on whether there is any evidence of processing of emotional stimuli at low levels of visibility above and beyond that of neutral stimuli. The stimuli used in this type of research tend to be quite restricted – faces in the most part, with some use of emotional scenes and body postures. Emotions themselves in this research field also tend to be quite narrow in definition – fear, anger and disgust are the most prevalently used negative emotions. However, it is perhaps worth some consideration of the general theory of emotional processing before turning to the research that has specifically looked at the processing of low visibility emotional stimuli.

Emotions are seen as central to human experience, and the disruption of the experience of emotion in a person can lead to severe social and personal impairments. The value of emotions has been debated, with arguments that emotions have evolved and are comparable, in part at least, for many different species and are not specific to humans (Darwin, 1872). Evidence for emotions being a component of human experience that evolved early in human history comes from the areas of the brain associated with emotional processing – at least some are located in subcortical structures thought to be old in an evolutionary sense and comparable to areas of brains for other animals (LeDoux, 1992).

Paul Ekman has furthered the evolutionary view of emotions by showing that emotional expressions are universally familiar across different cultures including those relatively isolated from globalisation (Ekman & Friesen, 1971). That experiencing emotions, particularly emotions relating to potentially threatening situations, carries evolutionary advantage is certainly not difficult to believe. More generally, that emotions are a central component of society – e.g. happiness, shame, anger, and disgust – makes good sense. They help society develop rules, norms, structure, and perceived fairness.

Another angle to emotional theory is the nature of specific emotions themselves. Some researchers suggest there are discrete emotions, e.g. Ekman (1972) proposed six discrete categories: anger, fear, disgust, sadness, surprise and happiness. Ekman later (1999) updated the discrete categories to include seventeen different emotions. An alternative view to separate emotions is some form of continuum. In 2011, Lovheim came up with a three-dimensional model of emotions that vary along continuums of three different neurotransmitters: dopamine, serotonin and noradrenaline. For example, low serotonin and dopamine, but high noradrenaline leads to distress whereas high serotonin and noradrenaline with low dopamine reflects surprise. Whether as a continuum or as discrete categories, emotions are dynamic and can be expressed in several different forms.

Another dimensional aspect of emotion is that of valence and arousal. Valence refers to positive vs. negative emotions (happiness vs. anger), whereas arousal is the degree of reactiveness to stimuli, including physiological responses like increased heart rate. A lot of research that follows in this chapter, and in this thesis in general, manipulates and focuses on emotional valence, whilst controlling for or ignoring arousal. However, section 1.7 of this chapter will refer back to arousal effects of emotional stimuli, particularly in emotional priming experiments.

It is important to note here that emotions can be discussed from several points of view: the personal experience of emotion, emotional reactions, the communication of emotions, and the perception of emotions displayed by others. It is this latter point that is of primary relevance to this thesis – the perception and processing of the emotion components displayed in emotional expressions from faces and body postures. I will focus on this element in the rest of this section.

Duncan & Barrett (2007) reviewed neurobiological data that suggests emotions heavily influence many areas of cognition (e.g. perception, memory), and many brain regions overlap between the processing of emotions and the processing of cognitions like perception and memory. Duncan & Barrett (2007) use this as evidence that emotion is a form of cognition, and not necessary a separate process. If this is so, it will be difficult, if possible at all, to disentangle emotional processing from other cognitive processes.

Research linking emotion with cognition dates back many decades. In the 1970's in particular the view was expressed that emotion required extensive information processing, and arose from cognition (e.g. Lazarus, Averill, & Optin, 1970). In other words, there is no emotion without cognition. However, Zajonc (1980) caused great controversy by claiming that emotion could be independent, and indeed precede, cognition. The primacy of affect hypothesis, as it was called, led to a wealth of debate and research into emotions and the processing of affect which is, to this day, not yet resolved. The research generated by the ideas of Zajonc (1980) was reviewed by De Houwer & Hermans (2010). One of main critiques of De Houwer & Herman's (2010) review of Zajonc (1980) was that he did not differentiate between 'effortful' cognition, and automatic cognition. In other words, cognition is not always effortful in the way Zajonc (1980) assumed, and so the debate about the link and possible overlap between cognition and emotion remains.

Zajonc's (1980) suggestion that the processing of emotions may be more automatic, and faster, than 'ordinary' cognition is related to the content and over-arching theme of this thesis: is there evidence the processing of emotions is relatively spared from disruption to V1 or disruption to visual awareness compared with processing of neutral stimuli? The two issues are different, but both are essentially asking whether emotional stimuli are treated differently compared with non-emotional stimuli, and whether emotions may be processed in the brain preferentially in terms of speed and resources.

Following on from the brief review of emotion (s), the next section will outline the processing of subliminal and low visibility emotional stimuli, and how subliminal emotional stimuli may impact behaviours.

1.6. Subliminal emotional stimuli and their impact on behaviour

There is emerging evidence that emotional stimuli may have a processing advantage and lower threshold for awareness. Kousta, Vinson & Vigliocco (2009) used a word/non word task (with conscious presentation), and found faster responses to emotional than to control words. There were no specific valence effects, but a general processing speed and accuracy advantage for emotive words. Gaillard, Cull, Naccache et al. (2006) found that when emotive words were compared with neutral words, the threshold measured for awareness was lower if the words were emotional. Both of these studies suggest a processing advantage for emotive words, with some degree of semantic processing occurring sufficiently rapidly to influence awareness. The flip side to the presence of an emotional advantage is that emotional stimuli can reduce performance at other tasks. The presentation of an emotional face, that is task irrelevant, immediately before a visual search array containing letters can disrupt participants' ability to complete the search array (Gupta & Raymond, 2012).

Using fMRI, Lang, Bradley, Fitzsimmons et al. (1998) found there was generally more activity in visual areas to emotional pictures than neutral (stimuli taken from the IAPs). The use of fMRI does not allow any temporal inference to be made, but it seems emotional pictures trigger a greater *amount* of activity in the visual system than neutral pictures.

Emotional stimuli may have speed and visibility advantages over their neutral equivalents. But when stimuli are subliminal, does any of that advantage remain? Murphy & Zajonc (1993) found that using a prime of an emotive face influenced how participants rated an emotionally neutral ideograph presented immediately afterwards. We saw earlier that a non-emotive stimulus can impact free choice on trials where a target has no specific response required (e.g. Schlaghecken & Eimer, 2004; Kiesel et al., 2006), but here the prime did not simply bias choice towards a response-congruent stimulus, it nudged participants' evaluative response to an unrelated target.

Whilst subliminal affective primes may influence an emotional judgment of neutral stimuli, the effects may be task dependent, only influencing responses to emotive decisions and not linked to a general arousal effect (e.g. always faster following a negative prime, regardless of the target discrimination being made). Houwer, Hermans, Rothermund & Wentura (2002) used affective words as both primes and targets, where participants categorised the words as positive or negative. The emotion of the prime led to congruency effects, however there was no impact of emotional primes on a control task that did not require semantic processing. This suggested that the impact of an emotional prime is not one of general arousal, but specific to some form of congruence. Houwer, Hermans, Rothermund & Wentura (2002) conclude that affective priming is a form of response conflict. When the prime emotion is irrelevant to the task, it has very little impact on response times.

Tamietto & de Gelder (2007) used subliminal emotional faces as primes and visible, emotional faces as targets, and found you can prime the recognition of fear but the same is not true for primes with happy expressions. They concluded that negative processing is stronger, more effective and/or pre-tuned when compared to positive emotion. The design used by Tamietto & de Gelder (2007) used very similar stimuli as primes and targets – both were faces. The priming effects from this type of design could reflect basic perceptual priming e.g. of the curve of the mouth that forms the relevant expression. If this were true, the priming effect would not be emotional, per se. Using stimuli that differ perceptually, but convey the same emotion, would avoid this explanation of priming effects. Andrews, Lipp, Mallen & Konig (2011) used primes of emotional faces (happy and angry schematic and photographs of faces) and targets of positive and negative words. When the primes were masked, no significant effect of the prime faces were found on responses to the target words. When unmasked, the primes did influence responses. This study, then, found no significant effects of a subliminal emotional face on performance. This contrast in findings between Tamietto & de Gelder (2007) and Andrews et al. (2011) could reflect the different emotions – Tamietto & de Gelder (2007) only found priming for fear, and Andrews et al. (2011) did not use fearful expressions. The differences could also reflect a limitation to priming, in that priming may only be found when there is some featural/perceptual overlap between primes and targets. If so, this would argue that the priming effect does not reflect semantic processing of the emotions present in the primes. More research is needed in this field o clarify the presence of emotional priming effects, and the mechanism(s) that drive any found effects.

Emotional priming effects are not limited to words/expressions that are generally perceived as emotional by all. Banse (1999) used primes of loved ones' names and faces, the participant's own name and face, and control stimuli. The results displayed priming effects in the form of more positive ratings of neutral stimuli following a loved ones' name/face compared to the participants' own name, or the controls. Banse (2001) went on to find that when participants discriminated between positive and negative affective stimuli, priming with a photo or name of a disliked person created a 'reversed' priming effect (NCE) such that responses to positive target stimuli were actually faster than if the prime had been of a loved ones face or name. This again indicates that priming can sometimes result in NCEs, and suggests this is truer under lower visibility (masked) conditions than in higher visibility conditions (unmasked).

An extension to this, Frings & Wentura (2003) used a Big Brother logo as a prime for affective target images. The extent to which the logo generated priming effects was correlated with how long participants had spent watching Big Brother in the previous week – e.g. the longer participants spent watching Big Brother, the stronger the positive priming. Therefore, affective priming can occur with initially neutral images that have undergone reasonably high exposure and positive associations.

The influence of emotional primes is not restricted to influencing people's response times and ratings of neutral stimuli. Dimburg, Thunberg & Elmehed (2000) presented emotional faces unconsciously, and found congruent activity in facial muscles in the participants. That is, when a happy face was presented subliminally, the participant showed small levels of activity in muscles used to smile. This suggests that subliminal activation reaches a level where there are automatic emotion-related reactions to the images, involving low level mirroring of the emotion in question.

The question of to what extent you can influence someone by subliminal presentation has important implications for media and ethics in a range of situations. In particular, if there was a way to influence the effectiveness of an advertisement, through subliminal messages, is it ethical to allow companies to do this? Could people in voting campaigns win elections through cleverly designed adverts – not just designed to persuade us, but to subliminally alter our opinion of them?

Bateson, Nettle & Roberts (2006) manipulated the image on an honesty box in a university kitchen. They found that the type of image present on the box influenced how generously staff members contributed to their teas and coffee. If the image was of a pair of eyes, they gave significantly more than if the picture was of a flower. The manipulation, whilst not unconscious in terms of presentation times and active disruption to awareness (e.g. masking), went unnoticed by the staff members in that they did not remember the changes. So does this mean that a manipulation this simple can alter how generous someone is? Can a pair of eyes, without us realising, manipulate how socially 'fair' we are?

Stranhan, Spencer & Zanna (2002) would argue that it is possible, with some constraints, to influence how effective an advert is. They used subliminal primes of words related to thirst (e.g. thirst, dry) and showed these could influence behaviour following an advert (more likely to rate highly/drink an advertised beverage). However, this was only the case in participants who were already experiencing thirst.

Winkleman, Berridge & Wilbarger (2005) used subliminally presented faces of varying emotive expression to ascertain the impact on the effectiveness of advertisements. Participants did a short block of a 'gender discrimination task', which contained either positive or negative expression primes (with the same emotion used throughout each block). They then were given a drink, and asked try some of it and to rate it. The emotion present in the prime influenced both how much people subsequently drank, and how highly they rated the drink's value. When the prime was a smiling face, people drank more and valued the drink to a greater extent than if the prime was a frowning face. This experiment is showing a general emotive effect, as the faces/expressions were not directly related to the drink itself (and occurred temporarily separate). The implication is that the emotion was processed semantically, and the positivity of a smile or negativity of a frown either altered people's perceptions of a product, or led to mood induction effects that altered the perceptions of the product. This study measured awareness of the prime faces by informing participants as to the primes presence (after the main experiment), showing them a trial with a prime face, and then asking participants to match the face to two options shown on screen. One of the two options would be the same face as shown – with the same emotional expression as the prime. The other face would show a different expression. Performance was at chance level. This is perhaps not the strictest of measures of awareness – it does after all rely not only on discrimination but also to some extent memory. However, it still implies at least low levels of visibility as the task should be relatively easy if the faces were visible. The main experiment findings appear to be a result of subliminal processing of emotional expressions.

Stewart & Schubert (2006) attempted to influence people's opinions of politicians by presenting the word RATS as a prime before a political advert. They did find a change in subsequent attitude as a function of the prime; however this did not extend to voting intent (planned behaviour). So it seems our political orientation, in terms of voting behaviour, is at least robust to one exposure of a negative prime. This experiment did not measure people's awareness of the prime at any stage, so although brief presentation was used (1 frame – refresh rate not specified) it is unclear whether this effect was subliminal.

A similar study was run by Weinberger & Western (2008), using subliminal primes of either the word 'rats', or a picture of Bill Clinton. Photos were then used of unknown politicians, and participants had to rate how they felt about them. Following a prime of 'rats', negative ratings were increased for the unknown politician. However, negative ratings were decreased following a prime of Bill Clinton. This study is suggesting that people's opinions about an unknown person can be influenced relatively easily. However, this study does not measure a voting behaviour, so the influence could be negligible on actual choice (as found by Stewart & Schubert, 2006). Also, the study was conducted over the Internet. Whilst this has the advantage of reaching a lot of people, it does suffer from lack of experiment control, which is quite crucial when presenting images subliminally.

1.7. Low visibility emotional priming: alternatives to valence congruence effects

The research summarised above has mostly assessed the processing of a subliminal emotional prime by examining whether there is evidence for an advantage in processing, or a bias in choice, towards an emotionally congruent target. This section will review evidence that an emotional prime may also affect target processing through arousal.

Using visible pictures taken from the IAPs, Rozenkrants, Olofsson & Polich (2008) found differences in the arousal of the stimuli affected ERPs to a greater extent than differences in valence. In particular, the arousal led to differences in a late positive component of the ERP. Hinojosa, Carretie, Mendez-Bertolo, Miguez & Pozo (2009) also used ERPs, but with a priming paradigm. They used positive words that varied in arousal as primes and targets, and measured ERPs. No effect of arousal congruency was found on RTs, but the ERPs showed a late positive component (450-550 ms) difference between congruent and incongruent trials when arousal of the prime word was high. A similar effect was not found for low arousal conditions. Here, then, is some evidence that differences in the arousal of emotional stimuli can lead to differences in neural activity. A similar experiment, with subliminal prime stimuli, has not currently been tested. But it certainly seems plausible that arousal could lead to effects on priming, and this could be independent of, and interact with, valence priming effects.

Research that has shown measurable responses to subliminal emotional faces in sympathetic facial muscle activity (Dimburg, 2000), and in SCR (Stone and Valentine, 2001), is also relevant here. These studies provide evidence that subliminal emotional stimuli can lead to physiological responses. These responses are in themselves an index of emotional arousal. Interestingly, SCR have been associated with the amygdala: Glascher & Adolphs (2003) found patients with damage to the amygdala showed reduced SCR to stimuli with high arousal levels in contrast with healthy participants (stimuli were taken from the IAPS database), and this reduction in SCR correlated with the volumetric damage to the amygdala – patients with higher levels of damage showed lower SCR responses. Arousal affects are, then, measurable following subliminal stimuli, and also relate to the network of processing of emotions that will be reviewed in a later section (processing emotions in the brain).

Adolphs, Russell & Tranel (1999) tested a patient with selective, bilateral, damage to the amygdala. The patient could still discriminate emotional valence, but not arousal (tested with faces, words and sentences). This arousal deficit was particularly so for fearful and angry stimuli. In support of this, Berntson, Bechara, Damasio, Tranel & Cacioppo (2007) tested a series of patients with damage to the amygdala, and found patients showed impaired arousal grading for negative stimuli, but not for positive valence stimuli (stimuli taken from the IAPs). However, Anders, Eippet, Weiskopf & Veit (2008) used fMRI in healthy participants and found the amygdala activity was more closely linked to valence than arousal. The role of amygdala activity for valence and arousal remains unclear, but it certainly appears plausible that this structure associated with processing emotions can distinguish between different arousal levels. (See later section for a more general discussion of processing of emotions.)

One relevant research tradition is that focusing on the power of emotional stimuli to draw attention. An example is the dot probe paradigm. This typically involved presenting a stimulus as a potential cue to a target's location. The targets usually involve a simple discrimination, with participants reporting the orientation of the dots, or the location of the dots on screen. Mogg, Bradley, de Bono & Painter (1997) found evidence of emotional words drawing attention to spatial locations, evidenced by faster performance for dot probes presented in the same location. This experiment did not use masked stimuli. A more recent experiment, by Sutton & Altarriba (2011) used masked and unmasked cues (emotional words) and found, regardless of the masking condition, a similar effect: subliminal emotional words drew attention to a spatial location (again evidenced by faster responses to target stimuli in the same location as the cue).

Another, quite extraordinary example of drawing attention was shown by Jiang, Costello, Fang, Huang & He in 2006. They used subliminal erotic images of men and women. Not only did they find evidence of spatial cueing of attention to the location of erotic images, but this was dependent upon the gender and sexual orientation of the participants: the attention of heterosexual men was orientated to the location of an erotic photo of a woman whereas the attention of a homosexual man was directed towards the location of an erotic photo of a man. It is not only the potential for threat, then, than can draw attention – erotic images of the gender we each find attractive can achieve a similar effect.

To conclude, the studies included in this section indicate that emotions can capture attention and lead to physiological responses independent of a second stimulus or any form of response conflict. Whilst there is considerable evidence showing congruence effects in the literature, to both emotional and non-emotional stimuli, there is also considerable evidence for ‘stand-alone’ effects of subliminal emotional stimuli (and to some extent at least non-emotional stimuli). The effect of presenting a subliminal emotional stimulus can be complex, and could – from an observing researcher's perspective – have several kinds of effect on the processing of a subsequent stimulus. These include:

- facilitation/interference with the processing of an emotionally congruent/incongruent stimulus (and with the possibility of not just PCE but NCE effects);
- non-specific arousal modulating processing (e.g. a general increase in processing efficiency, or a change in response criterion);
- attraction of attention to the emotional content of the stimulus, perhaps diverting resources from the task at hand and decreasing general processing efficiency, or activating a more cautious response mode.

Nor are these potential effects mutually exclusive.

1.8. Problem of measuring awareness

One considerable challenge facing the ‘unconscious’ priming literature is that of how to measure awareness. Studies considered thus far, and in the rest of this review, use a variety of different measures to ascertain visibility of images. Many studies in the past have simply asked participants to report their awareness of a stimulus (e.g. Kolb & Braun, 1995), or even assumed that brief presentation will result in subliminal processing (e.g. Stewart & Schubert, 2006). However, it is widely argued that subjective introspective reports are an inadequate measure of visibility. Indeed, Fisk & Haase (2005) found that use of a discrimination task – an objective measure – gave different levels of awareness to subjective self-report. The subjective nature of participants self-report, in the very least, can lead to considerable individual differences.

However, a contrasting view is that awareness, phenomenal consciousness, is defined only by subjective experience (e.g. Dehaene & Naccache, 2001). Can an objective measure sufficiently capture this? Can we know the experience of the participant without asking them? Using forced choice discriminations to assess awareness does not capture of the rich subjective experience, only the availability to the participant of some internal index correlated with the presence or properties of the invisible stimulus (as in the "gut feeling" reported by blindsight patients who lack phenomenal visual experience of the stimulus). However, it does tell us something about the experience of the observer, and it does this in a clear and objective manner that is (theoretically at least) less plagued by issues of participant interpretation. For example, a rating scale of subjective visibility is open to participant's interpretation of the various levels from which they can choose – low visibility to one participant may mean virtually invisible, and to another simply quite difficult to see. When an objective forced choice discrimination is made (was the face present or absent?) this is a lot less likely to occur.

Some researchers have argued that getting full 'suppression,' and being certain a stimulus is presented outside awareness, is not always necessary. Schmidt & Vorberg (2006) suggested that using the principle of a direct measure (e.g. presence or absence of the 'subliminal' stimulus) and indirect measure (e.g. congruence on RTs and errors in priming paradigms), allows the possibility of a dissociation such that a manipulation has differing effects on the two measures (not requiring a zero value of the direct measure.). That is, if you can show ability to detect the presence/absence of a stimulus does not correlate with the priming effect on a response to a subsequent stimulus, then this may be sufficient to investigate unconscious processing without reaching a point at which ability to detect presence is at chance. This is in principle the approach Eckstein & Perrig (2007), and Kiesel et al. (2006) took, by showing independence between two measures.

Wiens (2006) reviewed the literature on unconscious processing, and considered the issue of measuring awareness and gaining sufficient suppression. He also argues that it may not be necessary to get 'full suppression.' If awareness is actually a continuum, and not a binary situation (as he believes), then measuring several points on the 'awareness scale' and comparing this with other measures (e.g. priming) could be enough. This may be true to some extent (it still allows the teasing out of constructs that do not vary with awareness), but it does not solve the issue of *how* to accurately measure awareness.

Another option was put forward by Schmidt (2007), who proposed using an index of brain activity considered necessary for visual awareness (for example, a component of an EEG signal). Lamme (2006) suggests a similar principle, where neuroscience measures, for example a measure of feedback processing, should be utilised to a greater extent, and treated with greater importance. The idea of using quantifiable, identifiable measure from neuroscience does present a pleasing concreteness to an otherwise ill-defined issue. However, the exact measure to use within neuroscience is still not clear, and so the debate here is not solved. And of course using a measure of brain activity is often not an available or practical option in behavioural laboratories.

Perhaps a cautious approach to measuring awareness would be the best. If all researchers used similar objective measures of visibility, then this may reduce some of the noise in the literature and give a higher degree of consistency between experiments. Measures of awareness that are cautious include measuring performance in a task that requires limited perceptual processing. For example, on visibility test trials, one can ask participants to make a forced choice between the presence (or absence) of a stimulus, or about a very basic property (e.g. upright versus inverted face). One can also analyse discrimination in forced choice data of this kind using Signal Detection Theory's d' (e.g. Schlaghecken & Eimer, 2004). This essentially fits a discrimination model with normal distributions of signal and noise of equal variance to the probabilities of correctly choosing and incorrectly choosing one response alternative. It factors out possible changes in bias towards choosing one response alternative, which may be particularly important with emotional stimuli, and it scales discrimination ability in a more principled way than percent correct.

Another point to make concerns the use of the phrases ‘subliminal’ and ‘unconscious’ within the reviewed literature. If awareness is on a continuum, and if there is debate as to how awareness should be measured, and at what point a stimulus may be considered ‘outside awareness,’ is it appropriate to be using these terms? Given that many researchers are using methods of dissociating effects on a measure of visual awareness from effects on a performance index such as priming of RT perhaps we should to adopt the neutral phrase ‘low visibility.’ This may avoid making claims that are untrue of truly 'subliminal' presentation, and also be more accurate in reflecting the continuum that is awareness and human experience.

The research covered thus far in the review has established that awareness can be disrupted in a number of ways (damage to V1, TMS, visual masking), but this disruption does not always lead to a complete loss of ability to discriminate properties of a stimulus, or the ability of a stimulus to influence subsequent behaviour. In particular, there seem to be residual ability to discriminate, and be influenced by, emotional stimuli.

1.9. The processing of emotional stimuli in the brain

I now present a brief review of what we know about the processing of emotional stimuli in the brain, as a framework for considering possible processing routes from visual stimuli to processing of emotion that explain the processing of subliminal (or low-visibility) emotional stimuli. The amygdala has been heavily implicated in the processing of emotions under visible conditions and will be provide something of a focal point for this review.

Dagleish (2004) published a useful review of affective neuroscience. In this review, he outlined the historical theories of the processing of emotional stimuli in the brain from the early theory that the hypothalamus plays a critical role in the processing of emotions (Bard, 1928), to the role of the limbic system (hypothalamus, thalamus, amygdala, hippocampus and cingulate cortex) in processing emotions (MacLean, 1949), to modern day research that highlights several areas of the brain that may be involved. These areas include the prefrontal cortex, amygdala, the ACC, and to some extent the hypothalamus.

A meta-analysis was run on studies looking at the processing of emotions by Phan, Wager, Taylor & Liberzon (2002). They found that for visually presented emotional stimuli, V1 and the amygdala were most reliably activated across studies. There was also some evidence that the medial prefrontal cortex (including the ACC) may have been involved in the general processing for all emotions and all presentation modalities. A more recent meta-analysis by Kober, Barrett, Joseph et al. (2008) found a whole host of areas were involved in the processing of visual emotional stimuli, including: V1, amygdala, superior temporal sulcus (STS), thalamus and several areas of the frontal cortex.

Pitcher, Garrido, Walsh & Duchaine (2008) used TMS to investigate the role of the occipital face area (OFA) and the face region of the right somatosensory cortex (rSC) in emotional face processing in a delayed matching task (two faces are presented temporally separate, and participants have to respond whether they were the same or different). They found stimulating both areas impaired ability to match emotional expressions, but not ability to match identities, in comparison to control conditions of TMS to the vertex or no TMS. This provides evidence that facial expressions are processed in these two areas associated with face processing.

The number of areas associated with the processing of emotions from the modern literature, then, is large and distributed through several brain areas. So far, the following areas have been implicated in processing emotional stimuli in some form: amygdala, V1, STS, thalamus, numerous frontal areas including the ACC, OFA, and the rSC. The remainder of this section will focus on studies that have implicated the amygdala in processing emotional stimuli. The amygdala has been one of the areas most consistently associated with processing emotions, and has been linked to a variety of emotion related functions including: emotional influences on memory, implicit emotional learning, and emotional perception and attention (Phelps & LeDoux, 2005).

As the amygdala will be the main focus for the rest of this section, it is worth first outlining some basic information about the amygdala and its known functions. The amygdala nuclei are small structures located deep within the middle temporal lobes. They are comprised of several nuclei, and implicated in a number of studies looking at different functions. For example, the amygdala has been implicated in studies investigating fear conditioning (Killcross, Robbins & Everitt, 1997), attentional processes (Gallagher & Holland, 1994), and in the formation and consolidation of memories (McGaugh, 2004). The amygdala appears to have connections to many areas of the brain, and a meta-analysis completed by Robinson (2010) highlighted the extent of this: both the left and the right amygdala appear to have connections with parts of the frontal, temporal, occipital, and parietal lobes as well as the cingulate, the thalamus, and the caudate. It would be fair to say the role the amygdala plays in behaviour appears to be broad, and the amygdala has connections with a variety of cortical and subcortical regions. The rest of this section will focus more specifically on emotional processing.

Adolphs & Tranel (2003) studied patients who had suffered damage to the amygdala, and found impairments in recognizing, quite specifically, emotions from facial expressions, with ability to detect emotions from a scene (photo) taken from commercially available films comparable to controls. In particular, there was a disproportionate difficulty in detecting angry facial expressions, which were frequently mistaken for a smile. This result suggests the amygdala is playing a substantial role in processing emotions, particularly for negative facial expressions such as anger.

Further investigation by Adolphs, Gosselin, Buchanan et al. (2005) found that patients with bilateral amygdala damage showed an impaired ability to direct attention to the eyes when looking at faces, and this drives the inability to recognize some emotional expressions. If the patient was told to look at the eyes, then expression recognition was improved. This links in with evidence from autistic patients, who show reduced activity (compared to controls) in the amygdala, and a reduced ability to use eyes to judge emotional expressions (Baron-Cohen et al., 2000). This result could be hinting at an attention-related role for the amygdala, where the amygdala directs processing to aid emotion recognition.

If the amygdala is processing (or heavily involved in the processing of) emotions, then it would be logical that individual differences in emotions and the experience of emotions should be linked to activity levels in the amygdala. Pissiota, Frans, Michelgard et al. (2003) looked at the phobic effects in participants with specific phobias using PET. Half suffered a fear of spiders, and half snakes. Images of each were then used in the scanner, allowing the comparison of snakes vs. spiders in the two types of phobias. They measured an unconscious startle effect to the stimuli (through eye muscle activity), and found this correlated with blood flow in the left amygdaloid-hippocampal region and the medial ACC, and these effects were specific to the class of image towards which the participant reported being phobic. This provides more evidence that the amygdala is important for phobic stimuli processing (as is the ACC); amygdala activation is predictive of individual differences in the experience of (some) emotions.

Phelps, O’Connor, Cunningham, et al. (2000) showed participants unknown black or white faces, and measured an eye blink startle response and activity in the amygdala. Before the study, participants were also asked about their attitude to people of a different race to their own. The strength of amygdala activity correlated with the startle effect, thought to be an implicit racism measure, but not with expressed attitude towards different races. This effect, interestingly, went away if familiar other-race faces were used, especially if they had a positive association. This suggests there is an emotion related automatic response to unknown, other race faces that is expressed in a similar way in the amygdala and the startle response.

Costafreda, Brammer, David & Fu (2008) ran a meta-analysis of 385 published studies, to determine when amygdala activity was most likely to be found. They found that emotional stimuli had a higher chance of producing elevated amygdala activity than neutral stimuli, passive processing had a greater chance than active task instructions, activity was higher for displays of disgust or fear, and the left amygdala appeared to be active more for linguistic tasks while the right amygdala was more active for masked images. Hence it seems there are a variety of factors that influence amygdala activity. In particular, the amygdala seems to be tuned particularly to negative emotions under passive viewing conditions.

There is considerable evidence for the amygdala being important for processing stimuli conveying emotional expressions and threat related images from a variety of sources, and some evidence for the involvement of a diverse network of areas of the brain (e.g. ACC, thalamus, STS). With this in mind, the next section will consider potential processing routes for stimuli under conditions of low.

1.10. A dual route to visual processing?

Previous sections of this review have covered evidence that awareness can be disrupted through damage to V1, TMS to V1, and visual masking. Some ability to discriminate between stimuli remains, and this especially includes the discrimination of emotional expression. The amygdala is heavily implicated in the processing of emotions in general, but how does this residual emotional ability (when awareness is disrupted) come about? This section will focus on research on *where* in the brain activity occurs for emotional stimuli of which patients/participants are unaware, and also discuss processing *routes* in the brain that lead to these emotional blindsight results.

Morris, de Gelder, Weiskrantz et al. (2001) used fMRI on a blindsight patient during presentation of emotional faces to the blind part of the visual field. They found activity in the amygdala following the unaware presentation of emotional faces, and this activity co-varied with activity in the posterior thalamus and the superior colliculus. They conclude that sub-cortical structures are heavily involved in subliminal processing of emotions, and that there appears to be a colliculo-thalamo-amygdala pathway. Williams, Das, Liddle et al. (2005) conducted a similar experiment but using healthy participants with backwards masking of fearful faces. The result was very similar: interactions were found between the brainstem, thalamus and amygdala in relation to masked fearful faces.

The findings of Morris et al. (2001) and Williams et al. (2005) imply that despite the patients' damage to V1, or backwards masking in healthy participants, information from the eyes reached emotion processing networks via a sub-cortical pathway, and it is this pathway that was responsible for amygdala activation. However, some researchers have argued that blindsight is the result of residual functioning in V1, or so called ‘islands’ of spared function. Morland, Lê, Carroll, Hoffmann & Pambakian (2004) found 2 patients with damage to V1 showing blindsight type behaviour. When fMRI was used on both these patients, it was found their residual, unconscious ability was associated with activation of different underlying areas of the brain. One patient showed activity in the lateral occipital cortex, with no evidence of significant V1 activation. However, in the other spared areas of calcarine cortex were active. These results highlight a considerable problem with using patients with brain damage – they are all different, and similar behavioural patterns can be due to varying underlying activity. However, it also indicates that blindsight cannot always be explained by spared function, and residual ability must, sometimes at least, be attributed to other processing routes/brain areas.

A variant of blindsight is Riddoch syndrome, so called as the condition was first reported by George Riddoch in 1917. This syndrome is essentially blindsight following damage to V1, but with some residual awareness of motion and colour. Schoenfeld, Noesselt, Poggel, et al. (2002) studied such patients, and found (using fMRI connectivity analysis) that the sub-cortical areas implicated above in the residual processing in blindsight project directly to V5, which then propagates to V2/V3. In other words, they thought that the activity in visual areas is reached through a ‘back door’ (i.e. through subcortical structures) and spreads backwards through visual cortical areas. This processing is sufficient for some (restricted) level of awareness in patients with Riddoch syndrome. Whilst V1 plays a highly important role in awareness, Riddoch syndrome patients appear to suggest some awareness can be achieved (although it appears very limited) without processing in V1.

Patient data can be difficult to interpret. The Morland et al. (2004) study provides a good example of how patients can vary as to the damage that occurred, and the spared functions. Even if patients are highly similar, they may not give us an accurate insight into 'normal' processing as they frequently suffered the damage years in the past leaving plasticity as an issue, and have taken part in dozens of previous experiments, with the possibility of learning of compensatory strategies. Experiments using healthy participants, then, are a welcome addition to this field of research.

Williams, Moss & Bradshaw (2004) used a binocular rivalry paradigm where different images were presented to each eye and compete for awareness, which typically leads to only the image presented to the dominant eye being noticed. They scanned participants in an fMRI scanner, whilst presenting images of faces either so they should be consciously noticed (e.g. to the dominant eye), or so they would not be consciously noticed (e.g. presented to the 'weaker' eye). The 'competing' images were of houses. The results contrasted aware with unaware conditions, and showed that conscious processing of fearful (but not happy) faces was associated with amygdala activation, whereas unconscious processing for both fearful and happy faces was associated with amygdala activation. Here is a clear, amygdala-related, difference between positive and negative emotion dependent on whether the images were presented consciously or unconsciously, suggesting differential processing related to level of awareness.

The amygdala could be responsible, or at least as part of a wider network, for the processing that results in priming effects. Dannlowski, Ohrmann, Bauer et al. (2007) took the priming effect resulting from negative prime images (facilitated responses to target negative images) and found a significant correlation with amygdala reactivity. This is a nice result, as it does shows not just a difference between two conditions, but a direct correlation with a behavioural measure showing a close link between unconscious processing of emotions and the amygdala. Whilst no causal inference can be made from this, using correlations provides evidence of a more dynamic relationship than an overall association (e.g. from a basic contrast).

The evidence reviewed thus far repeatedly suggests that low visibility emotional stimuli are associated with activity in sub-cortical structures. Given the evidence from blindsight patients, it appears likely (or at least, very possible) that these structures may process emotional information in the absence of input from V1. In other words, it is possible they receive information independent of striate cortex and that this processing can lead to the residual behavioural effects in the absence of awareness.

Le Doux (1996) developed a dual route theory for the processing of emotions in the brain. His theory suggests that the cortical route, which can lead to awareness of stimuli, is relatively slow and processes information high in detail. In contrast, the sub-cortical route is “a quick and dirty processing system” (Le Doux, 1996, pg. 163), which is old in evolutionary history. The benefit of such a system is predominately as a fast way to identify threat and begin an initial reaction (such as a fight or flight response). The proposed subcortical route involves a projection from the retina to the superior colliculus, which then reaches the amygdala through the pulvinar nucleus of the thalamus. A recent review article by Tamietto & de Gelder (2010) argued that this theory has a substantial amount of evidence.

The dual route theory appears to be very much in line with the research reviewed above which has indicated a prominent role for the amygdala and the thalamus. Research has also been summarised that has shown functional co-variation between these structures (Morris et al., 2001; Williams et al., 2006) indicating the presence of a pathway involving the amygdala, thalamus and superior colliculus. One particular prediction that follows from the dual route theory involves timing of activity – if the subcortical pathway really is rapid, and the cortical pathway relatively slow, and the two receive independent projections from the eyes, then there should be evidence of these differences (i.e. activity in the amygdala before striate cortex). Exactly this was found by Luo, Holroyd, Jones, Hednler, & Blair (2007) using MEG. Event-related gamma band synchronisations (thought to represent substantial numbers of neurons firing “in time” in any one area, and relating in this case to an event) were found to fearful faces as early as 10-20 ms in the thalamus area, and 20-30 ms in the amygdala. The striate cortex did not show such synchronisations until 40-50 ms. This, then, suggests the amygdala received a (rapid) input independent of the striate cortex.

The previously mentioned meta-analysis conducted by Robinson (2010) investigated the connectivity of the amygdala, finding many connections between the amygdala and cortical areas. It is well established from work with macaque monkeys that a direct connection exists from the amygdala to striate cortex (Mizuno, Uchida, Nomura et al, 1981; Amaral, Behniea & Kelly, 2003; Freese & Amaral, 2005). Even if, as suggested by LeDoux (1998), subliminal emotions are processed through a subcortical pathway this pathway could interact with the cortical route – it does not necessarily have to be entirely independent of cortical processing. The already mentioned study by Williams et al. (2006) found evidence of a pathway involving subcortical structures for subliminal fearful faces, but also functional connectivity between the amygdala and extrastriate cortex. Activity in the fusiform cortex (an area of temporal cortex associated with face processing) has also been found to correlate with activity in the amygdala during presentation of subliminal emotional faces (Vuilleumier, Richardson, Armony, Driver & Dolan, 2004). In addition, Jiang & He (2006) found activity in the fusiform cortex to subliminal (pattern masked) emotional faces, as well as activity in the superior temporal sulcus (STS) that correlated with activity in the amygdala. Finally, Sabatini, Penna, Franciotti et al. (2009) found activity, in response to angry faces, in the amygdala, anterior insula, fusiform gyrus and the STS. The evidence, then, indicates that a wide range of cortical areas are involved in subliminal processing of emotions, and these areas are frequently found to co-vary with the amygdala. However, many of the mentioned studies used fMRI to image activity, and fMRI cannot tell us the order activity occurred in – i.e. whether activity in the cortex occurs before or after activity in the amygdala.

The research reviewed above could lead to some scepticism concerning a subcortical route to emotional processing. Although there is substantial evidence for the amygdala playing an important role in emotional processing, there is also substantial evidence for a substantial role for cortical regions, and interactions between cortical and subcortical regions, in emotional processing. Fuel has been added to the fire by a recent review by Pessoa & Adolphs (2011), which gives a detailed critique of the dual route theory of emotional processing. The more convincing arguments put forwards by Pessoa & Adolphs will be considered.

The evidence for the presence of rapid emotional processing, and the evidence that structures in the subcortical route are rapid processing routes, are both disputed by Pessoa & Adolphs (2011). They point to evidence that areas of the cortex have been shown active at very short latencies following an emotional stimulus, for example in orbitofrontal cortex 100-150 ms following a stimulus (Kawasaki, Adolphs, Kaufman et al., 2001). In addition, Pessoa & Adolphs point to mixed findings in the literature in regard to emotion specific activity occurring rapidly in subcortical structures. Krolak-Salmon, Henaff, Vighetto, Bertrand & Mauguiere (2004) conducted a depth electrode study in ten patients and found activity in the amygdala modulated by the emotional expression present in a face occurred at 200-400 ms following the stimulus, which is not especially fast compared with the cortical response to emotional expressions. Pessoa & Adolphs make the additional critique that studies using EEG and MEG have fairly poor spatial resolution, making arguments about where any ‘rapid’ activity originated from dubious. These points are certainly valid – evidence for rapid processing of emotions in subcortical or cortical areas remains debatable. However, evidence against rapid processing of emotions seems just as questionable as the evidence for rapid processing of emotions –the jury is still out on the speed of emotional processing.

In terms of the ‘low quality’ processing associated with the subcortical route, Pessoa & Adolphs point to evidence that the amygdala receives inputs from areas associated with highly processed object information with what is thought to be relatively rich detail (e.g. anterior inferotemporal cortex: Amaral, Price, Pitkanen, & Carmichael, 1992). Pessoa and Adolphs also point to evidence that the amygdala shows category-specific responses (to animals versus natural scenes: Kreiman, Koch & Fried, 2000), with the categorisation (in the their opinion) requiring high spatial frequency information. These points do not, on their own, preclude the amygdala processing low spatial frequency information via a subcortical route in addition to high spatial frequency input(s) via a cortical route. Finally, Pessoa & Adolphs suggest that recent research looking at the processing of emotions shows that facial expressions rely on high spatial frequency information, particularly for fearful expressions (Smith & Schyns, 2009). These points of evidence give support to the critique of one of the supposed properties of the subcortical route – its low level processing.

The theory of a subcortical route to processing emotions, called ‘the standard hypothesis’ by Pessoa & Adolphs, generally involves a projection from the retina reaching the amygdala via the superior colliculus and the pulvinar of the thalamus. The role of the pulvinar is disputed on several terms by Pessoa & Adolphs: that it is a passive relay in the processing route, and that a pathway from the superior colliculus to the pulvinar, then from the pulvinar to the amygdala, actually exists. They list studies that have shown the pulvinar is both heavily connected with areas of the cortex (Stepniewska, 2004), and that the pulvinar may play a substantial modulatory (active) role in processing information from multiple areas (Sherman & Guillery, 1996), and possibly in the control of attention (Ship, 2004). This seems incompatible with the idea of the pulvinar mostly just passing information on to the amygdala, but it does not make it impossible. Perhaps somewhat more convincing is the argument that although there is evidence in primates for a connection from the superior colliculus to the pulvinar, and from the pulvinar to the amygdala, the areas of the pulvinar involved in each appear to be both different, and with no lateral connections. The very existence of the route in question, then, is less than certain.

The importance of the amygdala itself in emotional processing is also open to some debate. The role of the amygdala in emotional processing will be covered in a later section, but it is the case that patients have been reported with bilateral amygdala impairment who show very little or no impairment for the latency at which an emotional face broke into awareness in a continuous flash suppression task when compared with control participants (Tsuchiya, Moradi, Felsen, Yamazaki, & Adolphs, 2009).

Instead of a subcortical route, Pessoa & Adophs suggests a multiple waves model of emotion processing with a much greater role for the cortex. They suggest that there are 'short cuts' in the visual system that may enable rapid emotional processing, and may also allow the direction of attention and processing resources to emotional and salient stimuli.

Pessoa & Adolphs interpretation of the literature is open to debate. de Gelder, van Honk & Tamietto (2011) responded to the points raised by Pessoa & Adolphs (2011), noting many counter arguments and weakness in Pessoa & Adolphs suggestions. Firstly, in favour of the existence of a subcortical route, they point to experiments that have shown co-activation in human participants between the amygdala, superior colliculus and the pulvinar (Tamietto & de Gelder, 2010). They also point to evidence the tree shrew has connections from the superior colliculus to the pulvinar, and from the same part of the pulvinar to the amygdala (Day-Brown, Wei, Chomsung, Petry & Bickford, 2010). de Gelder, van Honk & Tamietto (2010) also acknowledge that there are interactions between visual cortex and the amygdala, but argue there is not evidence these interactions relate to emotion, and may in fact be relating to other visual processing, e.g. motion. They also claim there is no evidence activity in extrastriate precedes or is independent of amygdala activity, which they count as evidence against there being multiple independent routes.

For the argument Pessoa & Adolphs (2010) make about high and low spatial frequency, de Gelder, van Honk & Tamietto (2010) argue that both high and low frequency information likely contribute to 'aware' processing of emotional expressions, the argument that when unaware the amygdala route depends more on low frequency information, and there is evidence for this from fMRI experiments (Vuilleumier, Armony, Driver & Dolan, 2003).

de Gelder, van Honk & Tamietto (2011) also suggest that the argument the subcortical route is rapid does not simply mean it will necessarily be the first to show emotion relevant activity. They distinguish between first detectable activity and the first to lead to action. I.e., the subcortical route will lead to faster behavioural and physiological responses. A final argument de Gelder et al. (2010) put forward is that the patient SM, given as an example of a patient with amygdala damage who can still process emotions, is not only a single case study but also has been shown to have deficit fear responses by previous research (Adolphs, Tranel, Damasio & Damasio, 1994) questioning the ability to use SM as evidence against the amygdala's role in processing emotions. Given the mixed evidence in the literature it would perhaps be best not the throw the baby out with the bathwater for the moment, but the Pessoa & Adlophs (2011) article certainly gives food for thought on the relevance and validity of the proposed subcortical emotional processing route.

A final point to consider here in relation to the role of alternative routes in visual processing is the possibility that residual ability in TMS studies and backwards masking results from spared feed-forward processing in V1; i.e. that TMS and backward masking selectively disrupt recurrent feedback (and hence phenomenal awareness) leaving feedforward processing (and hence activation of emotion, meaning, etc.) relatively unaffected. This is unlikely to account for blindsight as there is destruction of tissue in the area of V1 corresponding to the scotoma. It is theoretically possible, but unlikely that TMS could have such a selective effect, as the (only causal) evidence for association between recurrent feedback and awareness is that TMS to V1 in the appropriate (late) interval disrupts awareness. The effects of masking are, of course, hard to directly localise to one cortical area, let alone to one phase of activation of that area.

The precise areas of the brain (and the specific pathways) involved in subliminal processing of emotions remain unclear. It does seem likely there is an input from the retina to the amygdala (via the superior colliculus) that is separate from striate cortex, and that through this route processing of emotional valence in the absence of awareness is possible. More research is needed on the processing of low visibility emotional stimuli, in particular the role of striate cortex, and more generally the role of cortical and subcortical areas, their interactions, and specific roles in subliminal processing.

1.11. A brief word about subliminal processing of non-emotional stimuli

The research summarised in the section above gives a picture of a dynamic and interactive set of brain areas involved in subliminal processing of emotions, with a key role for subcortical structures. But research cited earlier in this review found evidence that non-emotional properties of low visibility stimuli can also lead to behavioural effects (e.g. Kiesel et al., 2006; Schlaghecken & Eimer, 2004; Vorberg et al., 2003; Weiskrantz et al., 1974). This thesis examines differences between the processing of neutral and emotional stimuli, and so some discussion of brain areas involved in processing non-emotional stimuli is relevant here. Are they essentially the same as emotional stimuli, or are there some differences?

Moutoussis & Ziki (2002) used fMRI during unconscious and conscious presentation of (neutral) faces and houses. During unconscious processing there was generally less activity than during conscious processing. That is, faces and houses activated their associated areas of the occipito-temporal cortex in both aware and unaware conditions, but to a different extent. This study suggests conscious and unconscious processing differs only on the basis of quantitative differences and do not involve different processing routes. With respect to the amygdala, Suslow, Ohrmann, Bauer, et al. (2006) found that activity in the amygdala correlated with ability to detect the presence of faces (regardless of emotion), suggesting that for non-emotional stimuli, amygdala activation relates to the level of visibility (whereas for emotional stimuli, the amygdala activity remains fairly robust). Both of these studies, then, highlight the possibility that, for non-emotional stimuli, the difference between aware and unaware may be better characterised quantitatively than qualitatively.

How is it that blindsight patients can discriminate, above chance, non-emotive stimuli if the differences between unaware and aware conditions are purely quantitative? One possibility is for a direct projection from the LGN to V5 (that bypasses V1), which could then back-propagate to ‘earlier’ visual areas and spread to extrastriate cortex. Evidence already reviewed has suggested the possibility that V5/MT receives information from the eyes in blindsight patients with Riddoch syndrome (Schoenfeld et al., 2002; Barbur, Watson, Frackowiak & Zeki, 1993). There is also support for this idea from animal studies, with Sincich, Park, Wohlgemuth & Horton (2004) reporting a direct projection from the LGN to V5 in macaque monkeys. Some form of alternative route, then, could explain blindsight patients’ residual processing ability for non-emotional stimuli, although this alternative route appears to be different to the processing of emotions.

As with emotional stimuli, what brain areas mediate the processing of subliminal non-emotional stimuli is far from clear. However, it does appear that the processing of non-emotional stimuli may arise through a direct projection to V5 (via the LGN) mostly involve visual cortex. More research to investigate non-emotional stimuli under low visibility conditions is needed, and the differences between emotional and non-emotional stimuli require some clarification. However it does not appear likely that the same route is involved for both emotional and non-emotional stimuli.

1.12. A brief word about visual attention

This review, and most of this thesis, focuses on visual awareness. A related, although separate, topic worth mentioning here is the literature on visual attention. A discussion of subliminal emotional stimuli being able to attract attention was mentioned earlier, but this section will explicitly compare attention and awareness.

Koudier & Dehaene (2007) note that there are two possible ways conscious awareness can 'fail.' These are through a lack of sufficient strength of activation from a stimulus presented to an attended location, or through a failure to attend to a stimulus that may otherwise have sufficient strength. That is, attention is the 'gate keeper of consciousness.' Dehaene, Changeux, Naccache, Sackure & Serggent (2006) defined three types of processing: conscious, subliminal and preconscious. Subliminal refers to conditions like those discussed in most of this review. Preconscious refers to situations where there is a failure of attention – the stimuli in question may be processed to a sufficient level to gain awareness, but it has not been allowed through the gateway of attention.

The separation of attention and awareness is not particularly new. Lamme (2003) notes there is evidence from psychological and neuroscience perspectives that stimuli that are unattended can later become available to awareness, and attended items can fail to reach awareness (for example in masking paradigms). This links into a distinction noted in by Block (1990), who categorised two types of awareness: phenomenal and access. Access awareness does require attention, and is generally reportable/remembered. Phenomenal awareness on the other hand involves *experience* of the stimuli, but not necessarily the ability to report it and does not require attentional resources.

Whilst awareness and attention are two different processes, they are highly related and can influence each other. If a stimulus does not have enough strength to reach awareness, this review has shown it can still be processed meaningfully and influence behaviour. However, research suggests that spatial attention amplifies subliminal processing. Lachter, Forster & Ruthruff (2004) found less priming was evident for a lexical decision task if the prime location was not attended to (i.e. prime and target were at different locations). In fact, they went as far as to say no evidence of priming was found when the prime was presented at an irrelevant (to the target) location on screen and not attended to. Sumner, Tsai, Yu & Nachev (2006) used a response priming paradigm where the prime was presented either above or below the fixation point, with the target also presented at one of the two locations, and not necessarily the same location as the prime. At the start of the trial, the location of the prime arrow was cued correctly or incorrectly, and this cue had no bearing on the location of the target. The authors found cueing attention to the location of a prime arrow increased the size of the priming effect. This effect was found to not simply be due to increased visibility for the ‘attended’ primes. Some priming remained present when the cues were not attended, so in this case attention appeared to *amplify* the priming effects.

Emotional stimuli have been shown to capture attention, and this is another source of the argument that emotions can be processed independent of attentional resources: if emotions were not processed outside attention, it is difficult to see how they could draw attention. Examples of the emotion capture effect include work by Ohman, Lundqvist, & Esteves (2001) who found schematic angry faces are found faster in a search array, arguably due to an attention capturing effect. However, the emotional capture effect has been questioned by Huang, Baddaley & Young (2008) who found task instructions altered the emotion capture effect in an attentional blink task. When participants were asked to search for a specific semantic category from words in a rapid serial visual presentation stream, emotional distractor words reduced participants’ reports of targets relative to neutral distractor words. In contrast, when targets were identified on a perceptual basis, there was not difference in target reporting when distractors were emotional or neutral. Emotional stimuli’s ability to capture attention, then, appears to be limited and potentially linked to task specific conditions.

Brain imaging research has led to a mixed picture as to the relation between attention and emotional processing. The research was reviewed by Silvert, Lepsien, Fragopanagos et al. (2007), who note there has been research finding unattended emotional stimuli can lead to similar levels of activity in the amygdala (e.g. Vuilleumier, Armony, Driver & Dolan, 2001), and others that show unattended emotional stimuli show a reduction compared to attended (e.g. Pessoa et al., 2002). They suggest the research that found no difference between attended and unattended emotional processing did not have a high enough perceptual load in the distracting task (e.g. what it was that stopped participants attending to the emotional stimuli). They presented faces and houses in the periphery of vision and found unattended fearful faces led to stronger activation in the amygdala than neutral faces. However, this finding was limited to trials when a concurrent task with peripheral stimuli had relatively low demands (discriminating whether the tilt direction of two houses was the same or different). When the task had higher demands (discriminating whether the identity of two houses was the same or different) the difference between fearful and neutral faces in the amygdala disappeared. The authors linked this effect to attention – the higher load task required more attentional resources than the lower load. It could also be argued that the load manipulation called upon other resources (e.g. working memory), but either way the findings implicate that areas associated with processing emotions can be modulated by the demands of a concurrent, emotion irrelevant, task.

Overall, then, debate continues whether unattended emotional stimuli lead to similar or markedly lower levels of activity than attended stimuli in regions of the brain thought to process emotion. However, those that show a reduction in activity for unattended conditions tend to show just that – a reduction, but not an abolishment of activity relating to the emotional stimuli. Some effect of unattended emotional stimuli remained suggesting that attention amplifies emotional activation, not that attention is required for processing at all. Of course, it could be that the remaining emotional processing evident in priming tasks and in neural activity is due to some attentional resources remaining and being allocated to the emotional stimuli – this as argued by Silvert et al. (2007), and cannot be discounted.

Attention, then, modulates subliminal processing and the relationship between preconscious and subliminal processing is an interesting and developing research topic.

Either a lack of attention, or a lack of stimulus strength, can lead to a lack of awareness. The experiments in this thesis were designed to keep attention relatively constant across conditions to selectively investigate low visibility processing.

1.13. Summary

In summary, it seems that V1 plays an important role in visual awareness. Disruption to processing in V1 following brain damage, or through application of TMS, can lead to reduced (or absent) awareness. Visual masking, in particular backwards masking, can also disrupt awareness. Whilst the precise neural locus of visual awareness is far from clear, it seems apparent that feedback processing, and V1, play substantial roles in visual awareness. Whether this is part of a larger reverberation network, and/or if there are specific brain regions/routes that are particularly important for visual awareness, remains debateable.

When awareness is disrupted, some residual effects on behaviour from the disrupted stimuli remain. Particularly when the stimuli are emotional, it seems that they can speed up (or slow down) responses to a subsequent stimulus with a congruent value, influence two alternative forced choices, and have some impact on preference and ‘liking.’

It appears that sub-cortical structures such as the amygdala play an important role in processing of emotion in low visibility stimuli, and there is now considerable evidence that a direct projection from the retina to subcortical structures (a “hot line to the amygdala”) is responsible (or at least important) for the ‘residual’ ability to process emotional information. However, the number of routes to emotion processing remains debateable, with evidence for one route, two relatively separate routes, and several routes that process stimuli in parallel.

Research to date has used a variety of methodologies to study visual awareness and emotion processing. There is still debate, however, as to the nature and requirements of visual awareness, processing of emotion in the brain, the conditions and nature of low visibility processing, and the impact of low visibility processing of emotional stimuli on our behaviour. In the experiments reported in this thesis I used either TMS or visual masking to disrupt processing and awareness of body posture and face stimuli. The aim was to search for evidence that processing of emotional content is privileged – i.e. can survive disruption of visual processing that impairs discrimination of featurally equivalent non-emotional content, consistent with the “hot line” account. The experiments used several behavioural measures:

- 2AFC discrimination of emotional and non-emotional properties of a low visibility image
- the impact of a low visibility prime on RT to a following visible target as a function of emotional congruence,
- the impact of a low visibility prime on forced choice to a neutral target.

Processing in V1 was locally disrupted using TMS in two experiments reported in Chapter 3 and 4. Chapter 5 reports 2 experiments that compared masking functions for discrimination of emotional and non-emotional properties of faces. Chapter 6 presents 4 experiments exploring the effect of an emotional prime presented under low visibility conditions on the decision about a subsequent target stimulus.

CHAPTER 2

Transcranial Magnetic Stimulation: basic principles, possibilities, and applications to research in visual processing

---

2.1. Introduction

This chapter introduces the method of Transcranial Magnetic Stimulation (TMS). It is of particular importance to this thesis as TMS has been applied with interesting results to visual processing and awareness, and TMS is used in the experiments reported in the following two chapters. This chapter will cover the background to TMS: the specific niche of TMS within neuroscience, how TMS works, various types of TMS and protocols that can be used, and the main advantages and disadvantages.

2.2. The neuroscience 'niche' of TMS

TMS is a relatively new technique in neuroscience that can be used to create 'virtual lesions' in the brain. It gives the ability to create participants temporarily equivalent to patients with localised cortical damage in the lab, with specific effects to the brain that are either very short-lived or, with certain paradigms, last for as much as an hour. Other modern methods in neuroscience (e.g. fMRI, EEG, MEG and PET) have all increased our understanding of how the brain works and how it is that cognition comes about. However, they link a function to a particular brain region in an essentially correlational manner: when a person does X, part of the brain (Y) is active. This does not give us causality, nor tell us what it is, specifically, that Y contributes to X. TMS is of particular use due to its ability to interfere with (or sometimes sensitise) specific brain areas directly. By pairing TMS with behavioural paradigms, it is possible to stimulate specific areas and measure the (direct) impact on the task performance or behaviour in question.

The versatility of TMS for studying cognitive neuroscience is high, and if used correctly TMS could further the quest for understanding the brain and behaviour considerably. The value of TMS is not going unnoticed – the use of TMS in published research has increased in recent years from around 600 publications between 1994-1998, to around 3000 between 2003-2008 (Rossi, Hallett, Rossini & Pascual-Leone, 2009). TMS has many advantages, but also some limitations and potential problems that will be outlined in subsequent sections. First it is important to cover how it is specifically that TMS works.

2.3. How TMS works

TMS works on the basic principle that a rapidly changing magnetic field can induce electrical currents (electromagnetism). A capacitor stores a large charge, which is released to a coil of copper wire. This creates, very briefly, a magnetic field up to 2 Tesla in strength. If the coil of copper wire is held next to the scalp, current will be induced in cortical tissue up to 3cm from the scalp (Rossi et al., 2009).

Some of the first TMS machines used monophasic pulses, which show a fairly rapid rise to the maximal induced current (peaking at around 0.1 ms) which then decays and reaches baseline around 1 ms after the initiation of the pulse (Walsh & Pascual-leone, 2003). This type of pulse was impractical for use with repetitive TMS designs, as it required too much energy to be used with pulses in rapid succession. Biphasic pulses were developed that had the advantage of a faster rise time to the maximal electrical field, followed by a rapid negative dip (the second 'phase'), with all impact of the pulse finished around 0.3 ms from the beginning (Walsh & Pascual-leone, 2003). The advantage to a rapid peak in electrical field is that neurons are not good at storing charge, and the impact of a relatively slow increase in field will be minimal and wasteful. The biphasic pulse uses less energy, and manages to recycle about 60% of the electricity it uses into the next pulse (Walsh & Pascual-leone, 2003). This, then, allows for a rapid succession of pulses making repetitive designs more practical.

There is a general consensus that TMS induces weak electrical fields that can raise activity (via depolarisation), and lead to action potentials (Walsh & Pascual-Leone, 2003). This can interact with a cell's existing level of 'activity,' so an action potential is more likely in a cell that is active above a resting state (Nagarajan, Durand & Warman, 1993; Pitcher, Walsh, & Duchaine, 2011). As the impact of the induced current will only cause depolarisation of an axon if the electrical charge is different across the cell membrane, it requires the induced field to vary across the axon (Nagarajan et al., 1993; Walsh & Pascual-leone, 2003). If an axon lies in such a way that the induced field does not vary across the axon, there will be no effect of the stimulation. As a result, not all cells that lie within the ‘area of stimulation’ will be directly stimulated by the pulse.

2.4. Spatial resolution of TMS

Depending on properties of the coil, the area of stimulation can be quite broad (for example, if a round coil is used) or relatively precise (if a figure-of-8 coil is used). It is thought that the diameter of the area of stimulation with a figure-of-8 coil is around 1-2cm (Pitcher, Walsh & Duchaine, 2010). The depth of stimulation is, currently, limited to the outer part of the cortex, and (as already noted) thought to extend up to 3cm from the scalp when the coil is placed against the head (Rossi et al., 2009).

To give an example of how precise the impact of TMS can be, Pitcher, Charles, Devlin, Walsh & Duchain (2009) showed distinct, selective impacts of TMS to three adjacent areas in visual cortex: the occipital face area (OFA, MNI coordinates: 45, -80, -12), extrastriate body area (EBA, MNI coordinates: 50, -72, -2), and the object area (OA, MNI coordinates: 44, -75, -6). TMS to each of these areas, which were functionally located using fMRI for each participant, resulted in disrupted discrimination performance for face, body, and object parts respectively. The triple dissociation following TMS to three adjacent areas of cortex, that showed selective disruption to each area in turn, is an impressive example of the spatial resolution of TMS in action.

Functional spatial specificity can also, to some extent, be measured when stimulating the visual cortex. When TMS is applied to early visual processing areas, such as V1, it can induce a visual percept. This is known as a phosphene, and is typically experienced as a small flash of light. The exact size and intensity will relate to the output of the machine. As V1 is topographically organized, varying the stimulation point on V1 can vary the location of the phosphene in an individual’s visual field. Phosphenes can also be used as a tool for mapping V1 in relation to the visual field (Fernandez et al., 2002). They also allow a measure of whether V1 has been successfully stimulated: a manipulation check of sorts that is rare for brain stimulation experiments.

Another point worth a brief mention here is that so far TMS has been discussed as if it affects only the local area of cortex that lies underneath the coil. However, there is evidence that the effects of TMS could be more wide spread. The brain is not made up of isolated patches of cortex – areas of the brain are interconnected and part of broader networks. For example, the visual system, as covered in Chapter 1, includes several regions of the occipital, temporal and partietal cortex that are connected with feedforward and feedback connections. In the context of TMS, then, it is likely that stimulating one area can have knock-on effects in other functionally connected parts of a network. Early work into the possible wider reaching effects of TMS was carried out by Paus, Thompson, Comeau, Peters & Evans (1997) by pairing stimulation with a measure of brain activity using PET. This work found that stimulation of the frontal eye fields led to changes in activity in the frontal eye fields, but also in the superior parietal and medial parietal-occipital regions – i.e. those regions previously thought to be functionally related/connected to the frontal eye fields. This work is a good example of how TMS can have wider ranging effects than is often acknowledged in papers, but it also highlights how TMS can be used as tool for researching functional connectivity in the brain.

2.5. Stimulation protocols and temporal resolution

Single pulses of TMS can be used to investigate processing that occurs at very specific times, in an 'online' design, where pulses are given just after (or sometimes just before) a stimulus to which the participant must respond, so as to interfere with particular processes. Though the pulse itself lasts for less than 1 ms, this does not necessarily mean its impact is limited to the same extent – a single pulse could have effects that last considerably longer. To gain some measure of the temporal specificity of single pulses of TMS in practice, the best approach is to see how specific the effect of single pulses of TMS can be on behaviour.

A few studies have used single pulses of TMS to V1 to disrupt visual processing, and these were reviewed in Chapter 1. The studies have generally found that pulses to striate cortex as little as 10-20ms different in their onset can have markedly differing impacts on perception (Amassian, Cracco, Maccabee, et al., 1989; Corthout, Uttl, Walsh, Hallett & Cowey, 1999; Corthout, Uttl, Ziemann, Cowey & Hallett, 1999). If experiments can find differing effects separated by such short intervals, then this suggests TMS is really quite temporally specific. Although we cannot state with any certainty the precise duration of the effect of a single pulse on processing, we can say the pulse itself lasts less than 1 ms, and a difference as little as 10 ms in pulse onset can influence the effect of TMS on behavioural measures, which suggests an effective duration of the order of tens of ms.

As well as using TMS to deliver an effect at a very specific time, a short train of repetitive TMS (rTMS) pulses can be used to spread the effect over a more extended interval. This can be useful in cases where little is known about exactly when processing occurs in a region, the temporal properties are of less interest, or the impact of a single pulse is not sufficient for a measurable effect and a more robust interruption is required. In the previously mentioned study by Pitcher et al. (2009), pulses were delivered over a window of 500 ms, at a frequency of 10Hz, on each trial when a stimulus was presented. This design gives comparatively course temporal information (compared to single pulse designs), but it can give a fairly robust disruption effect, as shown by Pitcher et al's (2009) triple dissociation.

It is also possible to use much longer trains of rTMS pulses before the presentation of the trials of an experiment, to generate an effect sustained through those trials, and measure the impact on behaviour: an 'offline' design. Offline designs have the advantages that researchers do not need to know when during a trial to try to disrupt processing, and that distracting effects on performance of auditory, tactile and muscular artefacts of the TMS pulse (see below) are avoided. A participant with a prolonged TMS disruption is more similar to a neuropsychological patient. It is also not necessary for participants in an offline study to keep their head still through the whole experiment. Offline rTMS is technically much less problematic than online TMS if the research also uses fMRI or EEG measures – although techniques have been developed to combine these with online designs. The main disadvantages of offline designs, apart from the lack of temporal precision, are the higher dose and greater risk of seizure from truly repetitive TMS (see below), and the greater difficulty of contrasting performance in conditions with more than one TMS site.

A form of rTMS called Theta Burst Stimulation (TBS) has recently been developed for offline designs. This uses a more complex series of pulses over a relatively short period to produce a surprisingly long-lasting effect. The protocol of the pulses can be continuous or intermittent, but the most frequently used at the moment, and with greatest reported success, is continuous TBS (cTBS) consisting of 3 pulses at 50Hz repeated at 5Hz. Huang et al. (2005) compared the effects of several TBS stimulation protocols, including cTBS, to the motor cortex and measured the impact over time on Motor Evoked Potentials (MEPs). The main findings are shown in Figure 2.1 (reproduced from Huang et al., 2005). They found that following 40 seconds of cTBS, the MEPs were reduced. This reduction was most pronounced at around 18 minutes following the stimulation, and was mostly gone around 60 minutes after stimulation. The same pattern was not found with intermittent TBS (iTBS), which actually showed the opposite result – MEPs were increased. Different pulse designs, then, can influence the resulting effect on the cortex.

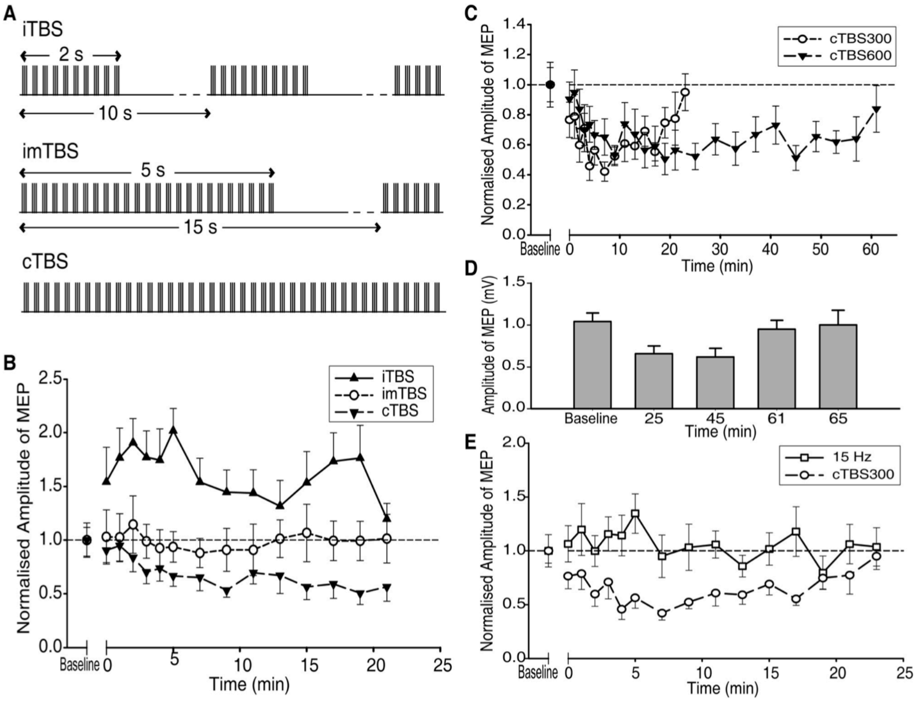

Figure 2.1. Paradigms of TBS and Their Effects on MEPs (A) Graphical illustration of the three stimulation paradigms. (B) Time course of changes in MEP amplitude following conditioning with iTBS, cTBS, or imTBS. (C) Comparison of the effects of cTBS given for 20 s (300 pulses; cTBS300) with the same paradigm given for 40 s (600 pulses; cTBS600). Reproduced from Huang et al. (2005).

cTBS has been applied by researchers to the striate cortex. Franca, Koch, Mochizuki et al. (2006) found the phosphene threshold (level of stimulation intensity required to elicit a phosphene) was increased by around 10% following 40 seconds of cTBS, with no impact following iTBS. To examine the effect on visual discrimination, Waterston & Pack (2010) measured performance in a variety of 2AFC discrimination tasks (e.g. horizontal or vertical lines in a Gabor patch). rTMS at 1Hz for 20 seconds, and cTBS for 40 seconds both improved performance at the task, relative to stimulation of the vertex. This facilitation was found across the visual field, not localised to the area of disruption. This finding is in contrast to single pulse designs finding disruption to discriminations following TMS to V1 (e.g. Amassian et al., 1989). It is also surprising that the effect was not localised within the visual field. The impact of cTBS on V1, then, remains unclear.

One possible explanation for the findings of Waterson & Pack (2010), an explanation the authors themselves noted, was that the repetitive stimulation led to noise reduction in the striate cortex. This idea draws on the basic principle that any area of the brain likely contains ‘background noise’ of neuronal firing, and/or neurons may respond to other information and tasks. If this is true, then it makes sense that by reducing the level of activity in that brain region, and by extension reducing this noise, performance at a task may actually get better. This idea leads to findings in direct contrast with the lesion hypothesis suggested earlier – that by reducing activity in an area of the brain the effect of stimulation on activity, and so the resulting behavioural or physiological effects, will work in a similar manner to a brain lesion. In reality both ideas are not mutually exclusive, and which one is true could well depend on the strength of stimulation, the brain area being stimulated, and specific behavioural or physiological measure used.

2.6. Main advantages and disadvantages of TMS

If the use of TMS is compared to research with patients, even patients with very circumscribed lesions, then TMS has clear advantages. Not only does TMS allow for study of the timing of the relevant processing, but also allows participants to act as their own controls. It is not possible to measure ability and performance in patients without the deficit, while participants in TMS experiments can be tested before, during, and after stimulation giving baseline measures of performance and allowing within-subjects designs for many experiments. Also, patients rarely have damage to one isolated area, but frequently show damage to multiple areas and/or show multiple deficits. In general, then, TMS allows for much greater control than using patients, and the ability to make causal inferences allows for the testing of hypotheses generated through reports of patients and results from fMRI and EEG experiments. This considerable advantage, of permitting hypothesis-driven research, was noted by Walsh & Pascual-Leone (2003, pg. 93): “One can now select the lesion for the question rather than select questions determined by the availability of the lesion.”

One important consideration in TMS designs is the need for a control site. If no control stimulation site is used, it is difficult to know whether the measured effects are due to the impact on neural processing of the pulse at the location stimulated, or a general reaction to the sensory stimulation (i.e. the noise of the pulse, sensation of the pulse on the scalp, effects of muscle twitches induced by the pulse). The ideal control site is one similar in terms of sensation and potential ‘annoyance’ to the participant: for example, stimulation of areas close to muscle (especially in frontal and temporal regions) can lead to twitching in muscles of the face or neck that can be irritating and distracting. If behaviour is affected by TMS in the experimental but not in the control location then it is unlikely the cause is anything other than the specific stimulation. Choosing the control condition can be difficult, and certainly warrants careful consideration when designing TMS experiments.

TMS does suffer some potential limitations, risks and downsides. Only the outer 2-3 cm of brain (i.e. the surface of the cortex and part of the cerebellum) is accessible to TMS as the effective field penetrates only a short distance from the coil. The peak sound intensity produced by a TMS pulse has amplitude high enough that it risks damaging hearing (but the "click" of the pulse is so brief that subjectively it does not sound very loud). TMS studies protect against damage with earplugs, and exclude any participants that could have more sensitive hearing (e.g. albinos, a family history of congenital deafness).

The more serious risk from TMS is that of inducing an epileptic seizure. Using a single pulse technique, there have been no reported incidents of induced seizures. However, the risks increase with repetitive pulse designs. Whilst this risk is still very small, researchers must always screen against any potential risk factors in individual participants, be careful to select stimulation protocols that are safe, and limit the number of pulses given in any one session or 24 hour period. Examples of exclusion criteria for participants are: any family history of epilepsy, any incident of head trauma or a neurological condition (e.g., meningitis), recent taking of illegal drugs or some prescription medications, consumption of alcohol in the last 24 hours, and consumption of significant amounts of coffee (more than 2 cups) earlier that day. If strict exclusion criteria are used, and a stimulation protocol chosen with care then the risk from TMS appears to be minimal. To ensure newly developed (and existing) protocols are safe there is a regular review regarding safety and TMS, and updated advice for exclusion criteria when screening participants (e.g. Rossi et al., 2009).

2.7. Implications for the present studies

TMS can be used to create very brief or longer lasting disruption to very precise areas of the cortex. TMS is proving to be a valuable tool in cognitive neuroscience, and in particular for investigating visual processing and awareness. One particular advantage of TMS to V1 is the ability to localise in subjective visual space the impact of the pulse of TMS on cortical processing through use of phosphenes. This can allow us to study the effect on behaviour of a disruption of visual processing well localised within the visual field.

One particular benefit of using phosphenes in this way is as follows. Because the stimulation site has to be fixed for at least a block in online designs, and for the whole session in offline designs, most TMS studies are constrained to use between-block contrasts between performance with stimulation in the site of interest and performance with a different (control) site of stimulation, or performance in a sham- (or no-) stimulation condition. Participants may adjust their performance or arousal levels to the perceived difficulty of the current block, or differences in the sensations they expect, leading to potential confounds between the effects of such adjustments and the effects of stimulation. However, when V1 is stimulated with a figure-of-eight coil to produce disruption to processing in a limited area of the visual field localised by the phosphene, stimuli can from trial to trial be presented randomly either within this area or in the equivalent area in the opposite hemifield, the latter trials serving as the no-disruption control. This kind of design was used in Experiments 1 and 2. Experiment 1 (described in Chapter 3) used an on-line design. Experiment 2 (described in Chapter 4) attempted a theta-burst offline design.

CHAPTER 3

TMS to V1 spares discrimination of emotive relative to neutral body postures

---

3.1. Abstract

This study used TMS to examine the role played by striate cortex (V1) in processing the emotional content of visual stimuli. Participants learned to discriminate two sets of body posture images. For half of each set, the posture's emotional significance (threat versus pleasant) provided a redundant cue for the discrimination; the other half were emotionally neutral. Stimuli were briefly presented at a lateral location in the visual field where a TMS pulse produced a phosphene, or at a control location in the opposite hemifield. A TMS pulse 70 - 140 ms after stimulus presentation at the phosphene location impaired discrimination of the neutral stimuli with little effect on the emotional stimuli; the two classes of stimuli were equally discriminable when presented at the control location. The results are consistent with the proposal that recognition of emotionally salient patterns such as social threat may be accomplished by a sub-cortical "hot line" bypassing V1.

3.2. Introduction

In this article we report an experiment on the contribution of cortical area V1 to the processing of the emotional significance of visually presented body posture images. LeDoux (1996) proposed a dual-route account according to which, in addition to the geniculo-striate route, a subcortical route plays an important role in visual perception of biologically critical emotions such as social threat. The suggestion, developed by later researchers, is that this subcortical route, including components of the emotion-processing network such as the amygdala (which receives information from the retina via the thalamus) provides crude but rapid and relatively automatic discrimination of threat and possibly other emotionally salient visual events and objects, and that this has provided an adaptive advantage in facilitating rapid responses to potential dangers (see Tamietto & de Gelder, 2010, for review). In a recent critique Pessoa & Adolphs (2011) challenge this account (with its dominant role for the colliculus-pulvinar-amygdala route); they argue that evidence on connectivity and processing speed in the cortex suggests multiple rapid parallel routes for emotion processing in the brain, and propose a more modulatory role for the amygdala. In response, de Gelder, van Honk & Tamietto (2011) argue that evidence from fMRI co-activation of sub-cortical structures with amygdala, spatial frequency specific effects present only in subcortical structures, and studies of patients with damage to the amygdala, all imply a substantial role for a sub-cortical route to activation of amygdala.

We address here only the issue of whether processing via primary visual cortex is required for discrimination of emotional stimuli. To ask this question, one can examine whether the processing of such stimuli is selectively spared by damage to V1 or by temporary suppression of its function with transcranial magnetic stimulation (TMS). Both tests have been used to address the contribution of V1 to emotion processing, and also, in rather more studies, its contribution to phenomenal visual awareness. Although our interest is in emotion processing, some consideration of the literature on visual awareness is appropriate.

Inspired by the recovery of some degree of functional vision in a rhesus monkey whose striate cortex had been ablated (Humphrey, 1974), Weiskrantz, Warrington, Sanders & Marshall (1974) studied a patient (DB) with localized damage to V1, resulting in an area of the visual field in which he was unaware of stimuli, and unable to report them. However, when asked to make a forced choice about properties of these invisible stimuli (such as the horizontal or vertical orientation of a line), he performed significantly above chance. This residual perceptual ability is commonly referred to as "blindsight." Several studies have now looked at the processing of emotional information presented in the "blind" part of the visual field in blindsight patients, and obtained forced choices better than chance about the emotional expressions of faces (de Gelder, Vrooman, Pourtois, & Weiskrantz, 1999; Morris, de Gelder, Weiskrantz & Dolan, 2001) and body posture images displaying emotion (de Gelder & Hadjikhani, 2006). Patients have even been fear-conditioned to line drawings of airplanes presented to the blind part of the visual field and paired with a shock (Hamm, Weike, Schupp, et al., 2003).

Sahraie, Weiskrantz, Barbur, et al. (1997) used fMRI to assess unconscious processing in a blindsight patient. The participant had to respond with forced choices to properties of images in a blind or aware part of the visual field. For stimuli of which the patient was aware, there was substantial activity in the neocortex. For stimuli of which the patient was unaware, activation shifted to more sub-cortical structures such as the superior colliculus. Also using fMRI, Morris, de Gelder, Weiscrantz & Dolan (2001) measured brain activity in a blindsight patient during presentation of fearful faces to the blind part of the visual field. They found activity associated with the fearful faces in sub-cortical structures (the amygdala and the thalamus). From this, they inferred that a sub-cortical pathway is responsible for residual processing. The amygdala responds to both positive and negative emotions (Williams, McGlone, Abbott & Mattingley, 2005; Yang, Menon & Eliez et. al., 2002; Garavan, Pendergrass, Ross, Stein & Risinger, 2001); it is not yet known whether sub-cortical emotional activation of amygdala is effective for all emotions or limited to threat-related images.

Other evidence for processing of emotional visual stimuli of which a subject is phenomenally unaware has come from experiments requiring healthy participants to categorize the emotion in briefly presented masked images (de Gelder & Hadjikhani, 2006; Morris, Öhman & Dolan, 1999; Öhman, Carlsson, Lundqvist, & Ingvar, 2007), and from masked priming paradigms (Kouider & Dehaene, 2007). Although such research continues to generate controversy, the idea that it is possible to process emotional (and indeed other) attributes of subliminal stimuli is no longer heterodox (e.g. Pessoa & Adolphs, 2011; Tamietto & de Gelder, 2010). But of course behavioral evidence for a dissociation of processing and awareness does not directly support claims about anatomical pathways. One interpretation is that processing sufficient to prime processing of a subsequent target, or bias a guess, is carried out in exactly the same pathway but at levels of activation too low to activate a broader network associated with awareness (Kouider & Dehaene, 2007). In Lamme's (2001) theory of visual awareness, "feedforward" processing of visual information through V1 to extrastriate areas is followed by feedback from higher levels to V1, and it is only this recurrent processing that engenders visual awareness.

Similarly, a persistent concern with evidence from blindsight patients is that there might be spared function in V1 sufficient for basic visual discriminations, but not for awareness. Morland, Lê, Carroll, Hoffmann, & Pambakian (2004) scanned two blindsight patients with fMRI. In one, they found activity in V1 in response to an ‘unconscious’ image, suggesting the blindsight in that case could be explained by spared function, but the second patient did not show any evidence of this. These results highlight the problems of using patients: cases of localized cortical blindness are rare, different cases have idiosyncratic features, and it is possible that spared function could explain the behavior. The use of Transcranial Magnetic Stimulation (TMS) in healthy participants offers the prospect of a greater degree of control over the suppression of function in V1, as well as the ability to investigate the time course of visual processing.

Amassian, Cracco, Maccabee, et al. (1989) varied the interval between a briefly flashed letter string, and a TMS pulse to V1. If the pulse was applied in a time window 80-100 ms after stimulus onset, it could prevent a participant from either reporting the letters, or being visually aware of them. Later studies have indicated that TMS can disrupt processing of letters within two time windows following a stimulus: one early on, estimated by different researchers to be 10-60 ms post-onset and a later window around 100-140 ms (Corthout, Uttl, Walsh, Hallett & Cowey, 1999; Corthout, Uttl, Ziemann, Cowey & Hallett, 1999; Kammer, 2007). The earlier window tends to be less robust (present in fewer participants). These effects have been shown not to be caused by eye blinks from the TMS pulse (Corthout, Uttl, Juan, Hallett & Cowey, 2000). The apparent presence of two distinct time windows in which TMS to V1 can disrupt processing is one source of Lamme's (2001) proposal that visual awareness arises from recurrent feedback to V1 from extra-striate cortex. Fahrenfort, Scholte & Lamme (2007) report that the difference in the occipital ERP signal to masked and unmasked presentations of a textured square emerges around 109-141 ms after target onset, suggesting that the recurrent ‘feedback’ stage is disrupted not only by TMS but by backward masking.

When TMS is applied to V1, most people experience a phosphene, a localized flash or patch of visual sensation. As well confirming the topographic relationship between V1 and the retina/visual field in humans (Fernandez, Alfro, Tormos, et al., 2002), this makes it possible to determine the location within the visual field of the effect of a TMS pulse. Most studies applying TMS to V1 have used a circular coil, which stimulates a fairly broad region of underlying cortex. We used a figure-of–eight coil to stimulate a more circumscribed area, and exploited phosphenes to locate the stimulated region within the visual field so that we could deliver a brief stimulus to a location where processing was visibly disrupted by TMS.

If the emotional valence of a visual stimulus can be processed via a subcortical pathway without processing in V1, then a task requiring emotional discrimination ought to be relatively immune to TMS that disrupts discrimination of other visual properties. This possibility was investigated by Jolij & Lamme (2005) using simple schematic faces (emoticons). Stimulus displays consisted of four such faces arranged in a small (2°) square: one sad or happy face (differing only in downward or upward mouth curvature) and three neutral faces. Participants made a two-choice decision about either the location of the emotive face (left or right of fixation) or the emotion (aggressive/happy), while one pulse of TMS was given at a variable interval following the stimuli. At a pulse onset of 110 ms, the face could not be located to one side better than chance, but emotion discrimination remained above chance. This is *prima facie* evidence that performance at a control task can be disrupted, whilst performance at emotion judgments is relatively unaffected. However, the use of schematic faces distinguished only by simple features such as upward or downward curvature of the mouth leaves open the possibility that participants were simply detecting, without locating, curvature values, not processing emotion. And the evolutionary rationale for the "hotline" theory makes it hard to understand why the sub-cortical pathway would be able to detect a salient expression or mouth shape but not be able to locate it, though it may be noted that the distance between left and right location for these stimuli was quite small.

The present study asked whether disrupting processing in V1 with TMS impaired processing of the emotional valence of a body posture image[1] in the same way as it disrupted processing of other visual properties. The typical strategy for asking this question (exemplified by Jolij & Lamme, 2005) is to compare the effects of TMS on performance of two tasks, one requiring discrimination of emotion, one of non-emotional properties. The problem is to match the difficulty of the two tasks and the discriminations required. For example, if the perceptual discrimination required for the emotion task is easier, it may, although processed via V1, be less vulnerable to the degradation to V1 representation resulting from TMS than the non-emotional discrimination. In this kind of paradigm the two tasks also have to be run in separate blocks (to avoid the additional complexities of task-switching), raising the possibility of differences in effort, attention, arousal, etc., between blocks.

We therefore adopted a different approach, to allow a within-task and within-block comparison of the effects of TMS over V1 on processing of emotional and neutral images matched for difficulty of discrimination. Participants first learned to classify

---

[1] Much research on processing of emotion expression has used facial expression, but the detection of emotion in a distant person's posture is no less adaptively important than reading a closer person's facial expression (cf. de Gelder & Hadjikhani, 2006).

eight body posture images into two apparently arbitrary response categories of four images each. For two images in each category, their emotional valence (threat versus pleasant) provided an additional, redundant, cue for the categorization. The other two images in each category were similar but emotionally neutral, and chosen to produce an equivalent level of categorization performance to the emotional images in the absence of TMS. For each participant, we found a coil location over the occipital scalp such that a TMS pulse generated a lateralized phosphene in the opposite visual hemifield, thus establishing localized perturbation of V1. Participants then attempted to perform the learned categorization with a single TMS pulse following the stimulus on every trial. The posture image was displayed unpredictably either in the phosphene location (where V1 processing was known to be disrupted), or in the equivalent location in the opposite hemifield (as a control condition, with no direct effect of the pulse on V1 processing in that location, but with the auditory and somatosensory consequences of the pulse still present). The pulse onset time varied unpredictably between 70 and 140 ms following stimulus onset. We ensured central fixation by occasionally requiring reports of a tiny digit presented centrally and discriminable only with foveal vision.

If the rapid processing of emotional valence is immune to disruption of V1 processing, we would expect to see a selective impairment to the discrimination of neutral stimuli presented in the phosphene location. Performance on emotional versus neutral stimuli presented to the control location in the opposite hemifield enabled us to assess how well their discriminability was matched in the absence of TMS. Note that, although the stimuli were certainly hard to see, we make no claims about, and did not formally assess, visual awareness.

3.3. Method

3.3.1. Participants

Thirty-four students from the University of Exeter participated. Five did not complete the experiment, as phosphenes could not be satisfactorily induced. (Two reported phosphenes ipsilateral to the coil location; three did not report any phosphene at all.) Of the remaining 29, who were tested over three sessions on different days, 15 were men and 14 were women; average age was 21. All gave informed consent, having read information sheets on TMS and on the experiment, and passed TMS screening checks. They had normal (or corrected-to-normal) vision. Participants were paid up to £24, with an element of the payment depending on their performance.

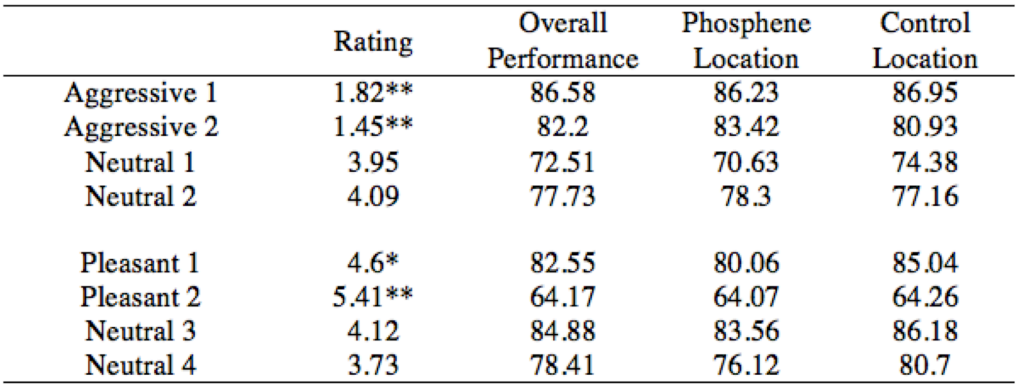

	Rating	Overall Performance	Phosphene Location	Control Location
Aggressive 1	1.82**	86.58	86.23	86.95
Aggressive 2	1.45**	82.2	83.42	80.93
Neutral 1	3.95	72.51	70.63	74.38
Neutral 2	4.09	77.73	78.3	77.16
Pleasant 1	4.6*	82.55	80.06	85.04
Pleasant 2	5.41**	64.17	64.07	64.26
Neutral 3	4.12	84.88	83.56	86.18
Neutral 4	3.73	78.41	76.12	80.7

* = p < 0.05, ** = p < 0.01

Table 3.1: Average ratings and average discrimination performance (% correct) for each of the eight stimuli used. A low rating (1) means very aggressive, a high rating (7) very pleasant, and a mid-rating (4) neutral. All four of the emotional stimuli were rated as reliably different from a neutral rating – none of the neutral stimuli were reliably different from neutral ratings.

3.3.2. Stimuli

The images used, shown in Figure 3.1, were monochrome photographs of one person in eight different poses, taken on the same camera, at the same exposure, etc. The actor was shown what to do with a demonstration and instruction. The background of each image was removed and the face pixilated using Adobe Photoshop Elements. The images included two aggressive and two pleasant postures, and four matched neutral

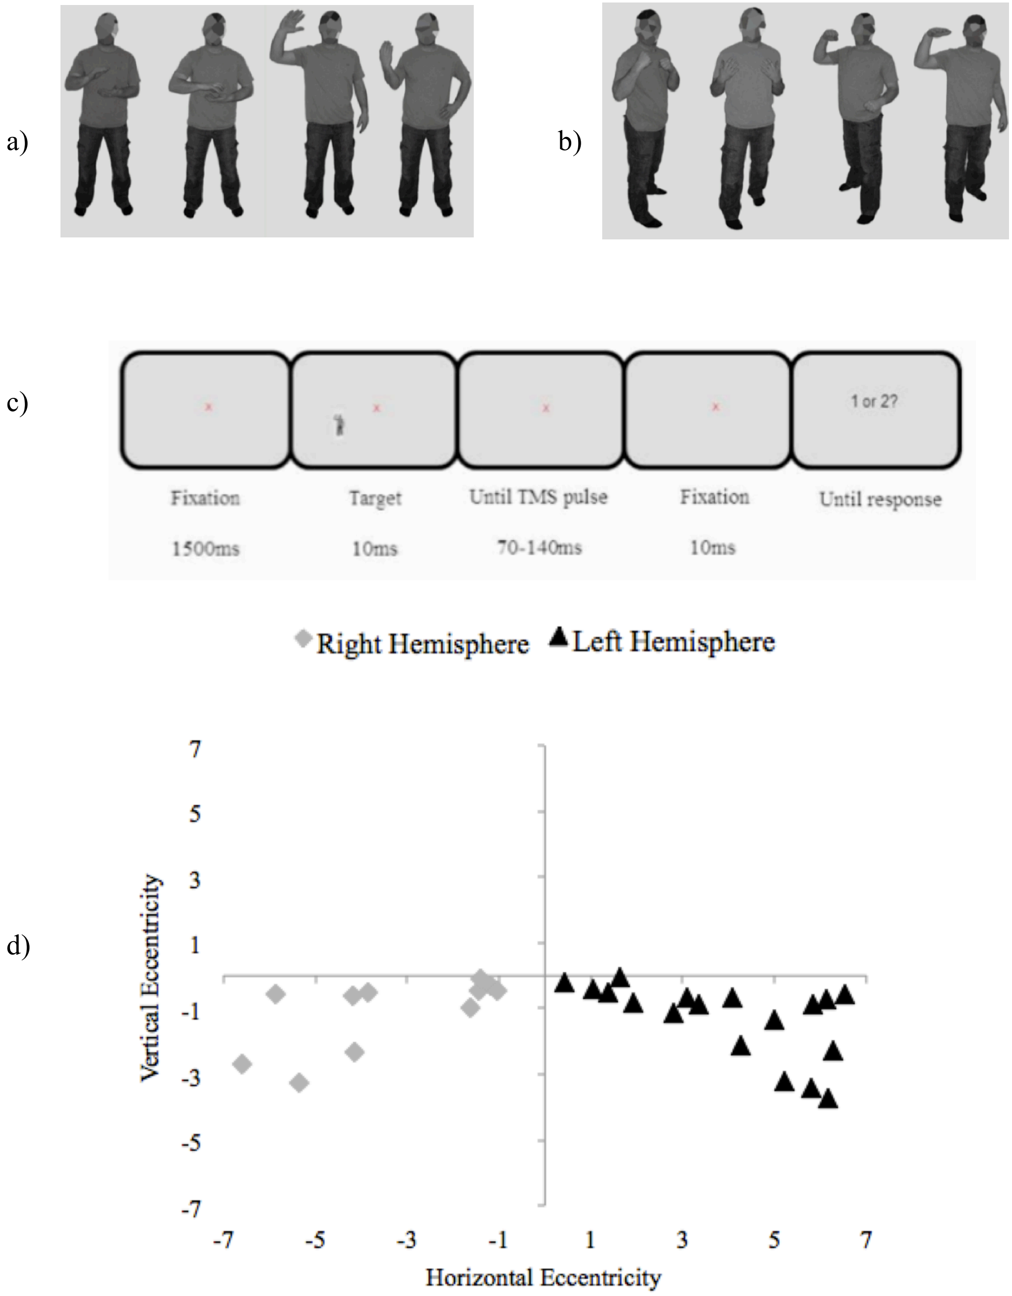

Figure 3.1: The body posture images making up the two categories: (a) the aggressive postures and two similar neutral postures; (b) the pleasant postures and the two similar neutral postures. (c) Sequence of events on a standard trial. (d) Coordinates (in degrees of visual angle from fixation) of each participant's phosphene location.

postures designed to be similar to the emotional postures but without obvious emotive content. Figure 1a shows the aggressive and their matched neutral postures, and Figure 1b the pleasant and their matched neutral postures.

An online questionnaire was used to check the emotional valence of the postures. Thirty people rated the images (presented in random order) on a scale of 1-7, with 1 representing ‘aggressive,’ 4 representing ‘neutral’ and 7 representing ‘pleasant’ expressions. The average ratings can be seen in Table 3.1. The neutral faces had average ratings close to 4. On average, both the aggressive postures were rated as clearly aggressive; one of the pleasant postures received a relatively strong pleasant rating, and one a less extreme pleasant rating.

3.3.3. Materials

The experiment was run on a PC using *Presentation* (version 14.3, Neurobehavioral Systems), and a 19” CRT screen with resolution of 1024 x 768, set to a 100Hz refresh rate. Participants sat 70 cm from the screen, head stabilized by a chin rest. The images subtended a visual angle of 2° vertically. A Magstim Rapid 2 stimulator and 70mm figure-of-eight coil were used to generate the TMS pulse.

3.3.4. Procedure

To induce and locate a phosphene a TMS pulse was applied at 70% of stimulator output while the participant fixated a central fixation point on the screen. The coil was held approximately vertically against the back of the head, to the right or left of the midline, with the handle held downwards at an angle of 45° away from the vertical across the midline. To avoid suggestion, very little information was given about what a phosphene would look like, and where in the visual field it was likely to occur. Once a phosphene had been obtained, the stimulation site was adjusted until the phosphene covered a region clearly to the left or right of fixation. Participants moved an example image of a body posture to the location on the screen of the phosphene using arrow keys[2]. The eccentricity of this region varied somewhat over participants, but was between 0.47° and 8.27° with a mean of 4.12° (275 pixels) and a SD of 2.27°. For 11 participants the region was on the right side of the screen, and for 18 participants it was on the left[3]. Individual participants’ phosphene locations are plotted in Figure 1(d). The pulse

---

[2] Most participants reported the phosphene covered the majority of the stimulus, but the disruption of neural processing caused by TMS is likely to extend well beyond its visible consequence.

[3] Although we attempted to alternate between hemispheres over participants, it was not possible to get a phosphene in both hemispheres in everyone, hence we ended up with unequal numbers of participants stimulated on the left and right.

amplitude was then varied to find the threshold at which a phosphene was no longer seen. For the rest of the experiment, an output of 120% of phosphene threshold was used, which averaged 71% of stimulator output.

Participants were next taught to discriminate between the two categories, each of four images, shown in Figure 3.1. Category 1 contained the two images for one emotion (aggressive) and the two similar neutral images (Fig 3.1a). Category 2 contained the two images for the other emotion (pleasant) and the two similar neutral images (Fig 3.1b). Assignment of categories to left and right response keys was balanced over participants. To start with, each image was displayed once centrally in a random order, and the participant told whether it belonged in category "1" or "2". At no point were descriptions such as "aggressive", "happy", "pleasant" or "emotion" used. There followed 80 practice trials on which an image was displayed centrally for 500 ms, the participant responded, and feedback was displayed ("Correct" or "Error"); the images were presented 10 times each in a random order. All participants performed with accuracy close to 100% well before the end of practice. Initially during practice, the emotional images were discriminated marginally easier (4% difference in accuracy), but the difference diminished with practice (2% difference in accuracy). When they came in for the second and third sessions, they were shown the images again one at a time on screen (with no time limit) to refresh their memories.

For the main experiment, on 89% of the trials one image was presented laterally, either at the *phosphene location* established as described above, or at the equivalent co-ordinates in the opposite hemifield (the *control location*). The events of a trial are depicted in Figure 1(c). One biphasic TMS pulse was given per trial, always over the same phosphene-eliciting scalp location established for that participant. The pulse followed the stimulus onset by a *pulse-onset time* of 70-140 ms (varied in 10 ms steps). Both the hemifield in which the image would appear, and the pulse onset, were unpredictable; trial sequences were newly randomized for every participant. The display "1 or 2?” appeared 150 ms after the stimulus onset to prompt a response, and remained on screen until a response was given: participants were encouraged to take their time and respond as accurately as possible. A new trial began immediately after the previous response, with no trial-by-trial feedback. At the end of each block, participants were informed whether their accuracy had exceeded 70%, in which case a bonus of £0.07 was awarded.

At the beginning of the trial, participants were instructed to fixate on a central dot. To ensure participants maintained fixation, a random 11% of the trials were fixation-test trials. On these trials, no posture image was presented nor was a TMS pulse given. Instead, at the fixation point, a small digit (0.2° tall) between 1 and 6 was presented for 50 ms, and the participant responded with the left key for odd numbers, the right for even. The small size of the digit ensured that discriminating its value was impossible if it was not displayed in the fovea, and hard even when it was. The task was deliberately made to be difficult in order to ensure participants fixated at the center of the screen – if the task were too easy, participants might be able to perform the task without fixating. An additional bonus of £0.07 was awarded for each block for which correct reports exceeded 70%.

After training, the experiment consisted of a total of 1860 trials (52 per combination of emotional/neutral stimulus, phosphene/control location of the stimulus, and pulse onset time, plus 196 fixation control trials), spread over three one-hour sessions. In order to keep the stimulation site the same, during the first session the location of stimulation was saved to the participant's MRI scan (available from another research study) using Brainsight. This location was then found and used on subsequent sessions.

3.3.5. Data analysis

Performance at the fixation task was analysed as percent correct, and a one-way t-test carried out to determine if performance was reliably greater than chance. As this task was designed to be difficult, performance above chance is indicative of participants maintaining fixation.

For the rest of the data, performance for each condition was indexed as d′. This was calculated with the formula: $d' = z(H) - z(FA)$, where H is p(correct) for Category 1 postures, and FA is p(error) for Category 2 postures; probabilities of 0 or 1 were replaced by 0.019 and 0.981 (halfway between 0 or 1 and the nearest possible value, as suggested by Macmillan & Kaplan, 1985). The d′ values were then submitted to an ANOVA with the factors of stimulus type, stimulus location and pulse onset time. The critical interaction of interest was between stimulus location and stimulus type, with the prediction that neutral but not emotional body postures would be affected by the stimulus location on screen.

Note that throughout the thesis I do not show error bars on any graphs except those showing the effects of a single variable. In a within-subject factorial design, there are several error bars that could be plotted, corresponding to the MS Error (interaction with Subjects) for each of the main effects, their interaction(s), or some pooled error term. Instead, all reports of the means of important effects and interactions in the text are accompanied by ± the standard error of the mean (SEM) contrast.

It is worth mentioning that I do not report effect sizes. Effect sizes are usually reported to indicate that the effect investigated by the experiment contributes a substantial fraction of the variance in the measures. However, this thesis pursues effects that we expect to be small (e.g. discriminability of masked stimuli); the question is whether they can be detected at all.

3.4. Results

3.4.1. Fixation check and stimulus eccentricity

Performance on the fixation test ($M = 77.4\%$, $SEM = 0.03$) was substantially greater than chance (50%), $t = 10.414$, $p < 0.001$, with all participants above chance (range = 54-98%). This indicates that the occasional digit probes were effective in maintaining fixation. Over individuals, performance on the fixation test also correlated significantly with overall performance level, $r = 0.88$, $p < 0.001$: those who found the fixation task difficult also found the main task difficult. As both are difficult perception tasks this is not surprising, but if some participants had deliberately fixated the lateral part of the field to improve their performance on the posture categorization trials, this would tend to produce an inverse correlation.

The eccentricity of the stimulus (which varied somewhat over participants) correlated significantly and inversely with overall performance, $r=-0.45$, $p<0.05$. This too is unsurprising, as acuity declines with distance from the fovea. Indeed, the existence of this effect provides additional evidence that participants were fixating centrally.

3.4.2. Classification performance

The critical interaction is shown in Figure 3.2. For images displayed at the control location in the hemifield opposite to the phosphene location, performance was almost identical for emotional and neutral images. But when the image was displayed at the phosphene location, discrimination of the neutral stimuli was clearly impaired, while performance on the emotional stimuli was only slightly reduced. A repeated-measures ANOVA was run with the factors: stimulus location (phosphene versus control), stimulus type (emotive versus neutral) and pulse onset time (70 – 140 ms); a between-subjects factor of gender was also included (but interacted reliably with no other factor). There was a significant interaction between stimulus location and stimulus type, F (1,27) = 7.18, p < 0.05 (interaction contrast: M = 0.20, SE = 0.08). There was a significant simple main effect of stimulus type for presentations at the phosphene location, F (1,27) = 7.39, p < 0.05, but no hint of one at the control location, F < 1. There was a significant simple main effect of stimulus location for neutral stimuli, F (1,27) = 10.96, p < 0.01, but not for emotive stimuli, F (1,27) = 1.11.

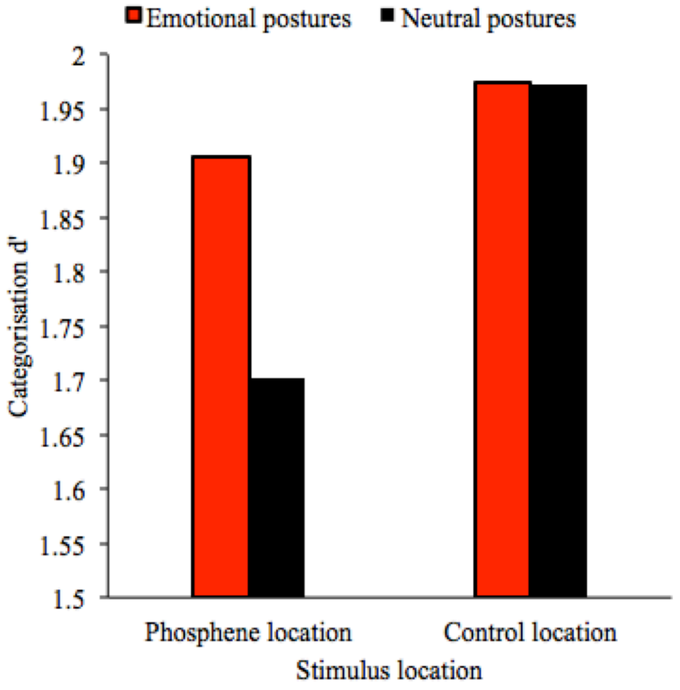

Figure 3.2: Group mean categorisation performance (d’) for each stimulus type (emotive and neutral) at each stimulus location (phosphene or control location).

Figure 3.3 shows the effect of pulse onset time for each combination of stimulus location and stimulus type. The inferior performance on the neutral stimuli under TMS can be seen across the whole range of pulse onsets, and the interaction between stimulus type, stimulus location and pulse onset was not significant, F (7, 189) = 1.41. Thus the critical interaction depicted in Figure 3.2 was not limited to any specific time window within the range of pulse onsets tested. If for the neutral stimuli we examine the difference in discrimination between the phosphene and control locations as a function of pulse onset time, some reduction in the size of this difference is discernable beyond an onset time of 120 ms, but the interaction between phosphene/control location and the linear component of pulse onset time did not approach reliability, F (1,27) = 2.12, p =

0.16. Nor was there any evidence that the lack of effect is due to averaging over individual functions with windows of maximum effect at different onset time. This would predict that fitted quadratic coefficients of this function for individuals would generally be negative, but with the peak at different places. The estimated quadratic coefficients were in fact negative for only 15 out of the 29 participants, clearly no different from chance.

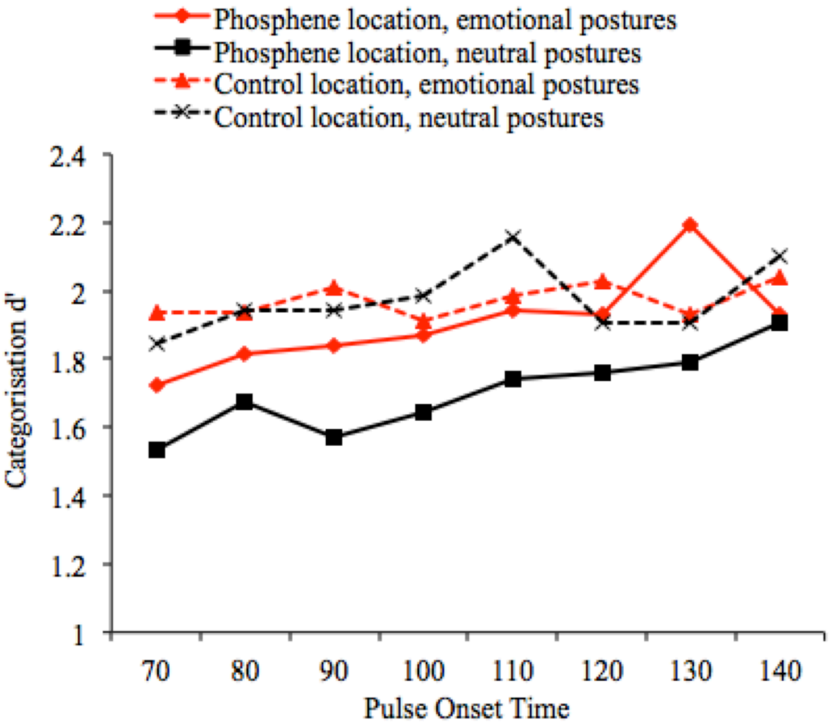

Figure 3.3: Group mean categoisation performance (d') as a function of stiulus location, stimulus type, and pulse onset time.

Figure 3.4 shows the relation between individuals' mean d' overall, and the size of the interaction contrast whose average is depicted in Figure 3.2. There is little sign that the critical interaction is restricted to the participants showing good discrimination, and the correlation was not significant, r = 0.147, p = 0.447. We make no claims about awareness (though most participants certainly found these brief lateral stimuli hard to see and discriminate). But if we assume the visibility of the stimuli is correlated over individuals with their overall performance, then Figure 3.4 does not suggest that the critical interaction in Figure 3.2 depends on the stimuli being either generally visible, or subliminal.

As noted above, the aggressive images were on average rated somewhat more

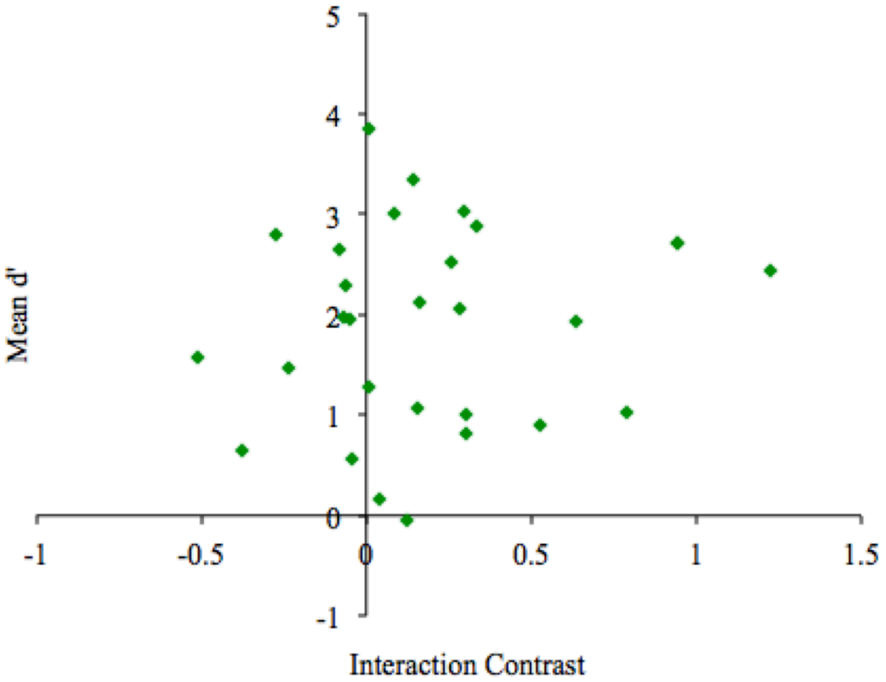

Figure 3.4: Relation between overall categorisation performance and the size of the interaction depicted in Figure 2: interaction contrast = (phosphene location emotive – phosphene neutral) – (control location emotive – control location neutral).

negatively than the pleasant ones were rated positively. Nevertheless, given the idea that a sub-cortical route might be limited to threat-detection, it is of interest to ask whether the critical interaction was markedly different for aggressive and positive images. Hence a further ANOVA was run with the variables stimulus category (postures in Fig 3.1a versus postures in Figure 3.1 b), stimulus type and location, this time on percent correct (as d′ cannot be estimated for one category alone). The mean percent correct can be seen for each stimulus and condition in Table 3.1. There was no significant interaction between location, stimulus type and stimulus category, $F < 1$, and the size and direction of the critical interaction contrast (emotional/neutral x location) was comparable for the aggressive postures ($M = 1.6\%$), and the pleasant postures ($M = 1.9\%$), though neither was statistically significant on its own due to the reduction in power (and possibly the use of percent correct).

The hemisphere stimulated (and the resulting phosphene hemifield) was varied between participants. To check for any laterality effect, an ANOVA was run with hemisphere stimulated as a between-subjects factor, and stimulus location and stimulus type as within-subjects factors. (One left hemisphere participant was excluded as her phosphene location was so close to fixation – 0.47° – that TMS may have activated the opposite hemisphere as well.) The overall interaction between stimulus location and stimulus type remained significant, F (1, 27) = 5.92, $p < 0.05$. The pattern was much the same for both groups: the interaction contrast was somewhat larger with the left hemisphere stimulated (M = 0.23; SE = 0.08) than with the right (M = 0.16; SE = 0.07), as may be seen in Table 3.2, but the difference was nowhere near significant, $F < 1$.

	Phosphene Location		Control Location	
	Emotion	Neutral	Emotion	Neutral
Left	1.71	1.56	1.71	1.78
Right	2.22	1.94	2.4	2.29

Table 3.2: Average discrimination performance (d') for each stimulus type and stimulus location, for participants stimulated in the left hemisphere, and the right hemisphere.

3.5. Discussion

Participants learned to classify body posture images into two arbitrary sets. For half the images in each set, the posture's emotional valence (aggressive/pleasant) was an additional redundant cue for the discrimination. Participants then classified such an image displayed briefly and unpredictably either in the same lateral location as a TMS-induced phosphene, or in a control location in the opposite hemifield. Control of fixation appeared satisfactory. When displayed at the control location, emotional and neutral stimuli were equally discriminable. When the stimulus was displayed in the phosphene location – the part of the visual field where cortical processing was disrupted by the TMS – the discrimination of the neutral postures was impaired, while the discrimination of the emotional postures was significantly less affected, and the small effect was not significant.

Discrimination of the posture's emotional attitude (threat versus pleasant) thus appears relatively immune to the disruptive impact of a TMS pulse on the processing in V1 that mediates discrimination of the non-emotional content. This is consistent with the idea (LeDoux, 1996, Liddle, Brown, Kemp, et al., 2005, Morris, Öhman & Dolan, 1999, Tamietto & de Gelder, 2010) that visual processing of information specifying biologically critical emotions, such as social threat, enjoys a sub-cortical route, bypassing V1, to components of the emotion-processing network (possibly including the amygdala). This evidence complements, using complex natural images, Jolij & Lamme's (2005) demonstration using highly schematic faces that TMS disrupted judgments of the location in the field of an emotional singleton to a greater extent than judgments of the singleton's emotional valence. The disruption of the ability to locate the emotional singleton to left or right reported by Jolij & Lamme might, given their small (2°) stimulus array, suggest that the subcortical route supplies only relatively crude configurational or location information.

The obvious alternative account of the critical interaction to consider is that the emotional and neutral stimuli were both processed only (or equally) via representations of their features in V1, but discrimination of the emotional stimuli was for some reason more robust to the disruption of V1 processing caused by TMS. One reason for such robustness might be that the visual postures denoting threat or pleasantness were more meaningful, familiar or easy to label than the neutral postures, and that the resulting match to familiar templates meant that less precise V1 representation was needed was to discriminate them. Another possibility is that the perceptual features carrying the emotional valence were somehow more robust to V1 degradation than the features by which the neutral stimuli were discriminated, though it is not obvious what the different features might be, or why their processing would enjoy this immunity.

We believe any variant of the proposal of differential robustness to V1 degradation is rendered implausible by the identical performance on emotional and neutral images when they were presented in the control location, where V1 processing was not disrupted by TMS. Discrimination of these briefly displayed parafoveal images was certainly not at ceiling: they were hard to discriminate for most participants. Nor does Figure 4 suggest that the critical interaction comes from participants at ceiling. That the emotion images were not easier to classify under clearly data-limited conditions at the control location, suggests that any greater "robustness" was giving the emotional images no advantage at this level of image degradation. It would seem somewhat implausible that the functions relating image degradation to performance would be such that only the additional degradation due to TMS would reveal the greater robustness of the emotional images.

If we accept the sub-cortical bypass account of the relative immunity of the emotional stimuli to TMS, the question arises why, given that emotional valence provided a redundant cue for the categorization, performance at the control location, with no TMS, was not superior for the emotional images. One possibility is that the availability of discrimination via the high-acuity geniculo-striate route somehow suppresses utilization of the sub-cortical route. In other words, if sufficient information is present from the higher-acuity geniculo-striate route, information from the lower acuity and potentially less accurate rapid subcortical route may be ignored. Another is that the emotional images were actually somewhat *harder* to discriminate via the cortical route than the neutral ones, and this difference cancelled out the advantage that would otherwise have been seen due to the redundant emotional cue in the no-TMS control location. Either way, the selective immunity of performance on the emotional images to TMS cannot be explained by their discrimination being easier or more robust to degraded presentation.

Yet another possibility to consider is that the sub-cortical bypass account is correct, but it is not emotional salience as such but something else — perhaps meaningfulness or familiarity — that is detected via a sub-cortical route, and the emotional stimuli were more meaningful or familiar. There are two problems with this interpretation. First, the possession of meaningfulness or familiarity did not map on to the arbitrary categorical distinction required for these stimuli, whereas positive versus negative valence did, and thus could serve as a redundant cue. Second, prior evidence and theory already makes a case for sub-cortical processing of emotion; sub-cortical processing of meaning more generally would be a more surprising claim.

Unlike some studies on the impact of TMS to V1 on visual awareness, we found no evidence for a critical time window for stimulation within the 70-140 ms range we sampled, though there was a modest but non-significant reduction in the size of the effect for pulse onsets 120 ms or more after the stimulus. There must presumably be an upper and lower bound, but there is no reason to expect the useful processing in V1 for categorical discrimination to be limited to first-pass or recurrent processing, even if the latter is associated with visual awareness. The difference between the findings of this experiment, and the findings of previous TMS studies, i.e. whether there is a specific critical time window or a more broad level of disruption (although 70-140 ms is not, of course, particularly broad), could be due to a number of factors. These include the overall visibility of the stimuli prior to disruption, the use of more focal stimulation with a figure of eight coil, or the properties of the stimuli themselves, e.g. real life body posture images as opposed to simple emoticons or letters.

The experiment did not aim to achieve unconscious presentation of the stimuli, nor assess visual awareness formally. However, the brief para-foveal presentations certainly rendered the stimuli hard to see as well as discriminate. Participants varied in their overall performance, in part because of variations in the eccentricity of the stimulus location. If we assume that visibility correlated with overall performance, the lack of any obvious relation between overall d' and the critical interaction suggests that the selective impact of V1-TMS disruption on discrimination of neutral stimuli is not dependent on visual awareness of the stimulus or its absence.

In relation to the possibility that the sub-cortical “hot line” to emotional networks is limited to threat-detection, the critical interaction did not seem to be limited to the aggressive images, even though the aggressive images were given more extreme ratings on average than the pleasant ones. However, the experiment was designed to measure discrimination between the categories, not compare performance across categories; this analysis could not use d', lacked power, and assumes the emotional and neutral stimuli were better matched across stimulus categories than the ratings suggest. The hemisphere stimulated also had no detectable impact on the critical interaction.

In summary, a single pulse of TMS to V1 after stimulus onset selectively impaired categorical discrimination of neutral body posture images, but discrimination of perceptually similar aggressive versus pleasant postures was almost unimpaired. The most obvious interpretation is that the discrimination of emotional content was spared because it could be processed via a sub-cortical route that does not depend on V1.

CHAPTER 4

The effect of Theta Burst Stimulation on V1

4.1. Introduction

The experiment reported in this chapter investigated the role of V1 in neutral and emotion processing of faces. As covered in Chapters 1 and 3, there is theory and evidence to suggest 2 (or more) routes in the brain for processing visual emotional information, with one (or more) independent of V1 (LeDoux, 1996; Tamietto & de Gelder, 2010).

In Experiment 1 (as reported in Chapter 3) I used a figure-of-8 coil used to stimulate a more specific area of the cortex than the above studies (which used round coils), and I exploited phosphenes to locate the stimulated region within the visual field. The design of Experiment 1 allowed the (brief) presentation of a body posture image to the precise location of disruption induced by the TMS pulse, or at a control location in the opposite hemifield. The results found that TMS disrupted processing of neutral body posture images presented at the location of the phosphene leaving emotional postures relatively unaffected. The experiment used a range of pulse onsets (70-140 ms following the stimulus), and found the disruption occurred across the whole range (although appearing to diminish at 120-140 ms).

The findings suggested that a single pulse of TMS delivered 70-140 ms following stimulus presentation can disrupt processing of neutral body posture images while leaving the discrimination of emotion in body posture images relatively unaffected. However, this brief effect of a single pulse is quite different to the permanent disruption found in blindsight patients that is permanent. In the experiment reported in this chapter, I used continuous Theta Burst Stimulation (cTBS) to V1 to try to induce a more prolonged level of disruption. cTBS should disrupt both feedforward and feedback processing in V1 and also has several practical advantages: TBS allows for a much simpler design with no varying of pulse onset during the experiment. Combined with the use of a control condition where stimuli are presented on the other side of the screen (as in Experiment 1) it allows the experiment to be run in one session. Also the participant is free from the annoyance/distraction of the pulse during the experimental trials, and retains the freedom to move their head to a greater extent.

The experiment reported here used stimuli of real life faces which were either emotionally neutral or displayed an emotional expression (anger or happiness). Faces were used with the aim of extending the findings of Experiment 1 on body posture images to faces, which are more commonly used in emotion perception research. Using faces also carries the advantage of being able to use a larger set of stimuli than is practical with body posture images, and allows the use of stimulus sets rated for emotional valence is previous research.

The basic idea of the experiment was similar to that of Experiment 1, in that emotion was used as a redundant cue to the discrimination. This time, however, the participant was not taught to discriminate two arbitrary categories, but asked to make a familiar discrimination by performing a gender classification task. Half of the stimuli had neutral expressions, and half emotional (happy or angry). Each emotion was paired (for that participant) with a gender; the emotion thus provided an additional (redundant) cue to the already familiar gender discrimination in the same way as it provided a cue to the newly learned response categories on which the participant was trained in Experiment 1. The specific pairing of genders with emotions was balanced across participants – half of the participants were shown angry male faces and happy female faces, the other half happy male faces and angry female. This design allowed for the direct comparison of the processing of neutral and emotional stimuli whilst the same task was carried out as in Experiment 1, with the advantage there discussed of not having to compare performance across tasks of different difficulty. Also as in Experiment 1, stimuli were presented at the location of a previously reported phosphene, or at the equivalent coordinate in the opposite hemifield.

In a cTBS within-subject design, one usually wants to compare performance during the period when processing is assumed to be disrupted by cTBS to performance without such disruption, in such a way as to avoid a confound with practice, fatigue etc. Although the contrast of disrupted site (localised by the phosphene) to control site in the opposite hemifield (as in Experiment 1) is in fact logically sufficient (if the disruption is effective) without a separate no-TMS control condition, the impact of cTBS on visual processing is not yet well explored. I elected therefore to use a before-during-after design to get an additional measure of the effect of cTBS on visual processing.

Each participant completed three series of discrimination trials: one before the cTBS stimulation, one between 8 and about 22 minutes after the stimulation, and one starting 65 minutes after the stimulation; I will refer to these as the first, second and third times. The first and third times acted as controls so that any differences found in the middle time– when the cTBS should affect processing in V1 – could be confirmed as due to the stimulation and not some other factor(s). More specifically, the second time was predicted to show disruption from the stimulation selective to the phosphene location on screen, and (relatively) selectively so for the neutral faces, as in Experiment 1. It is possible that during the first phase participants would still be learning the emotion-gender association, but this should be established by the end of the first phase. If emotional stimuli became easier only after sufficient experience, then there is still time 3 with which to compare performance.

4.2. Method

4.2.1. Participants

Sixteen participants (five men and eleven women; average age 21, SD = 0) took part in the experiment. All passed TMS safety screening questions before taking part. The experiment lasted around 2 hours, although participants spent a significant proportion of this time waiting between phases of the experiment. Due to this waiting, participants were paid a little less than the standard rate I used in Experiment 1: up to £15 (depending on performance).

4.2.2. Materials

The set of target stimuli used thirty-two identities from the NimStim database of photographs of posed emotional faces (The MacArthur Foundation Research Network on Early Experience and Brain Development, 2002), sixteen male and sixteen female. Each identity occurred once as a neutral face, and once as either a happy or angry face with the emotion contingent on the gender (i.e. happy female faces and angry male faces, or vice versa, with the assignment balanced over participants). The faces were converted to grey scale, and hair cues limited through cropping in Photoshop. Examples of the stimuli can be seen in Figure 4.1 (a). Each stimulus occupied 120 x 168 pixels, i.e. visual angles of 1.66 x 2.23°.

The experiment was run on a PC using *Presentation* (version 14.3, Neurobehavioral Systems), and a 19” CRT screen with resolution of 1024 x 768, set to a 100 Hz refresh rate. Participants sat 70 cm from the screen, with their head resting on a chin rest.

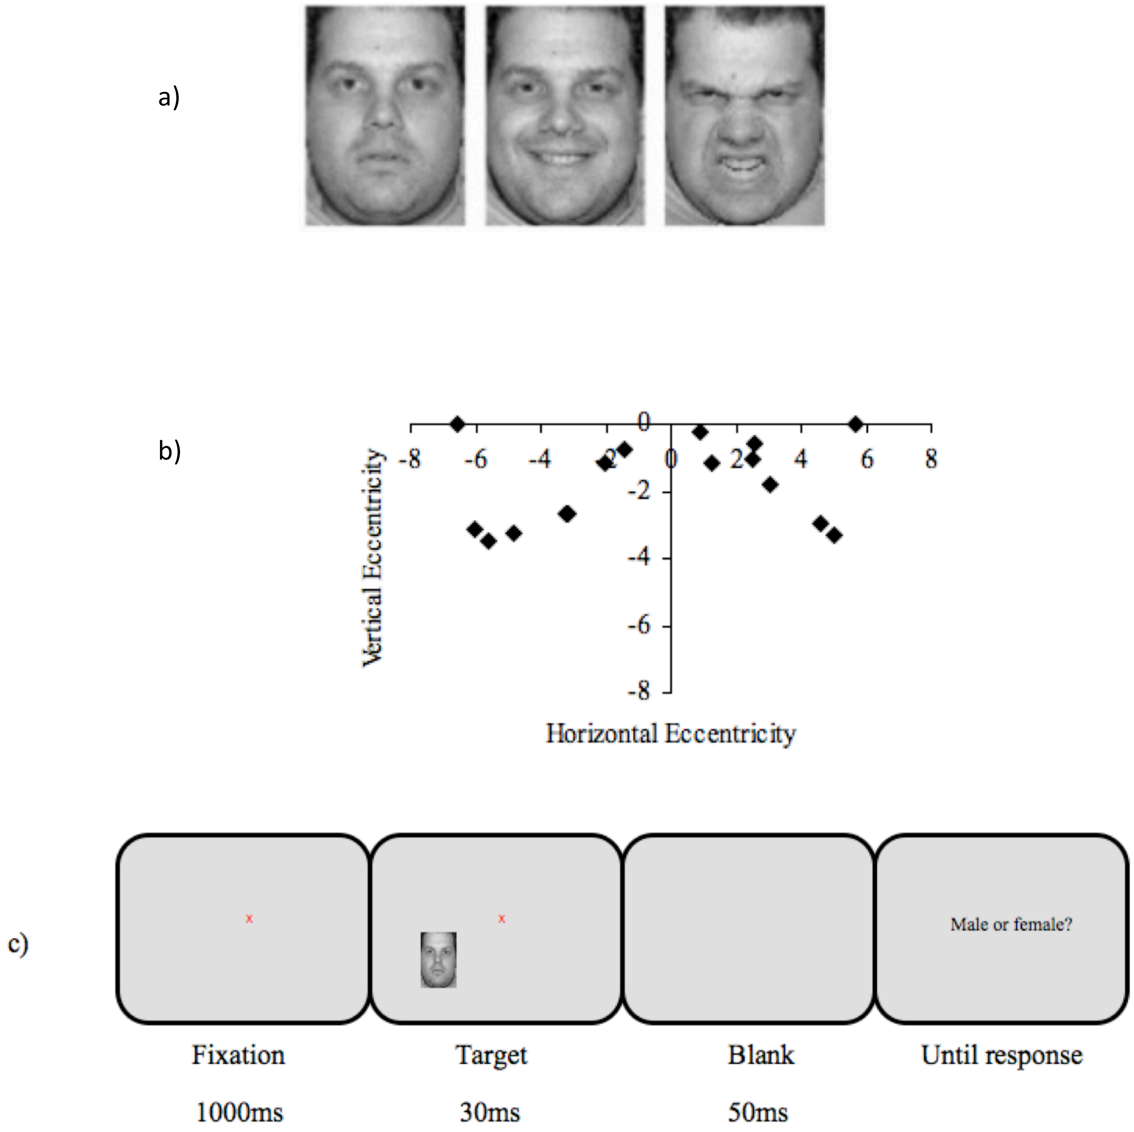

Figure 4.1: Examples of the stimuli used (a), the location of the phosphenes reported by participants (b), and a standard trial outline (c).

4.2.3. Stimulation sites and cTBS

To induce and locate a phosphene a TMS pulse was applied at 70% of stimulator output while the participant fixated a central fixation point on the screen. The coil was held approximately vertically against the back of the head, to the right or left of the midline, with the handle held downwards at an angle of 45° away from the vertical across the midline. To avoid suggestion, very little information was given about what a phosphene would look like, and where in the visual field it was likely to occur. Once a phosphene had been obtained, the stimulation site was adjusted until the phosphene covered a region clearly to the left or right of fixation[4]. Participants moved an example image of a (neutral) face to the location on the screen of the phosphene using arrow keys. The eccentricity of this region varied over participants, but was between 0.92° and 6.76° with a mean of 4.22° (SD = 1.97°). For eight participants the region was on the right side of the screen, and for the other eight participants it was on the left. Individual participants' phosphene locations are plotted in Figure 4.1(b), and were comparable for the left (mean X coordinate = 298, SD = 138; mean Y coordinate = -155, SD = 94) and right hemisphere (mean X coordinate = -231, SD = 127; mean Y coordinate = -100, SD = 89). Once a clear phosphene had been found, the pulse amplitude was varied to find the threshold at which a phosphene was no longer seen – the phosphene threshold (PT).

The motor cortex was found through stimulating at 60% of machine output, until a finger twitch was found. The stimulation output was then varied to find the threshold at which a twitch could be induced (motor threshold – MT). MT was measured as this is the most commonly used baseline for the cTBS output. I based the stimulation levels on PT, but measured MT to check I was not exceeding the usual stimulation intensity (80% of MT).

For twelve participants, the output for the cTBS was based on 65% of PT, and this did not exceed 80% of MT. For the remaining four participants, MT could not be measured (a finger twitch could not be found) and the cTBS output was based on 55% of PT. The mean stimulation intensity across participants was 34% of stimulator output (SD = 3). The cTBS protocol lasted for 40 seconds, and consisted of 3 pulses at 50Hz occurring every 200 ms (the same protocol used by Huang et al., 2005, and Franca et al., 2006). Participants were asked to close their eyes during stimulation to keep the activity levels in V1 relatively constant during stimulation.

4.2.4. Procedure

The discrimination test series was run 3 times, with the first (Time 1) before the cTBS. The cTBS was then given, participants waited 8 minutes, and then the test was run for the second time (Time 2). Participants waited again until 65 minutes had passed from the

---

[4] The location of the phosphene on screen would ideally be in a similar place for each participant. However, it was not possible to achieve comparable locations in all participants and thus some variability is present. This is not too much of a problem given the design is within-subjects, and the locations are comparable for the left and right hemisphere participants.

cTBS, and ran through the test for the last time (Time 3). Following the completion of Time 2 trials, PT and MT were checked again to measure any changes following the cTBS.

An outline of a standard trial from the experiment is shown in Figure 4.1 (c). A face was presented either at the previously determined phosphene location, or the control location on the opposite side of the screen. Participants had to judge the gender of the face and were not explicitly told about the added cue of emotion. Participants completed a short questionnaire following the last test, mainly asking about how they found the experiment, and whether they noticed ‘anything apart from the different genders’ in the faces. If participants said no, the experimenter then asked explicitly if they had noticed the emotions. If they had noticed emotions, they were asked if they found this varied with the gender at all. Of the sixteen participants, six noticed the presence of emotional expressions, and no participant realised the emotion was an additional cue to the decision. In addition to the gender task, on a random one in every 10 trials a digit was presented at the fixation point (instead of a face), and participants were asked to decide whether the number was odd or even. The number was presented briefly (60 ms) and was small so the task was impossible if participants were not fixating. As in Experiment 1, the purpose of the number task was to ensure (and reward) central fixation during the experiment.

Each discrimination test series consisted of 12 blocks of 48 trials: a total of 576 trials, consisting of 64 number trials and 512 face trials, lasting approximately 12-14 minutes. There were 64 face trials for each combination of location on screen x stimulus type (emotional, neutral) x stimulus gender, with each face identity occurring 4 times per condition. In each test series the design completely replicated every 3 blocks to allow the data to be broken down into quarters of the series to examine how any effect developed over time.

Before Time 1, one block of practice (48 trials long) was also given to familiarise participants with the task. After completing the first test series (Time 1), cTBS was applied at the phosphene-inducing location (to the right hemisphere for eight participants, to the left hemisphere for the other eight). After a pause of eight minutes, during which the participants read a book or paper they were asked to bring along, the second test series was given, lasting approximately 12-14 minutes. At the end of the second series phosphene and motor thresholds were again assessed. When 65 minutes had elapsed since the end of cTBS, the third test series was given. The TMS coil was not left in place during this time. After the phosphene was found, the location was marked on a bathing cap for relocation for the cTBS and PT check. The location of the phosphene on screen was the same both of these times. Bonuses were awarded for good performance at the end of each block: one for face discrimination, and one for the number task. Each bonus was awarded if performance exceeded 62% accuracy for each task, and was worth 4p.

4.2.5. Data analysis

The data were analysed as d′, in the same way as reported in Chapter 3. The d′ values were then entered into an ANOVA with the factors of stimulus type (emotional or neutral), stimulus location (phosphene location or control location) and phase of the experiment – before, immediately after, and delayed after stimulation. The critical interaction of interest was between stimulus location, stimulus type, and phase with the prediction that the gender discrimination for neutral but not emotional faces would be affected by the stimulus location on screen, and this would be specific to the time window immediately after stimulation.

An additional analysis was run, averaging performance for each specific stimulus across all conditions and all participants. This was done to ensure no specific stimuli were performed particularly well or poorly, which could add noise and artificially increase or reduce the effects of interest. This analysis can been seen in Appendix 1 (pages 221-222), but even with stimuli that were discriminated particularly poorly removed from the data, the pattern of results is almost identical to the one reported below.

4.3. Results

4.3.1. PT and MT changes

There were some differences between the PT before and after cTBS, but the mean difference was small: a decrease in PT following cTBS = -1.9, SEM=1.43. This effect of the stimulation may have been related to between-subjects factors: the location of the phosphene in the visual field (and so the area of V1 stimulated), the hemisphere stimulated, and whether the stimulation was based on 55% or 65% of PT. An ANOVA was run on the thresholds with the within subjects factor of time of PT (before or after cTBS), and the between subjects factors noted above (with the location of the phosphene included as a continuous covariate). The main effect of time of PT was not significant, F<1, nor did this interact with the distance of the phosphene from fixation, F<1. There was no significant interaction between time and hemifield stimulated, F(1,12)=1.94, nor between time and percent of PT, F<1. Overall, then, there is little sign of any systematic relation between properties of the stimulation and the measured PT change. The MT change was minimal (mean difference = 0.25, SEM=0.35) and was not significant, F<1.

4.3.2. Number task

Performance at the number task averaged 94% (SD=5), with all participants performing well above chance level (minimum=82%). Due to the difficulty of the task, such high performance indicates all participants were fixating during the experiment. In addition, if participants strategically fixated away from the centre, performance at the number task should drop with increasing distance of the face stimuli from the fixation point. The correlation between distance of the phosphene (and so stimuli locations on screen) from fixation and number task performance was not significant, although there was a trend in this direction (Pearson coefficient=-0.22). The high performance at the number task, and only a small and non-significant relationship between distance of the phosphene from fixation and performance at the number task, indicate participants were indeed fixating.

4.3.3. Face discrimination – overall results

Performance at the gender discrimination task was calculated as d'. The results can be seen in Figure 4.2. An ANOVA was run with the within-subjects factors of time, stimulus location, and stimulus type, and the between-subjects factors of version of the experiment (representing the specific pairing of emotion with gender, with version 1 representing angry male and happy female faces, and version 2 happy male and angry female faces). Overall, there was a main effect of time, F(2,28)=4.36, p<0.05, largely captured by the linear trend, F(1,14)=6.39, p<0.05, with performance increasing as time progressed. There was no significant main effect of stimulus location, F<1, indicating there were no overall differences between the locations. The main effect of stimulus type was marginally significant, F(1,14)=4.17, p=0.06, reflecting generally lower performance for the emotional than neutral stimuli (mean difference =0.15, SEM=0.11).

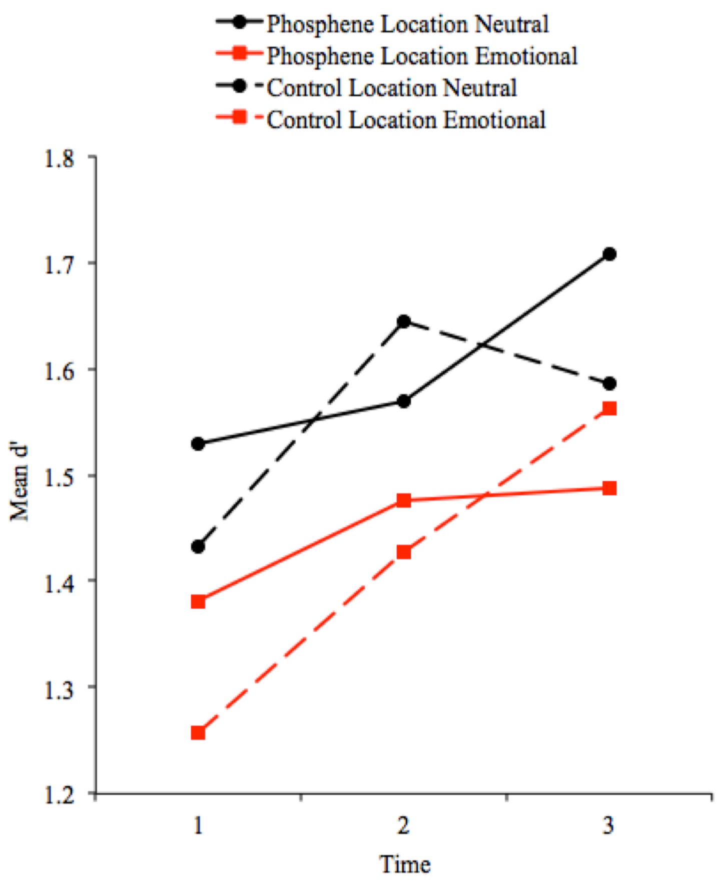

Figure 4.2: Group mean categorisation performance (d') for each time, stimulus type and stimulus location.

Stimulus type also interacted reliably with version of the experiment, F(1,14)=21.15, p<0.001. In Figure 4.3 the results are displayed for the two experiment versions separately. It appears that performance was generally higher for version 1 (mean difference =0.7), although there was no significant main effect of version on performance, F(1,14)=2.94. The highly significant interaction between stimulus type and version of the experiment appears to reflect higher performance for emotional than neutral stimuli in version 1 (mean difference =0.18, SEM=0.1) and higher performance for neutral than emotional stimuli in version 2 (mean difference=0.48, SEM=0.1). The pairing of the redundant cue of emotion with gender, then, in version 1 (male angry, female happy) slightly aided performance at the gender task, but for version 2 the pairing of emotion with gender (male happy, female angry) seemed to more substantially hinder performance at the gender task for emotional stimuli.

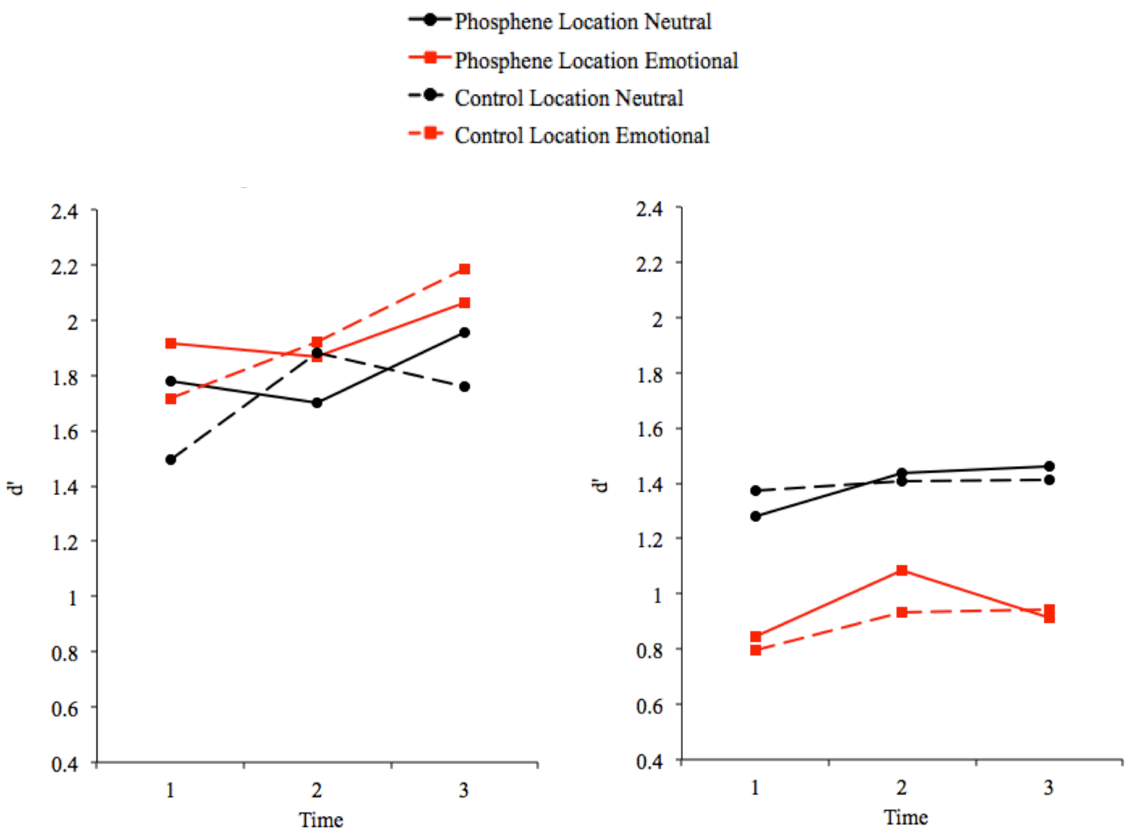

Figure 4.3: Group mean categorisation performance (d') for each stimulus location, stimulus type and time, for version 1 (left) and version 2 (right) of the experiment. Version 1 represents angry male and happy female faces, version 2 happy male and angry female faces.

As the experimental conditions were controlled across sets of 3 blocks (quarters of each test series) it was possible to analyse the development of any effects through each series. The results of this can be seen in Figure 4.4. This shows that performance had stabilised by the last quarter of Time 1. To determine whether any relatively transient effect of the cTBS could be detected within Phase 2, an ANOVA was run on the results from Time 2 including the within subjects factor of quarter. The interaction of interest – quarter x stimulus location x stimulus type was not significant, F<1, and did not reliably interact with version of the experiment, F<1. The interaction between quarter and location was also not significant, F<1, and did not reliably interact with version of the experiment, F<1. Overall, there was no significant main effect of quarter, F<1. Figure 4.4 indicates that there are two quarters for Time 2 that show reduced performance for the neutral faces at the phosphene location: 2 and 4. However, these are not continuous as would be expected if they were meaningful, and were matched with small advantages for the phosphene location for the other two quarters. Overall then there is no convincing evidence for any selective transient effect within Time 2.

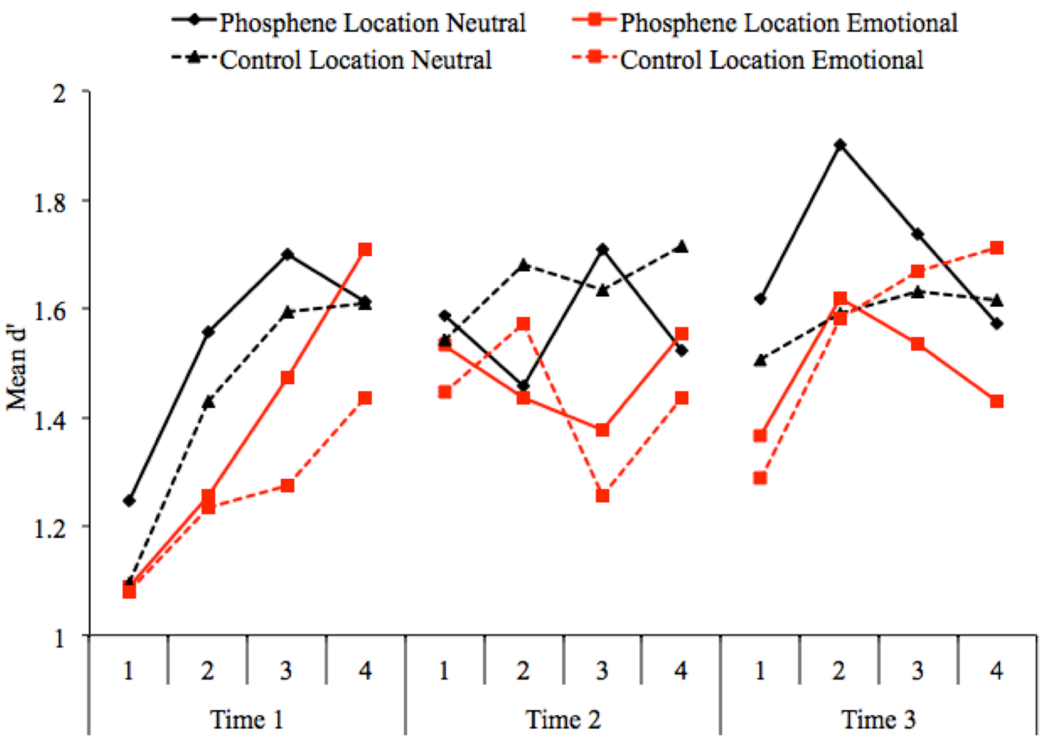

Figure 4.4: Group mean categorisation performance (d') for each of the conditions (stimulus type, stimulus location and time) split across quarters of each time set.

4.3.4. Impact of noticing emotion

As already noted, six participants noticed the presence of emotional expressions. This presents some problems as it leaves two groups of varying size, an issue for ANOVA. The two groups are plotted in Figure 4.5 to give a general measure of the differences. They appear to suggest that participants who noticed the emotional expressions generally performed the gender discrimination better than those who did not (mean difference =1.02) but otherwise the pattern of results appears to be highly similar. Whether or not emotions were noticed, then, did not appear to have much impact on the *pattern* of results.

4.3.5. Emotion specific effects

To check for any overall effect of emotional valence (e.g. differential arousal from angry and happy faces) Figure 4.6 shows the data broken down into the angry (and equivalent neutral) and happy (and equivalent neutral) faces. There did appear to be some differences between the two, namely higher performance for the phosphene location than control

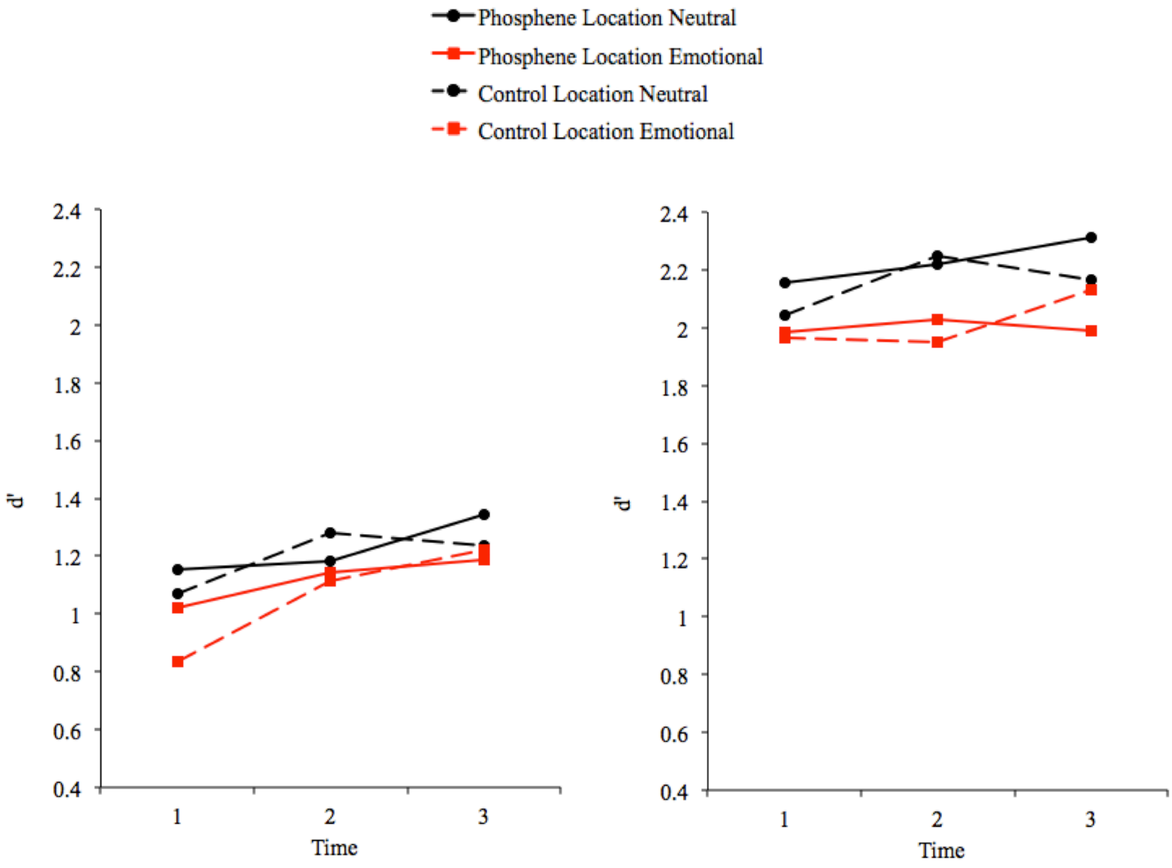

Figure 4.5: Group mean categorisation performance (d') for each condition for participants who noticed the emotional expressions (right) and those who did not (left).

location for the happy faces and neutral equivalents, and generally higher performance at the control location for the angry and neutral equivalent faces. An ANOVA was run on the data, with the within subjects factors of time, stimulus location, stimulus type, and emotion and the between subjects factor of version of the experiment. The main effect of emotion was not significant, F(1,14)=1.29, and did not reliably interact with version of the experiment, F(1,14)=1.34. The interaction between stimulus type and emotion was not significant, F<1 and did not reliably interact with version of the experiment, F<1. The interaction between stimulus location and emotion was not significant, F(1,14)=2.22, and did not reliably interact with version of the experiment, F<1. The interaction between time, stimulus type, stimulus location and emotion was not significant, F(2,28)=1.19, and did not reliably interact with version of the experiment, F<1.

Overall, then, there appeared to be very little effect of the specific emotion on the results.

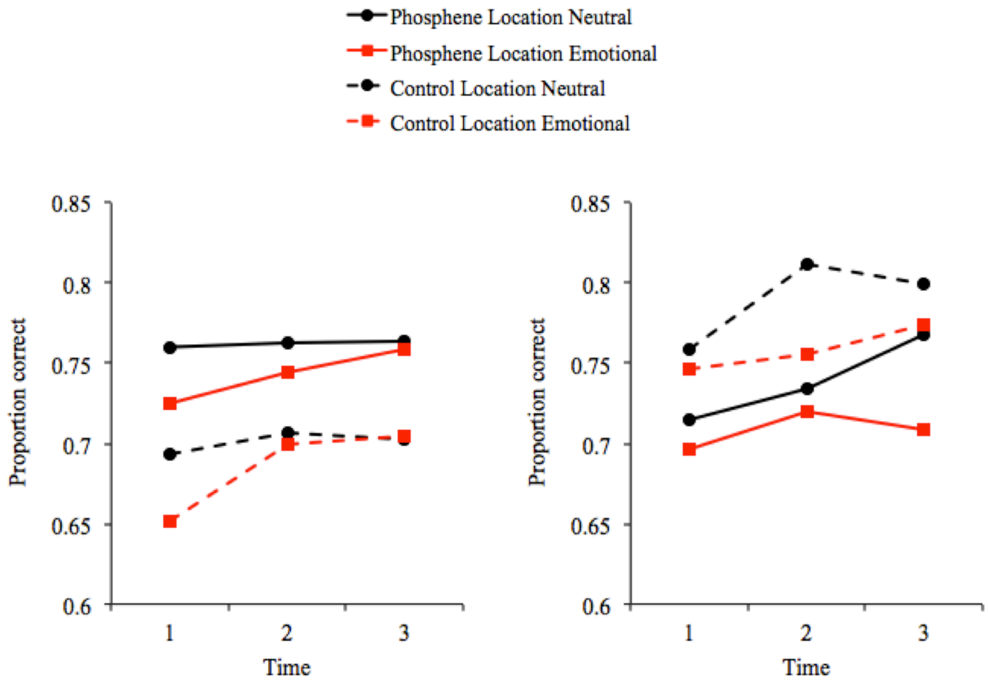

Figure 4.6: Group mean categorisation performance (proportion correct) for each condition for the angry (right) and happy (left) emotions and equivalent neutral images.

4.3.6. Face discrimination: impact of hemisphere stimulated

The hemisphere stimulated was controlled across participants but may have influenced the findings. The results are broken down by hemisphere in Figure 4.7. An ANOVA on the results, including the between subjects factor of hemisphere stimulated, was run. Overall, there was no significant main effect of hemisphere stimulated, F<1. The time x location x version of the experiment x hemisphere stimulated interaction was not significant, F<1. The time x location x type x hemisphere stimulated interaction was not significant F<1, but the 5 way interaction between time, location, stimulus type, version of the experiment and hemisphere stimulated was significant, F(2,28)=5.63, p<0.05. This complex interaction is shown in Figure 4.7. The pattern predicted -- selective depression of performance at time 2 for the neutral stimuli in the phosphene location -- is visible only in the right hemisphere group who completed version 1. However, there is certainly no obvious reason why the effect predicted should be obtained only for this combination of conditions, so it is very unlikely that this interaction is meaningful.

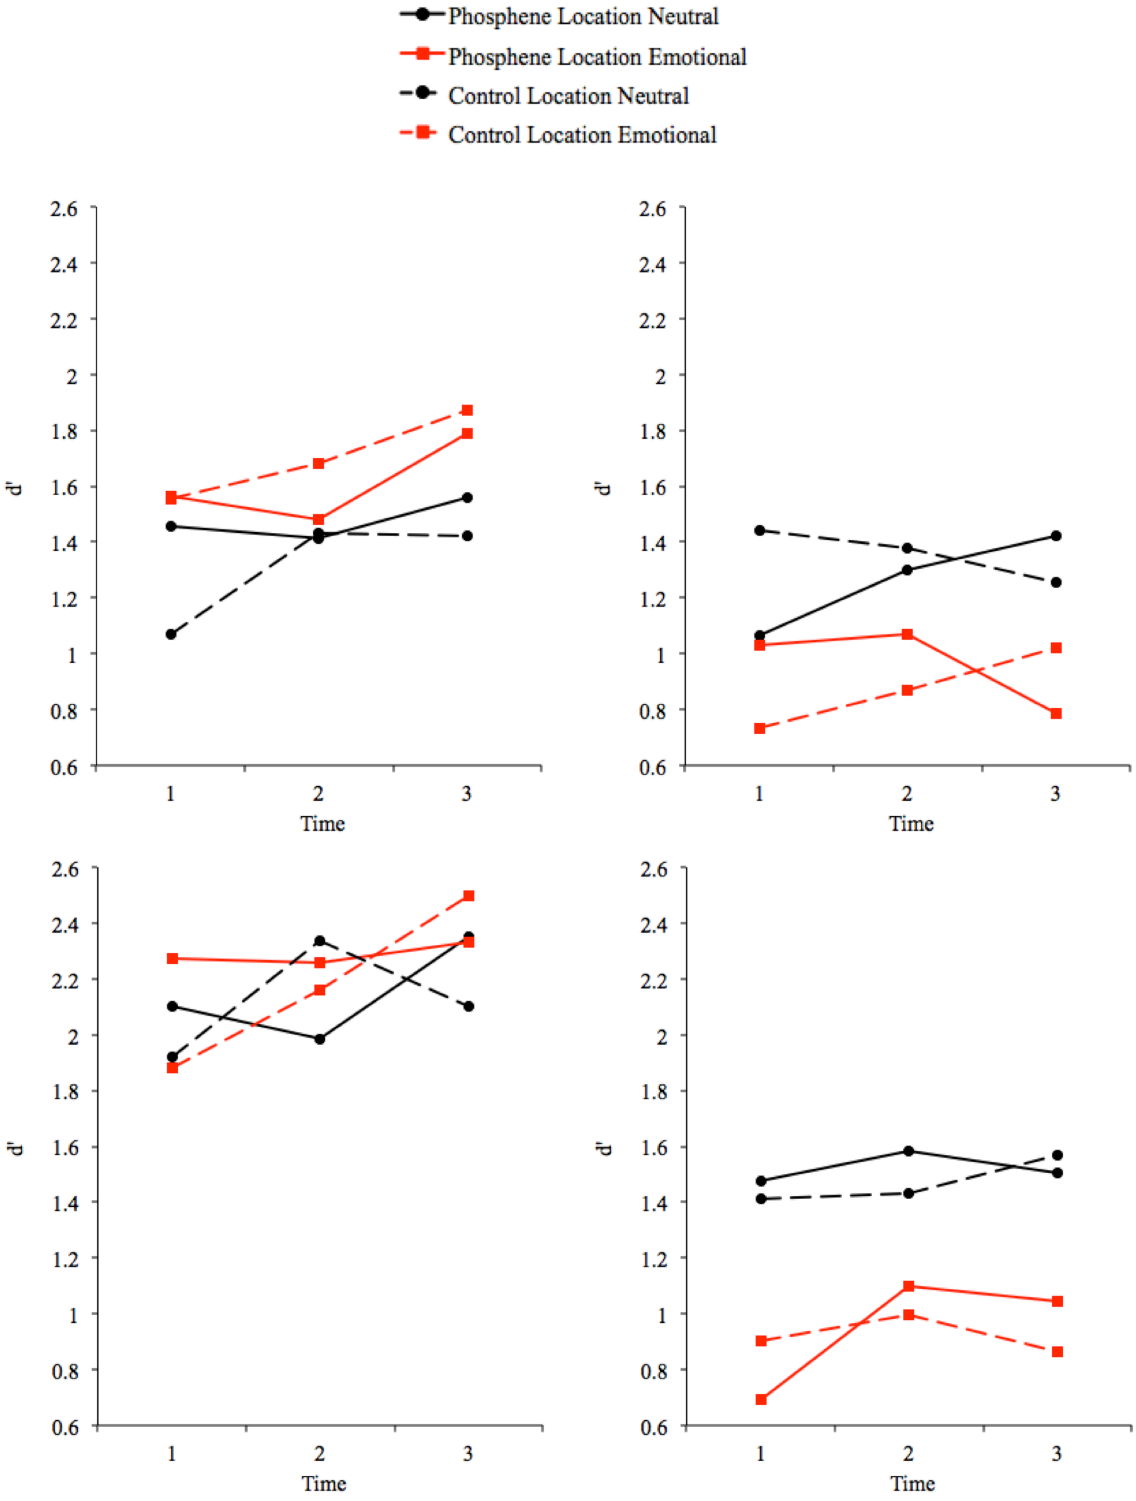

Figure 4.7: Mean categorisation proportion correct (d') for the right (top) and left (bottom) hemisphere participants who completed versions 1 (left) or 2 (right) of the experiment. Each graph represents data from 4 participants.

4.4. Discussion

This experiment aimed to disrupt processing in V1 using cTBS, and investigate the impact on neutral and emotional face processing. The pattern predicted, on the basis both of theory and the result obtained with on-line TMS in Experiment 1, was an impairment of performance after cTBS (i.e. during Time 2) for faces presented at the "fatigued" phosphene location relative to stimuli presented at the control location, but relative sparing from this impairment for stimuli for which emotion provides a redundant cue to the decision. It must be acknowledged that there is very little evidence for such a pattern in the data.

First, there was no clear indication that the cTBS disrupted performance on discriminations at the stimulated location in the visual field. Second, although performance showed a slight benefit from the redundant cue when happy correlated with female and angry with male, it showed a significant cost when the correlation was the opposite way around. Third, if we take mean Time 1 and Time 3 performance as the baseline for each combination of emotion/neutral and stimulus location, there was a suggestion in the overall data of the predicted pattern: relative impairment for neutral stimuli at the stimulated location compared to the other combinations. But further analysis showed that this pattern was seen only for the four participants given version 1 who were also stimulated in the right hemisphere. Not only was this not predicted; there seems no obvious post-hoc account of why the effect should be limited to this group.

Before I consider possible reasons for the absence of the predicted effect, I consider briefly several possible explanations for the significant interactions involving experiment version: i.e. the pairing of emotion with gender category. Firstly, performance was generally higher for participants who completed version 1. This could reflect mere chance – higher performing participants happened to be paired with version 1, and this is the source of the differences. This appears relatively unlikely given the interaction was significant. It is possible that the pairing of emotion with gender in version 1 (angry-male, happy-female) was somewhat more intuitive for version 1 participants, so that the correlation they were being invited to learn implicitly did not conflict with their prior assumptions. For example an angry male face may generally be more of a physical threat (men do after all tend to be larger than women) and, a priori, more associated with threatening emotions than female faces. This is obviously speculation, and other explanations due to simple perceptual properties of the stimuli (i.e. properties of a frown simply made male faces more visible, and a smile made female faces more visible, but the opposite pairings the opposite effect) cannot be ruled out.

As always, null results are very difficult to interpret. There are several reasons why the cTBS manipulation may not have been effective, either in detectably depressing processing in the phosphene location, or in having the hoped-for selective effect on neutral stimuli in that location. First, it may be that this form of stimulation does not robustly disrupt processing in V1. There are mixed findings in the previous literature - Franca et al. (2006) found that the same protocol as used here increased phosphene thresholds (suggesting some form of suppression), whereas Waterston & Pack (2010) found improved discrimination performance of Gabor patterns following the same protocol. These differences in findings might be explained by the difference between lesion effects and noise suppression, as discussed in Chapter 2. If both disruption to processing of visual stimuli, and noise suppression in the visual cortex enhancing performance result from cTBS to V1, it is possible this experiment induced both and the effects essentially cancelled themselves out. Such explanations cannot be ruled out without some measure or manipulation of these possible components.

Alternative reasons for the null effect include the possibility that the coil was placed at a suboptimal angle; although this is unlikely given the same orientation was used as Experiment 1 that did find disruption with single pulses. The stimulation protocol used may have simply been too weak at 65% of PT. That the cTBS had little or no effect on processing in V1 would be consistent with the observation of little (and no significant) change in the phosphene threshold following stimulation. However, as no sham or control stimulation condition was included, it is possible that the stimulation did affect PT, but practice effects, or the presentation of multiple images being flashed in the corresponding area of the visual field, mitigated any change in PT. Given there was very little change in PT at all let alone any significant change, this seems to be unlikely. However, it cannot be ruled out. Another possibility is that the effects of cTBS are more widely distributed in the visual field than the effect of a short pulse train, so that the opposite hemifield was also affected, though the overall lack of a clear dip in performance after stimulation does not support this.

A second possibility to consider is that cTBS to V1 is effective, but has a time course very different to that observed for motor cortex by Huang et al. (2005) – on which the timing of the present experiment was based. I started running the second set of test trials 8 minutes after stimulation. This may have been too early or (more probable) too late. But this seems unlikely due to the lack of any significant effect of quarter of Time 2 – if the experiment was started too early, some evidence of a significant effect in the first quarter of the second test series would be expected or, if started too late, some effect in the last quarter. An effect completely gone within 8 minutes from stimulation, or one that has not begun within 22 minutes from stimulation seems unlikely.

A third possibility is that there may have been a problem not with the stimulation, but with the task. The discrimination of gender may have been relatively robust to TMS disruption due to the specific features present in faces, and/or the visual discrimination required for the gender discrimination. Paired with relative weak effects of cTBS stimulation, the task that is fairly robust to disruption may have led to the (lack of) effect. It may be more difficult to impair the discrimination of faces with TMS, regardless of the task, than to impair the discrimination of the body posture stimuli used in Experiment 1. However, this is fairly unlikely given that the task was not particularly easy in general – participants were not performing at ceiling. A property that that is relatively difficult to discriminate in the task, but robust to TMS interference, seems relatively implausible.

As for the lack of sufficient effects of emotion or emotion advantage in the data, there are also several possible explanations. Participants may simply not have learnt the association between emotion and gender adequately. Unlike Experiment 1, where participants learned a new and arbitrary categorisation, here they performed a very familiar categorisation; the well-learnt associations between facial features or configurations and gender might have blocked learning of the new association of emotion to response category. No participant noticed explicitly that the emotions were an additional cue to the response, and only six noticed the presence of any emotions at all. Or the association may not have been learnt simply because the conditions were not suitable for learning to take place (i.e. the emotions were too difficult to distinguish in the parafoveal part of the visual field).

To summarise, the results from this experiment tell us little about the immunity of emotion processing to disruption of processing in V1. There was little evidence of any disruption to face processing following cTBS to V1, and there was even less evidence of privileged processing of emotions. There are many reasons why the experiment failed to provide evidence of an extension to the results of Experiment 1 using faces and offline cTBS. Regrettably the time available for this thesis project (Experiment 2 was one of the last run) did not permit a further attempt to obtain such evidence. Further research is certainly needed to ascertain the precise effect cTBS to V1 has on visual perception.

CHAPTER 5

Emotions are special: using visual masking to dissocaite processing of neutral and emotional stimuli

---

5.1. Introduction

This chapter reports two experiments that investigated visual awareness and emotion processing using masking paradigms. The experiments reported in the previous two chapters used TMS to disrupt processing in V1, and measured the impact on emotional and neutral body posture image and face discriminations. Some evidence was found that the processing of emotions is relatively spared from the effects of disruption with TMS when compared with the processing of neutral stimuli. The relative sparing of emotional stimuli could be due to incomplete disruption to V1, but it is unclear why emotion processing would be more spared than the processing of visually similar features of neutral stimuli. Alternatively, there could be another route in the brain that allows processing of the emotion of low-visibility stimuli that bypasses V1. A dual route model of emotion processing, distinguishing the geniculo-striate route and a subcortical route, has substantial evidence from the literature (Tamietto & de Gelder, 2010). Whether due to robust processing in visual cortex, or an alternative route(s), there is some evidence that emotional expressions may be special and enjoy privileged processing when compared with equivalent neutral stimuli. The experiments in this chapter took a different approach, using visual masking to disrupt perceptual processing and produce a range of levels of visibility. The main aim was to examine whether the performance of tasks requiring discrimination of facial emotional expression is, at low levels of visibility, privileged relative to performance of equivalent tasks requiring discrimination of other facial properties

Visual masking has been used for decades to investigate visual processing. In particular, it is used to study subliminal, supraliminal, and general low visibility processing, as well as the timing of visual processing and the conditions required for visual awareness. Through the use of varying stimulus, mask, and temporal properties, it is possible to generate a variety of conditions under which the level of processing of the image, and the impact of the image on other behaviours, can be tested. In particular, to suppress visual awareness, backwards masking (where the mask follows the target), and sandwich masking (where masks precede and follow the target), are the most common methods in use.

There are two ways we can assess what information the brain can extract from a briefly presented and masked image. One is to require explicit discrimination of particular properties of the image. The other is to test for the effects of those properties on the processing of a following visible target in a priming design. In this chapter I present two experiments using the former kind of design. In the next chapter I present four experiments using a priming paradigm. To some extent the masking experiments inform the priming experiments by establishing psychometric functions relating discrimination performance to the timing of the mask following the stimulus of interest. Through this the visibility levels in the different conditions were ascertained, important when attempting to use low visibility/subliminal presentation.

There have been very few masking studies comparing processing of emotional and neutral stimuli in two-alternative forced choice tasks (2AFC) when awareness is disrupted. Based on the findings of Experiment 1, where disrupting processing in V1 had more impact on discrimination of neutral than of emotional stimuli (Chapter 3), one might expect discriminations of emotional stimuli to be less affected by a backwards mask than discriminations of neutral stimuli.

The ideal experiment would compare performance in two discrimination tasks using the same stimuli. One task would require discrimination of the emotional valence of the stimuli. The other would require discrimination of an orthogonal, but emotionally neutral property based on the same visual features. The two tasks would be matched in difficulty (and below ceiling performance) for supraliminal presentation. One could then examine the discrimination in both tasks as the visibility is reduced by (e.g.) reducing the stimulus-mask SOA, looking for superior performance on the emotion discrimination at the short SOA.

In practice, this ideal is hard to achieve. Among the factors that need to be taken into account in choosing the non-emotional "control" discrimination are: the general potential confounds of difficulty, meaningfulness, and type of processing required (e.g. featural or configural). The experiments reported in this chapter used faces as target stimuli. Faces were chosen because they are good stimuli for the display of unambiguous emotional expressions with which all participants are very familiar. There is also a considerable amount of research looking at face processing, which may aid selection of the control (neutral) task and help match difficulty levels. For example, there is the well-documented face inversion effect (Yin, 1969). This is where the impact of inverting a face on people's ability to discriminate properties of the face is disproportionably affected by inversion when compared to control stimuli that we are also used to seeing upright, such as houses. This effect is commonly thought to be due to people using configural information (involving the relation of face parts) to process faces, which is heavily disrupted when a face is inverted leaving people to rely on featural processing (Yin, 1969). Inverting faces has been shown to influence (mostly delaying and attenuating) ERP components or effects associated with faces (Eimer & Holmes, 2002; Ashley, Vuilleumier & Swick, 2004). fMRI studies have found differences in activity in the fusiform face area (FFA) between upright and inverted conditions, with the level of activity in the FFA correlating with the size of the behavioural face inversion effect (Yovel & Kanwisher, 2005). Inverting faces, then, influences processing at least in areas of the brain thought to specifically process faces, and appears to dissociate between featural and configural processing.

The following two experiments aimed to compare discriminations of neutral and emotional (properties of) faces when the visibility of the target was varied by changing the stimulus onset asynchrony (SOA) between the target and a backward mask. Two different techniques were used to control for the potential differences between emotional and neutral discriminations apart from affective content. In the first experiment covered by this chapter (Experiment 3) I compared performance on an emotion discrimination task to performance on a control task, attempting to match or at least bracket the difficulty of the emotion discrimination with the control tasks. In Experiment 4 I followed the same strategy as in my TMS experiment (Experiment 2) and used the same non-emotive task throughout, but with an additional (redundant) cue of emotion to the decision.

Experiment 3 asked participants to discriminate between happy and angry emotional expressions, between ages (younger or older than ~40) and between orientations (upright or inverted). Orientation discrimination was used primarily as a baseline index of visibility; I expected it to be easier than emotion discrimination for supraliminal stimuli. Age discrimination was used as a task whose difficulty could in principle be adjusted to match that of the emotion discrimination, and which was similar to emotion discrimination in benefiting from configural processing: it uses a variety of sources of information from a face (e.g. wrinkles around the eyes or mouth, and thinning of the lips) in a similar manner to emotion. Age and emotion discriminations were carried out on upright and inverted faces, with the idea that discrimination of these properties in inverted faces would provide a feature based control task for comparison to the more configural processing of emotion and age in upright faces. All tasks used the same set of faces that varied in emotion and age – i.e. the faces in the orientation block could be angry or happy, young or old.

5.2. Experiment 3

5.2.1. Method

5.2.1.1. Participants

Twelve participants (5 men and 7 women) with a mean age of 22 years (SD=6 years) took part in the experiment. The experiment was split over three sessions of 50 minutes, and participants were paid up to £18 depending on performance.

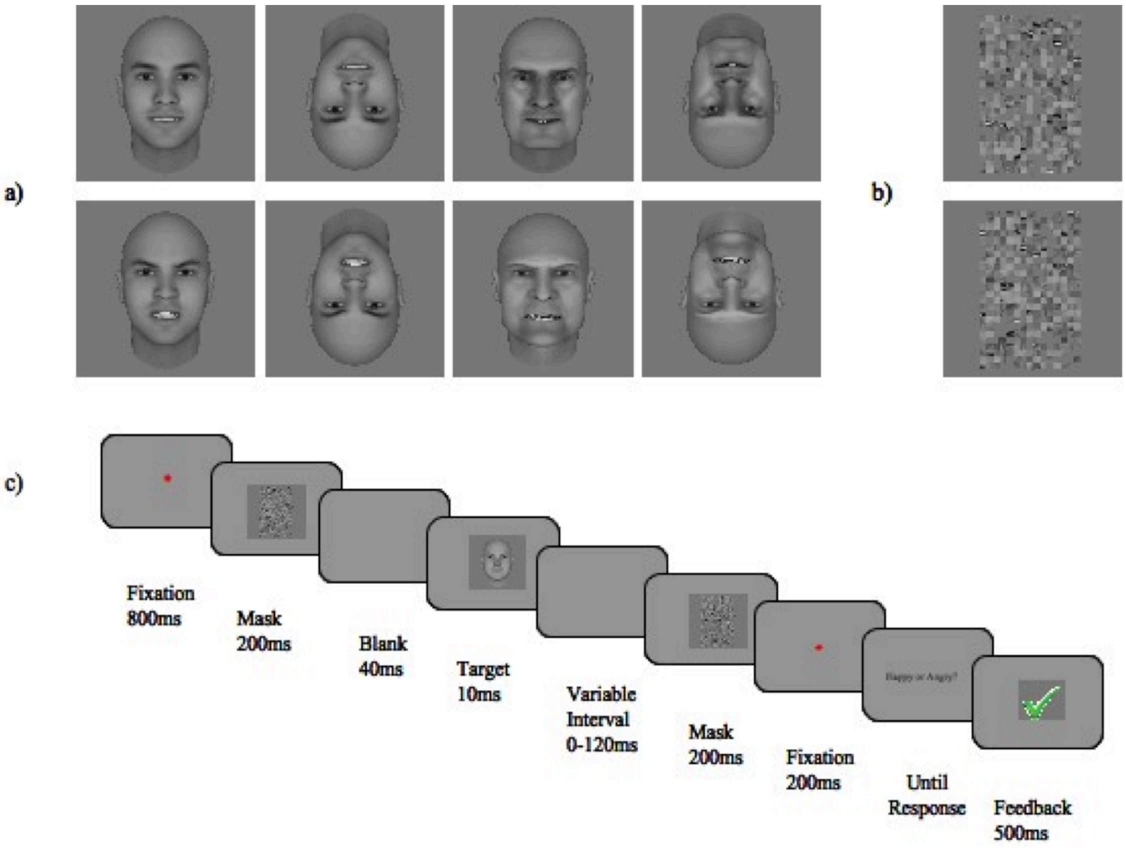

Figure 5.1: Experiment 3: (a) examples of the target stimuli (top happy, bottom angry, left young, right old), (b) examples of mask stimuli used as backwards and forwards masks, (c) standard trial outline.

5.2.1.2. Materials

The stimuli were computer-generated using the FaceGen application (SingularInversions), which allows one to create a facial identity and vary its apparent age and/or emotional expression (happy or angry). Computer-generated faces were chosen over one of the available collections of emotional face stimuli because the discriminability of the images on the age dimension needed to be controlled to try to match the difficulty of the emotion and age discrimination tasks. The age of the faces was varied dynamically on-line (see below) and included five "young" ages in the range 20-30, and five "old" ages in the range 50-60. There were a total of nine different face identities, giving 360 face stimuli in total. Examples of the face and mask stimuli can be seen in Figure 5.1a and 5.1b. The mask stimuli were created by scrambling the target faces using a plug-in for Photoshop (Adobe) that split the image into squares and randomly re-ordered them. Each mask consisted of 360 squares set on a grey background. There were a total of eight masks created, with four occurring as forwards masks and four as backwards masks. Both the face stimuli and the masks were 200x200 pixels in size (2.76 x 2.76°) and presented centrally on the screen throughout. The specific size of the face within each stimulus was 115x180 pixels (2.49 x 1.59°).

The experiment was run using *Presentation* (Version 14.3, Neurobehavioural Systems) with a 19” CRT monitor with resolution 1024x768, set at a refresh rate of 100Hz.

5.2.1.3. Procedure

The experiment consisted of 3 sessions of 50 minutes. Participants performed one of three discrimination tasks (of emotion, age or orientation) in five types of block:

- Discriminating emotion for upright faces,
- Discriminating emotion for upside down faces,
- Discriminating age for upright faces,
- Discriminating age for upside down faces,
- Discriminating face orientation (upright or inverted).

The (repeating) order of these block types varied across participants, as shown in Table 5.1. Of the 12 participants, 3 completed each of the possible block orders.

Block	Order 1	Order 2	Order 3	Order 4
1	Emotion Upright	Age Upright	Age Inverted	Emotion Inverted
2	Age Upright	Age Inverted	Emotion Inverted	Emotion Upright
3	Orientation	Orientation	Orientation	Orientation
4	Emotion Inverted	Emotion Upright	Age Upright	Age Inverted
5	Age Inverted	Emotion Inverted	Emotion Upright	Age Upright

Table 5.1: All 4 of the possible block orders in Experiment 3, varied between participants.

An example of a standard trial can be seen in Figure 5.1 (c). The target face was sandwich masked on each trial, with a constant interval (40ms) between a 200ms forwards mask and a 10 ms (one frame) target display, and a variable interval (10, 20, 30, 40, 50, 60, 70, 90, 110 & 130 ms) between the target and a 200 ms backward mask. On each trial feedback was given following the response in the form of a tick or a cross presented centrally for half a second, then the next trial began immediately. SOA and target type (was presented varied randomly through the experiment with the constraint of four trials per block at the longest SOA. The longest SOA was controlled to provide a sufficient (and equal) number of trials per block for the online adjustment of age discrimination difficulty (see below). The other SOAs remained unconstrained per block to keep most of the experiment fully randomised.

Throughout the experiment performance at the longest SOA (130 ms) was compared for the age and emotion (upright) tasks. The extent of the age gap between the younger and older faces was altered online during the experiment in an attempt to match performance on the age and the emotion tasks. At the end of each round of the five tasks, if the accuracy on the emotion task was higher than on the age task, the age gap was increased by five years in the next round (unless already at the maximum). If the accuracy on the age discrimination performance was higher, the age gap decreased (unless already at the minimum). At the start of the experiment the ages 25 and 55 were used. However, all participants ended up moving rapidly to the easiest discrimination available (20 and 60) by half way through the experiment, indicating the age task was generally very difficult.

In total each of the three sessions consisted of 25 blocks of 40 trials (3000 trials in total across the experiment) — five of each type of block. For each combination of block type, SOA, and response category (i.e. emotion, age or orientation category depending on the task) there were 30 trials. To control the frequency with which each specific stimulus was presented, each session of the experiment used three of the nine face identities. Assignment of identities to session of the experiment was controlled across participants. For each condition (e.g. emotional expression, age, orientation) in each session, one identity was presented once only, and the other two twice. Which identities were presented once or twice per stimulus type varied randomly for each condition and session and between participants.

At the end of each cycle through the block types (i.e. every five blocks), participants were awarded the bonus, worth 20p, if their average performance for the five blocks was above 55%.

5.2.1.4. *Data analysis*

The data from this experiment were analysed as per cent correct for each condition of the experiment. d' values were then calculated as previously described. For the inversion task, H is p(correct) for upright faces, and FA is p(error) for inverted faces; for the emotion task, H is p(correct) for angry faces, FA is p(error) for happy faces. ANOVAs were run to assess the impact of block type and SOA on performance.

5.2.2. Results

Discrimination performance, calculated as d', is shown as a function of SOA in Figure 5.2. An omnibus ANOVA was run with the within-subjects factors of block type and SOA. Overall, there were main effects for block type, $F(4, 44) = 8.76$, $p < 0.01$, and SOA, $F(9, 99) = 10.97$, $p < 0.001$. There was also a significant interaction between block type and SOA, $F(36, 396) = 4.61$, $p < 0.001$. Figure 5.2 a shows the relationship between block type and SOA: all block types showed some evidence of performance increasing with SOA, although this was strongest for the orientation blocks and was minimal for the age blocks. Performance for the orientation blocks was considerably higher overall than for the emotion and age blocks, although at the shortest SOAs the emotion tasks were performed better than the orientation and age tasks. This emotional advantage at the shortest SOA is highlighted by Figure 5.2 b, which represents discrimination performance for the emotion tasks and for the inversion task at 10 ms SOA.

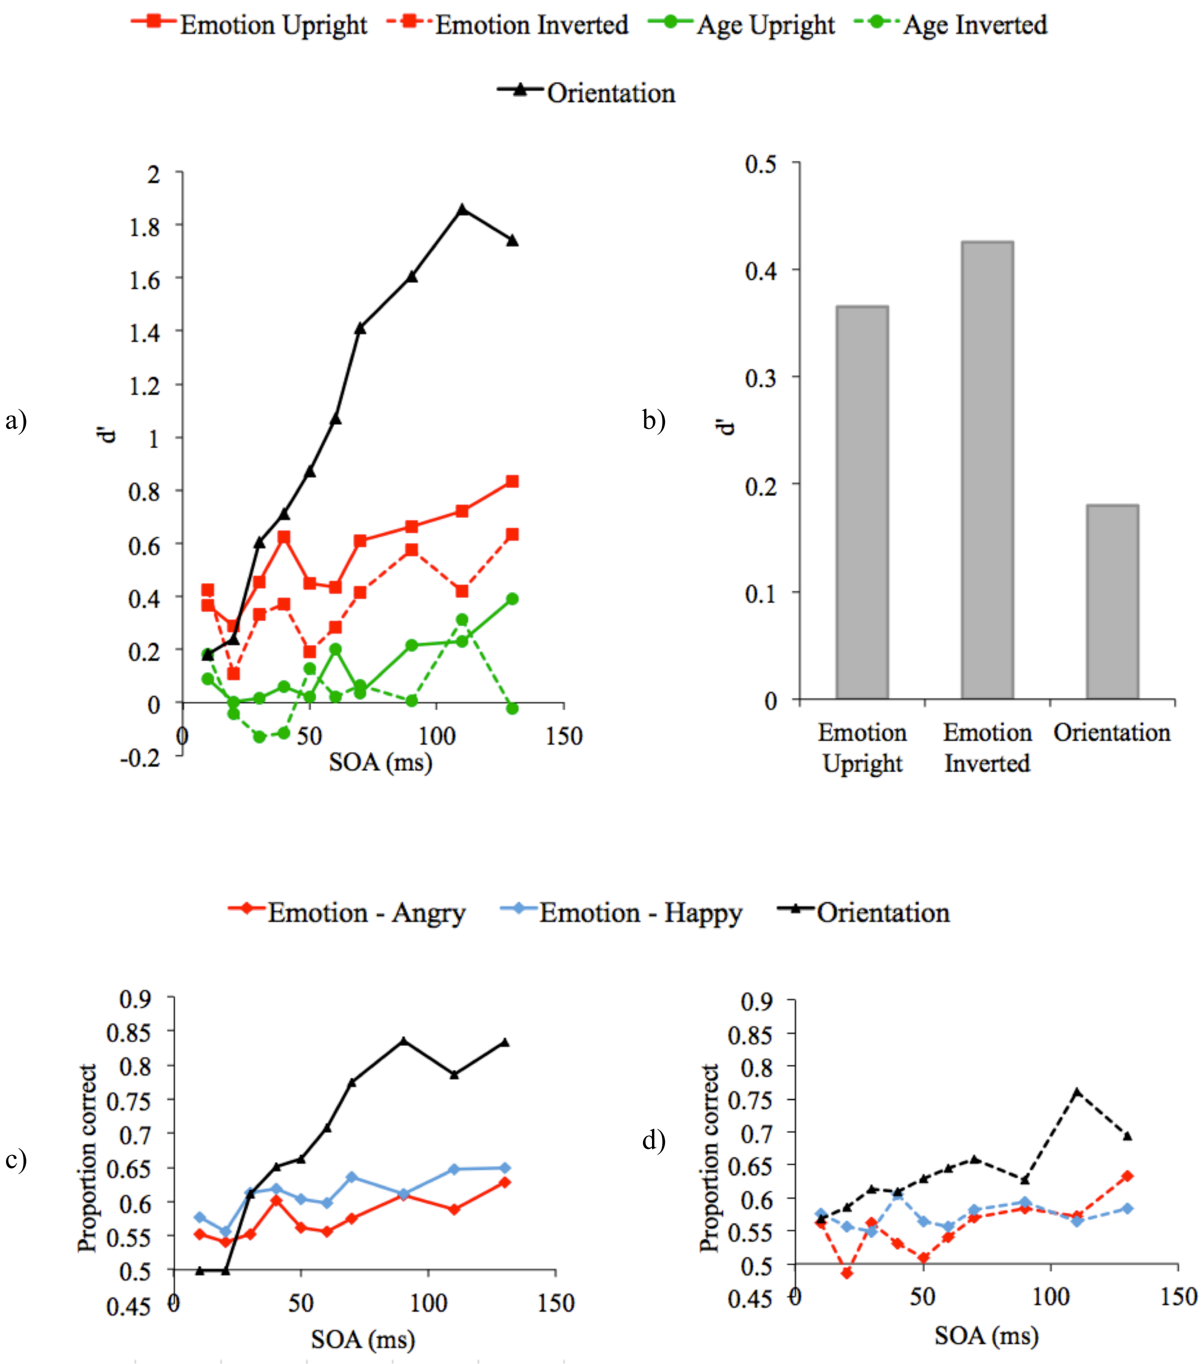

Figure 5.2: Group mean discrimination accuracy for Experiment 3, for all 5 tasks at all SOAs (a), the emotion and orientation tasks at 10ms SOA (b), and the emotion and orientation tasks broken down by emotion type (upright trials – (c), inverted trials – d)).

Task	SOA	Inversion	SOA*Inversion
Emotion	7.96*	6.35*	<1
Age	5.39*	2.94	1.75
Orientation	72.11**		

Table 5.2: F values for the ANOVA's run on each task in Experiment 3, with the factors of SOA (and inversion for emotion and age tasks). * indicates $p<0.05$, ** indicates $p<0.001$.

Separate ANOVAs were also run for each of the three tasks (emotion, age, and orientation discrimination) with the factor of SOA, and for the emotion and age tasks, orientation (upright versus inverted). As Table 5.2 shows, all three tasks show a significant increase in performance with increasing SOA. There was also a significant main effect of inversion for the emotion task, indicating better discrimination performance for upright than inverted faces, but the lack of a significant interaction with SOA indicates no detectable difference in the slope of the function. There was no significant impact of inversion or its interaction with SOA for the age task. Whilst inversion did appear to influence the discrimination of emotion, then, the effect was small and the inverted and upright emotional faces showed similar changes in performance with increasing SOA.

The feature of the results of most interest is the predicted advantage for the emotion discrimination task at the lowest level of visibility. Despite the orientation task being performed much better at longer SOAs, at a SOA of 10 ms, performance on the discrimination of emotion (for both upright and inverted faces) exceeded performance on orientation discrimination. Performance of the orientation and emotion tasks is shown separately for upright and inverted faces in Figure 5.2 (c and d). Planned comparisons between emotion and orientation discrimination performance at the shortest SOA were performed. The difference between the orientation and emotion tasks was not significant for upright faces, $t\ (11) = 1.04$ ($M = 0.19$, $SEM = 0.18$), and was marginally significant for inverted faces, $t\ (11) = 1.46$, $p = 0.09$ ($M = 0.25$, $SEM = 0.17$). The t-test comparing the differences between the emotion and the orientation task, with emotion performance averaged across the upright and inverted conditions, was marginally significant, $t\ (11) = 1.32$, $p = 0.1$. The difference between emotion performance for the upright and inverted blocks was not significant, t (11) = -.05 (M = 0.06, SEM = 0.12). As the expected finding was for the inversion task to be at chance, and the emotional discrimination above chance, one-way t-tests were run on the two emotional discriminations (upright and inverted) and the inversion task at the SOA of 10 ms. For the emotion upright, performance was reliably above chance, t (11) = 1.86, $p < 0.05$, as was performance for the inverted emotional discrimination, t (11) = 2.18, $p < 0.05$. For the inversion task, performance was not reliably above chance although it was marginal, t (11) = 1.42, p = 0.09.

The data were also analysed separately for each emotion, at each orientation. Here proportion correct was used as the measure of performance as d' cannot be estimated for just one response alternative. (Hence this analysis is vulnerable to the possibility of differential response biases in the two tasks.) The emotion specific proportion correct can be seen in Figure 5.2 c and d, with each emotion compared with the relevant measure of visibility – upright face trials for the orientation task were compared with the upright emotion condition, and inverted face trials for the orientation task were compared with the inverted emotion condition. The results indicate that for the upright faces in the orientation task with an SOA of 10 ms, performance (proportion correct M = 0.5, SEM = 0.03) did not differ reliably from chance, t (11) = -0.03. At this timing, the difference in performance for the angry upright faces and the orientation task was not significant, t (11) = 1.09 (mean difference = 0.07, SEM = 0.05). The difference in performance for the happy upright faces and the orientation performance was significant, t (11) = 2.07, $p < 0.05$ (mean difference = 0.08, SEM = 0.04). The difference between angry upright and happy upright faces was not significant, t (11) = 0.49 (mean difference = 0.02, SEM = 0.05). For the inverted trials, performance at the orientation task 10 ms SOA was 0.57 (SEM = 0.04), which was not reliably above chance, t (11) = 1.59. Performance for angry and happy faces in the emotion-inverted task was almost identical to the orientation task at 10 ms SOA (0.56 and 0.58 respectively).

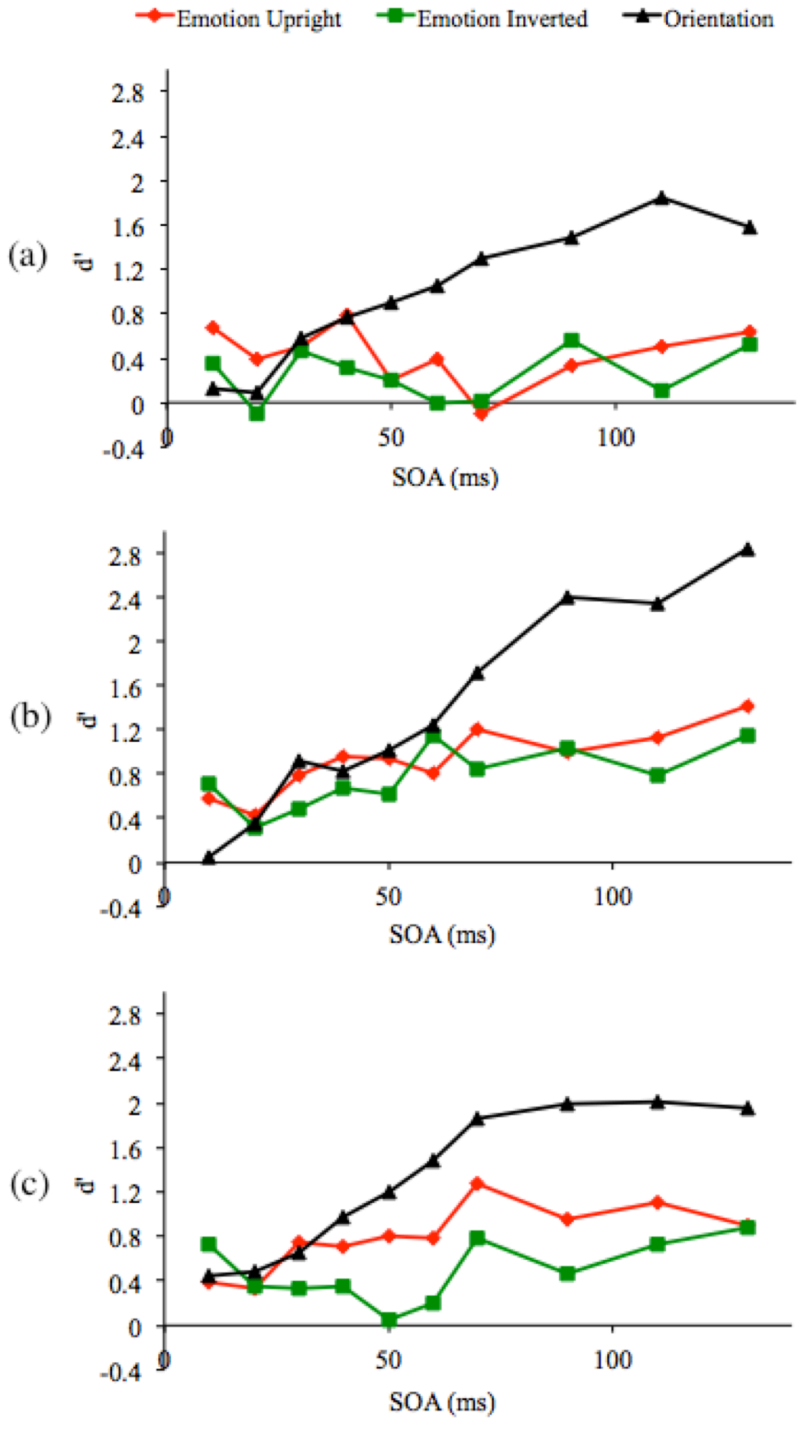

Figure 5.3: Group mean discrimination from Experiment 3 for the emotion and orientation tasks per session run, with (a) showing session one, (b) session two, and (c) session three.

The size of the emotion advantage at 10 ms was certainly small, and of borderline reliability. A further analysis examined its consistency with practice over the three sessions of the experiment. The results can be seen in Figure 5.3. Performance at the orientation task at SOA = 10 was at chance in the first and second sessions, but then improved. Performance was better than chance on the emotion task at SOA = 10 in all three sessions, but the effect of particular interest – better discrimination of emotion than of orientation — is most dominant in the first and second sessions. An ANOVA on the data with the factors of session, task (orientation, emotion upright and emotion inverted) and SOA for the times 10 – 40 ms (chosen as an emotion advantage was expected at the shortest SOAs) found that there was no significant impact of session overall, F < 1, and session did not interact reliably with SOA, F < 1, indicating that session did not materially alter the visibility of the stimuli. There was no significant interaction between session and task, F < 1, and no significant interaction between session, task and SOA, F < 1. The impact of session, then, is not significant despite a trend being apparent for larger differences between the emotion and orientation tasks for the earlier sessions.

5.2.3. Discussion

The results from Experiment 3 indicate that participants found the age discrimination (both for upright and inverted faces) very difficult, with performance lower than the emotion task in spite of the dynamic adjustment of the age contrast to try to match them. Clearly the age discrimination task in this experiment, even adjusted dynamically, was not a good control for the emotion discrimination task: all participants ended up with the maximum age difference available, but it was still much harder than the emotion discrimination.

The advantage for the emotion task over the orientation task at the shortest SOA provides some evidence for emotion processing undergoing less disruption with masking than the other tasks. Orientation discrimination was much the easier task at the longer SOAs, so to see higher performance on the emotion discrimination at the shortest SOA is highly suggestive of an emotional advantage under difficult viewing conditions. The difference between orientation and emotion at 10 ms was not significant overall, although there was a significant effect on accuracy, uncorrected for bias, for the happy upright faces, and there was a suggestion of larger differences between orientation and emotion at 10 ms in the first and second sessions of the experiment. Whilst the emotion advantage at 10 ms is not huge, in light of the difficulty of the task at higher visibility levels it is suggestive.

The relation between emotion discrimination of upright and inverted emotion faces, as a function of SOA, suggests visibility may affect the *type* of processing that decisions are based on – featural or configural. At longer SOAs performance for upright faces is generally superior, but at the shortest SOAs performance is similar for upright and inverted faces. This difference suggests that whilst performance at the higher levels of visibility may operate on a configural (or holistic) level, at the lowest levels of visibility performance may actually be based on features.

The emotion advantage at 10 ms was clearly apparent only in the first and second session of the experiment, diminishing in the last session due to improving performance at the orientation task. Although this effect of session was not significant, the pattern argues against emotion specific effects emerging only with practice. It is consistent rather with the idea that the emotion discrimination at the shortest SOA was relatively automatic, and based on a rather crude emotion detector that is either innate or develops with the (hefty) experience all participants would have had with emotional faces. In contrast performance for the orientation task did seem to increase with practice. It cannot be ruled out that practice played some role in performance, each session consisted of 1000 trials and substantial learning could take place.

To summarise, this experiment showed a clear and approximately linear masking function in all 5 tasks. This was particularly the case for the orientation task, which was performed poorly at the shortest SOA but rose rapidly with increasing SOA. There was some advantage for emotion processing at the lowest levels of visibility. The age control task did not work well as a control task, being performed very poorly throughout. A general configural processing explanation of the emotion advantages seems unlikely given the lack of difference between upright and inverted faces in the emotion task. The most straightforward explanation is that, at the shortest mask SOA, a relatively automatic feature-based emotion discrimination process is still activated by faces whose visibility is so poor that upright and inverted faces cannot be discriminated.

5.3. Experiment 4

The next experiment aimed to match emotional and neutral discrimination tasks better, and replicate an emotional advantage at low levels of visibility. To this end, a different approach to matching neutral and emotional task difficulty was used: using the same strategy as in Experiment 2 (Chapter 4)[5]. Just one task was performed throughout the experiment – gender discrimination. Half of the male and half of the female faces were emotionally neutral, while the other half had emotional expressions with the emotion providing an additional (redundant) cue to the decision (i.e. male angry faces and female happy faces). This allows for the direct comparison of the neutral and emotional faces, and the testing of the hypothesis that there will be an emotional advantage at the lower levels of visibility. As the discriminability of the response alternatives did not need to be manipulated in the same way as in Experiment 3's age task, an established set of real life faces (NimStim) was used instead of the computerised faces of Experiment 3. The main advantages to the use of an established set of stimuli includes the availability of a larger sample of identities, with more dissimilar faces, and more lifelike features and expressions, that have been tested and rated in previous research.

5.3.1. Methods

5.3.1.1. Participants

Sixteen psychology students took part for course credits. A small financial bonus (up to £2) was also paid depending on performance. The participant's average age was 21 years (standard deviation of 3 years), with four men and twelve women. Participants came in for two sessions of 45 minutes on different days.

5.3.1.2. Materials

Thirty-two different identities were taken from the NimStim face set for the target faces, sixteen male and sixteen female. Examples of the stimuli can be seen in Figure 5.4. The NimStim set is of colour photographs, but these were transformed to gray scale for the present experiment, and most of the surrounding hair removed through cropping to reduce simple featural cues to gender. Each face identity came in six

---

[5] It may appear odd that this experiment used the same design strategy was Experiment 2, which did not work. In fact, this experiment and Experiment 2 were run simultaneously, so the knowledge of its outcome was not available to guide the design of this one. Also, Experiment 2 appeared to fail to disrupt neural processing, and hence whether any emotion advantage could be found with this design was moot.

versions: a neutral, happy or an angry expression, upright or inverted. Eight mask images were created from the target faces in the same manner as Experiment 3, and consisted of 864 randomly reordered squares. The stimuli were 120 x 168 pixels in size (1.66 x 2.32°).

Figure 5.4: Examples of the stimuli used for the target faces (a), showing neutral (left), happy (middle) and angry (right) expressions, and examples of the mask stimuli (b), used in Experiment 4.

5.3.1.3. Procedure

The trial outline was essentially the same as Experiment 3. The SOA between the target and backwards mask varied between 10 and 130 ms (10, 20, 30, 50, 90, or 130 ms). There was only one task: throughout the experiment participants were asked to judge the gender of the faces, and the word ‘emotion’ was not used at any point. There were blocks where the faces were upright, and blocks where they were inverted. The block type alternated through the experiment, and participants were told at the start of each block which orientation would be present. Participants were given feedback on their response after each trial (a tick or a cross). All stimuli were presented centrally on the screen.

On half the trials, the stimulus displayed a happy or angry emotional expression as a redundant cue to the gender decision: on the other half, the stimulus had neutral expressions, with no extra cue to the decision. For eight participants happy faces were female, and male angry. For the other eight male faces were happy and female angry. Emotional and neutral stimuli were intermixed randomly throughout the experiment. Participants were not informed of the presence of emotional expressions, and were told only to discriminate between the male and female faces. Seven participants noticed the presence of emotional faces during the experiment, but only two of these picked up that emotions were an additional cue to the gender discrimination. The remaining nine participants did not notice the emotions at all and expressed surprise when told. As over half of the participants did not notice the presence of emotional expressions, this suggested that the stimuli were generally hard to see.

The experiment consisted of two sessions of 24 blocks (32 trials in each), with a total of 1536 trials in the whole experiment. Each combination of stimulus type (emotional or neutral), gender, orientation, and SOA occurred 32 times. For each SOA, each identity occurred once as a neutral face, and once as an emotional face. At the end of each block participants were informed that they had been awarded the bonus of 4p if performance was over 60% for the first two blocks, or if, thereafter, they had improved on their performance over the previous blocks (i.e. performing higher than the previous overall average) with improvement compared for each face orientation separately.

5.3.1.4. *Data analysis*

The data from this experiment were analysed as per cent correct for each condition of the experiment. For the emotional and neutral faces separately, d' values were then calculated as previously reported, where H is p(correct) for male faces, and FA is p(error) for female faces. ANOVAs were run to assess the impact of inversion, emotion, and SOA on performance.

5.3.2. Results

The impact on d' of SOA and the presence/absence of the redundant emotion cue are shown in Figure 5.5 for the upright and inverted face blocks. Discrimination was essentially at chance for the shortest 2 or 3 SOAs and then increased monotonically with SOA. The results were analysed with an ANOVA with orientation, stimuli type (emotional or neutral) and SOA as within-subject factors and version of the experiment (which emotion was assigned to which stimulus gender) as a between subjects factor. There was a main effect of inversion, F (1, 14) = 11.38, $p < 0.001$, with overall performance much higher for upright than inverted faces when there was any discrimination to speak of (mean difference d' = 0.19, SEM = 0.06); this did not interact reliably with version of the experiment, F (1, 14) = 3.05. There was a significant main effect of SOA, F (5, 70) = 16.27, $p < 0.001$, reflecting increasing performance with increasing SOA (as shown in Figure 5.5), and this did not interact reliably with version of the experiment, F < 1. SOA interacted significantly with inversion, F (5, 70) = 15.34, $p < 0.001$, suggesting the masking function differed between upright and inverted faces; evidently the function climbs more steeply for upright faces. The interaction between SOA, inversion and version of the experiment was marginally significant, F (5, 70) = 2.12, p = 0.07. There was no significant main effect of stimulus type, F < 1 (mean difference = 0.01, SEM = 0.04), and stimulus type did not reliably interact with version of the experiment, F < 1. The interaction between stimulus type and orientation was not significant, F (1, 14) = 2.49, and did not interact reliably with version of the experiment, F < 1. There was a significant interaction

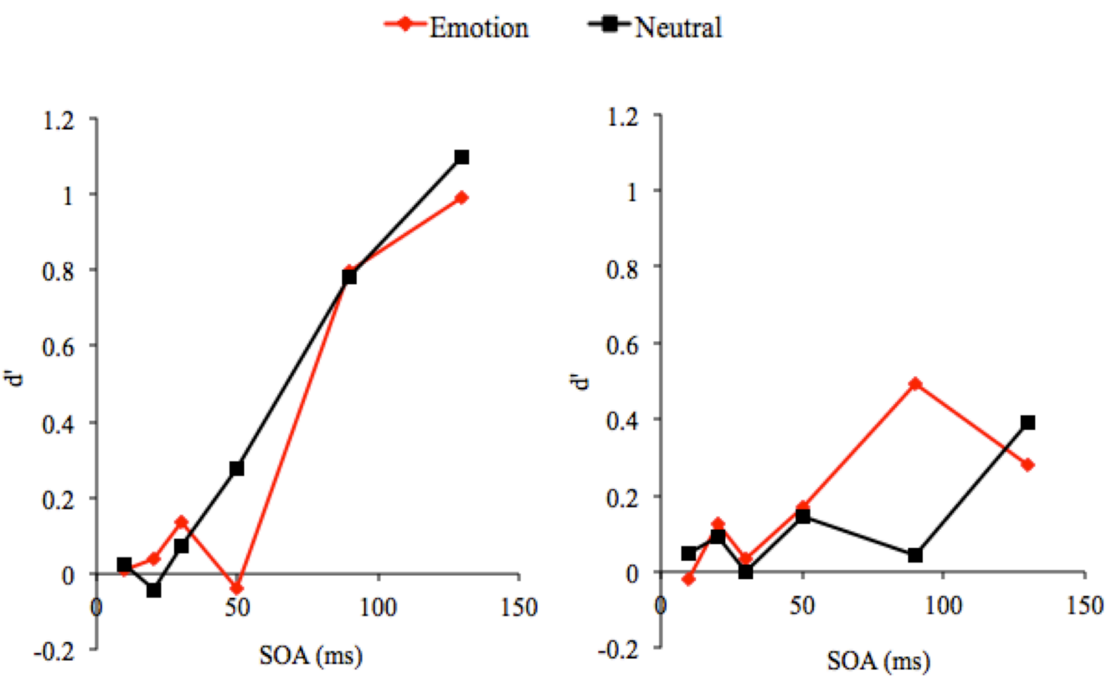

Figure 5.5: Group mean categorisation performance (d') for Experiment 4, for the emotional and neutral faces, for upright (left) and inverted (right) face blocks.

between stimulus type and SOA, F (5, 70) = 2.87, $p < 0.05$, and this did not interact with version of the experiment, F (5, 70) = 1.46. The interaction between stimulus type, inversion and SOA was marginally significant F (5, 70) = 1.95, p = 0.1, and did not reliably interact with version of the experiment, F (5, 70) = 1.81. The version of the experiment had no significant main effect, F<1.

As there was evidence in the main ANOVA for differences between the two face orientations, the data were analysed separately for each orientation to assess the impact of the conditions for upright and inverted faces in more detail. For the upright faces, there was no main effect of stimulus type, F < 1, no interaction with version of the experiment, F < 1. There was a significant main effect of SOA, F (5, 70) = 24.13, $p < 0.001$, and this did not interact reliably with version of the experiment, F < 1. The interaction between stimulus type and SOA was marginally significant, F (5, 70) = 2.02, p = 0.09, and did not interact reliably with version of the experiment, F (5, 70) = 1.76. There was no main effect of version of the experiment, F (1, 14) = 1.81. The marginally significant interaction between stimulus type and SOA most likely reflects the lower performance for the emotional than neutral stimuli at 50 ms SOA (mean difference = 0.31, SEM = 0.1), shown in Figure 5.5.

For the inverted faces, there was no main effect of stimulus type, F (1, 14) = 1.21, or interaction with version of the experiment, F < 1. There was a significant main effect of SOA, F (5, 70) = 3.2, $p < 0.05$, and this did not interact reliably with version of the experiment, F < 1. The interaction between stimulus type and SOA was significant, F (5, 70) = 2.72, $p < 0.05$, and did not interact reliably with version of the experiment, F (5, 70) = 1.54. There was no significant main effect of version of the experiment, F < 1. From Figure 5.5 there appears to be only one point at which emotions are performed considerably better than neutral stimuli: 90 ms SOA for the inverted faces. An emotional advantage was not predicted at this time, so a post-hoc analysis was run (Bonferroni corrected for the 6 possible different SOA comparisons). The difference between emotional and neutral performance for the inverted faces at 90 ms was significant with this correction, t (15) = 4.03, $p < 0.01$ (mean difference = 0.45, SEM = 0.11).

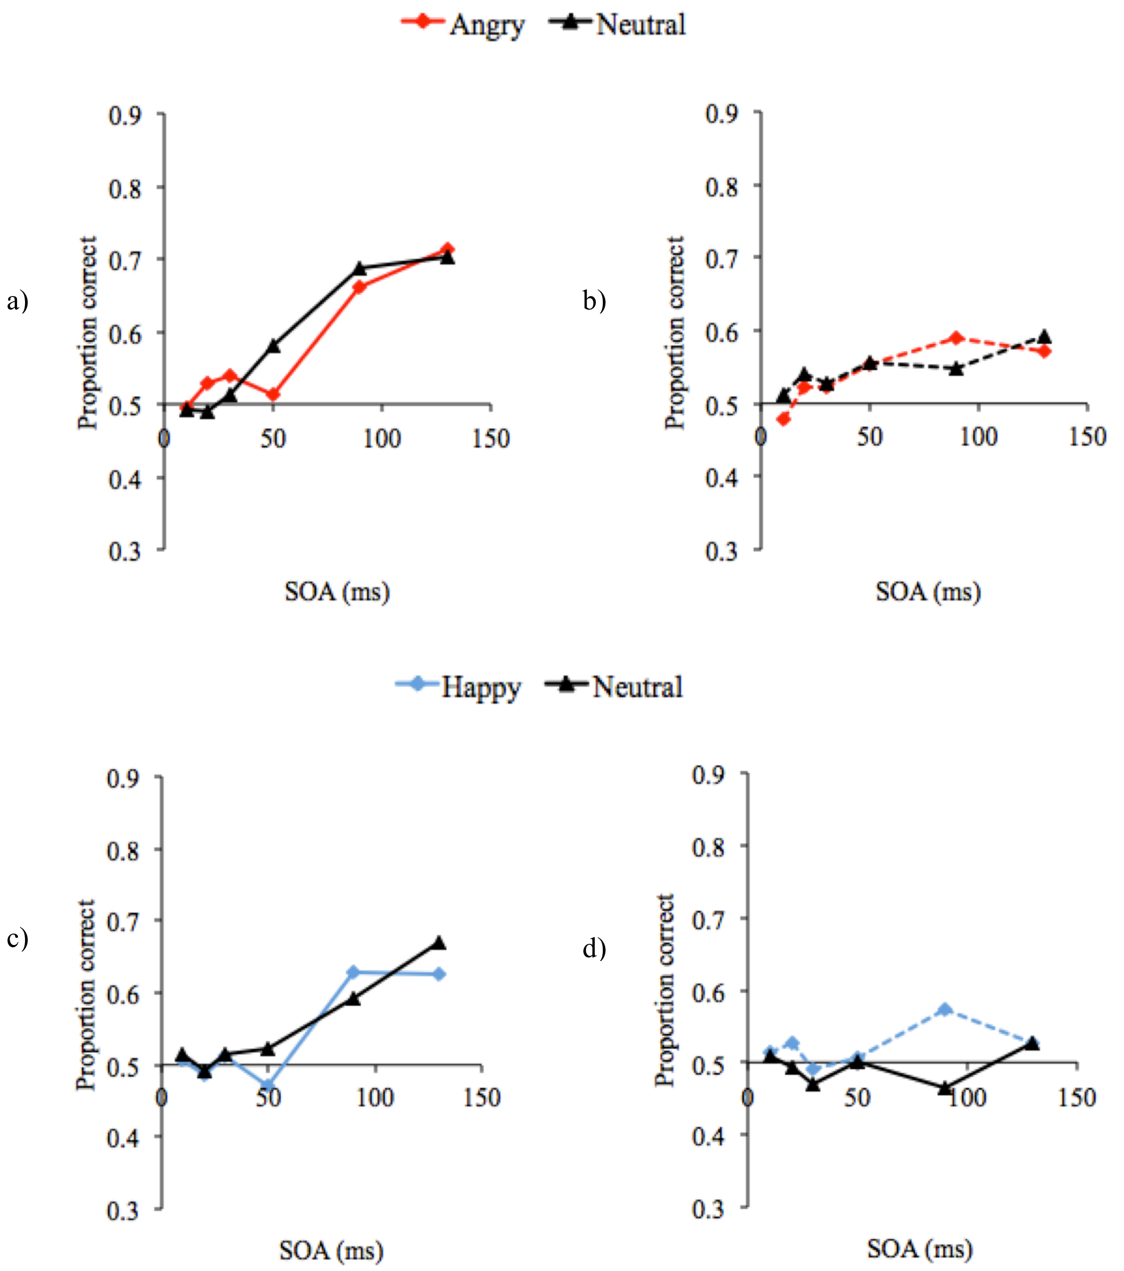

Figure 5.6: Group mean categorisation performance (proportion correct) for Experiment 4, for the upright (a & c) and inverted (b & d) faces for the angry and equivalent neutral faces (a & b) and for the happy and equivalent neutral faces (b & d).

The data were broken down into specific emotions (i.e. angry and happy) and the equivalent neutral stimuli, i.e. for participants shown angry male faces and female happy faces, the equivalent neutral for the angry trials would be male, and the equivalent for the happy trials would be female. Performance for each emotion and neutral equivalent is shown in Figure 5.6. Of particular interest here is whether there is evidence at the shorter SOAs for higher performance at either emotion compared with their equivalent neutral performance, and whether this interacts with inversion –

i.e. an interaction between orientation, stimulus type, emotion, and SOA. This had to be calculated as proportion correct as it involved analysing individual responses and hence d' cannot be computed. There was no significant interaction between orientation, stimulus type, emotion and SOA, $F < 1$, or significant interaction with version of the experiment, $F (5, 70) = 1.17$.

The interaction between stimulus type, emotion and SOA was also not significant, $F < 1$, but did interact reliably with version of the experiment, $F (5, 70) = 2.98$, $p < 0.05$. There is a small suggestion here that the separate emotions may interact with the results, although this interacts with the version of the experiment and is not reliably different between orientations. The interaction with version of the experiment hints that the effect varies for the male and female faces (as for version 1 participants the angry faces would be male, and for version 2 the angry faces female). Overall, as shown in Figure 5.6, there are no convincing differences between the emotional and neutral faces for either the angry or happy expressions and equivalent neutral faces. There is a slight emotion advantage for the upright faces at 20 and 30 ms for the angry compared to the equivalent neutral faces but this is very small (mean difference = 0.03, SEM = 0.03) and not at the same SOA as the effects found in Experiment 3. The persistent emotional advantage at 90 ms SOA for the inverted faces appears to be present for both happy and angry faces.

As it is likely that the association between emotion and response was learned (if at all) only gradually, it makes sense to analyse the data as a function of practice. Hence the data from the two sessions were analysed separately. The results from each session can be seen in Figure 5.7. For the upright faces, there does appear to be to be a small advantage at 30 ms for session 1, and at 20 ms for session 2, but these are small, not at the predicted time from Experiment 3 (10 ms) and in session 2 there is a larger effect in the opposite direction (50 ms SOA). For the inverted faces, the previously mentioned difference at 90 ms was present for both sessions. There was no substantial evidence, then, for any significant effect of session on the outcome.

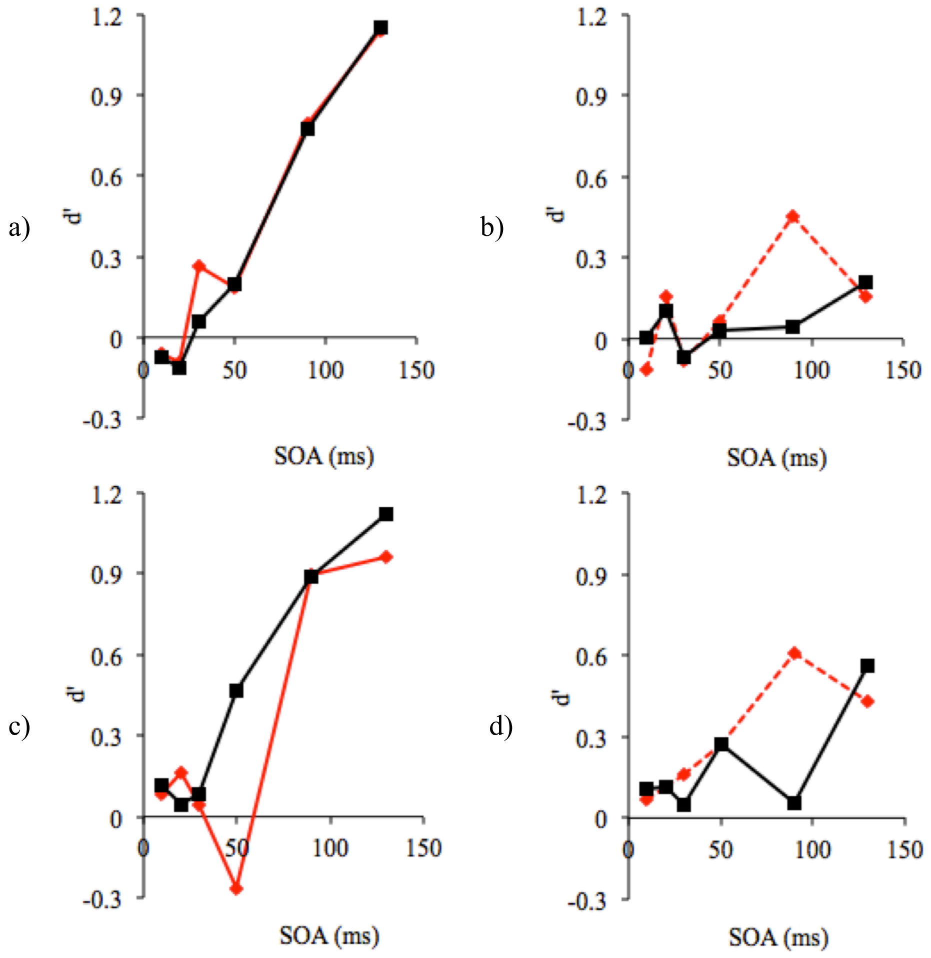

Figure 5.7: Group mean categorisation performance (proportion correct) for Experiment 4, for the emotional and neutral faces, for upright (a & c) and inverted (b & d) blocks, in session one (a & b) and session two (c & d).

5.3.3. Discussion

Overall there appears to be an appropriate masking function with performance improving from chance level with increasing SOA. This is true for both upright and inverted faces, although inverted faces are generally responded to less accurately. For the emotion faces, the emotional valence provided an addition cue to the discrimination. But there was little evidence of participants using this cue even at the longer SOAs, and hence discriminating the gender better, even when d' for the upright faces exceeded one. Nor was there consistent evidence for the gender of the emotional faces being discriminated worse than for the neutral faces (as might happen if the emotion grabbed attention). Hence the lack of any difference between neutral and emotional faces at short SOAs tells us essentially nothing about the privileged processing of emotion in subliminal stimuli.

The differences between the upright and inverted faces both as a main effect of orientation, and interacting with SOA, suggest that gender may well have been processed configurally. When the faces were inverted, performance was lower than when upright (when discrimination was notably above chance), suggesting when forced to rely upon featural information performance suffers compared to configural information.

The only point in the data where an emotional advantage was present, and considerable in size, was for the inverted faces at 90 ms. This effect was not predicted, but was significant with Bonferroni correction. It is not clear why the effect was found at this point in time. Moreover, the generality of the effect across different emotions (it was present for the happy and angry faces), and the size of the effect, makes it unlikely it is driven by a specific stimulus artefact. The source of the emotional advantage for inverted faces at 90 ms SOA is unclear, and so given the fact it was not predicted, and there is no evidence for a similar effect in the upright faces, I discount it as probably spurious.

Why did Experiment 4 yield no convincing evidence of an emotion advantage? One possible interpretation is that emotional stimuli do not carry any processing advantage and the effect of interest does not exist. This seems unlikely given the results of Experiments 1 and 3, and the previous literature on low visibility emotion processing. The more likely explanation is that participants did not learn to use the emotional cue (whether consciously or not). Gender discriminations are a common, socially relevant discrimination to be made and it is quite possible that all else was essentially ignored whilst the faces were processed (i.e. any emotional advantage was simply not used as relevant for the decision). In the associative learning literature, there is a phenomenon called ‘blocking,’ which is where a new, additional, discriminating cue is not learnt when predictive cues are already present (first noted by Kamin, 1969). To apply this to the design of Experiment 4, if there is information that is 100% predictive for the discrimination (gender), this may actually ‘block’ the learning of the other cue (emotion). It may also be the case that to be learned as a redundant cue, the emotion must be registered consciously. Given the generally low visibility of the stimuli and the fact that many participants did not even notice the emotion in the faces, these would not have been optimal conditions for such learning. (Note that in Experiment 1, where emotion was effective as a redundant cue, the TMS trials were preceded by a learning phase in which the stimuli were clearly visible.) Finally, it may be that when a control, neutral, task involving configural information is used, no advantage for emotions can be found.

5.4. General Discussion

In summary, only Experiment 3 provided some evidence that emotional stimuli carry some processing advantage when compared to neutral discriminations. This evidence was generally marginal in reliability, and was not supported with any evidence of a processing bias for emotions in Experiment 4.

There are alternative explanations to the emotion effect shown in Experiment 3. The emotion advantage may not have been specific to emotions, but a general meaningfulness effect. I.e., it was not that the discriminations were emotional, but that they were based on something (socially) meaningful and relevant. This could be looked at from two different angles. Firstly, it could be that the emotions were not used, and there was some alternative cue of meaning. As the faces only varied in the emotional expression shown, and there is no other ‘meaningful’ dimension apparent, this seems to be very unlikely. Second, it could be that the emotions did lead to the effect, but a similar advantage could be found with other ‘meaningful’ discriminations – i.e. the effect is not specific to emotions. This is a difficult point to address, as emotions are themselves one of the most meaningful cues possible in social stimuli. Another consideration of note is that there is evidence in the literature for emotions’ having some processing privilege, but the same is not generally true for meaningfulness in general. There are, then, ready neurological interpretations possible for an emotional advantage but not for a general meaningfulness effect.

Finally, it could be that it was again not the emotional content per se, but a bias for configural processing that led to a low visibility advantage. For Experiment 3, one could argue that the emotion discrimination relied on configural information, but the orientation task was done featurally. This is not clear, as it is certainly possible that the orientation task could be done configurally. Also, in Experiment 4, the only point at which there is an emotion advantage is for the inverted faces, which were presumably processed featurally. The results from the experiments in this chapter are not definitive on this point, as the emotional effects for the inverted faces are largely different to those for the upright faces. However, an interpretation of emotions being processed featurally at low levels of visibility, and so not due to configural processing, appears to be the most plausible from the findings of this chapter.

To sum up, it appears that there are differences between the processing of neutral and emotional stimuli when visibility is very low: emotional properties of faces are relatively privileged compared to neutral properties of faces. This effect was mostly present in Experiment 3, and was highly specific – there was no evidence for a general emotional advantage. The advantage gained in emotional stimuli appeared to be very short lived and to occur at precise periods of time. Whether from striate/extra-striate cortex or subcortical structures, there was some evidence that emotions were processed preferentially at very low levels of visibility – at least under some (very specific) conditions.

CHAPTER 6

Influencing behaviour with subliminal emotional faces: four priming experiments

6.1. Introduction

The experiments reported in Chapter 5 investigated processing of briefly presented faces, awareness of which was manipulated through varying the delay in the onset of the backward mask in a sandwich-masking paradigm. The experiments compared the effect of varying the backwards mask onset on discrimination of emotional expression to the effect on discrimination of orientation, and gender. Masking affected all tasks as expected, with performance increasing from chance or near-chance levels as the backward mask onset was delayed. In Experiment 3 there was evidence of better performance on the emotion discrimination task than on the orientation discrimination task at the lowest level of visibility, even though at higher levels of visibility the orientation task was clearly much easier. This provided some evidence of privileged processing of emotion for low visibility stimuli. However, this effect was quite small and of only marginal reliability. The attempt in Experiment 4 to demonstrate privileged emotion processing by using emotion as a redundant cue to gender discrimination failed, possibly because the participants simply did not take advantage of the emotion cue.

The experiments reported in this chapter used a different approach to investigate what information the brain can extract from low visibility (masked) emotional faces. In four experiments I tested the effect of low visibility, masked, emotional expressions on the processing of an immediately following visible target. This type of design – known as priming – frequently leads to congruence effects (e.g. faster responses when the prime and target are of the same valence than of a different valence). For emotional stimuli, it can also lead to arousal effects irrespective of the congruence of prime and target (e.g. negative prime stimuli may capture attention more than positive prime stimuli). These effects, and the literature on priming, were reviewed in Chapter 1. The main points from the literature that are relevant to this chapter are briefly recapped below.

Early work on priming found that a subliminal prime could influence how quickly people responded to subsequent stimuli (Marcel, 1983), but this work was later criticised on methodological grounds (Holender, 1986; Kouider & Dehaene, 2007). Criticisms of the Marcel (1983) study included: differences in the conditions under which visibility and priming were measured; taking the visibility measure before the priming trials and hence possibly confounding visibility with practice, and the small amounts of data used to estimate visibility. More significant evidence that subliminal stimuli can be processed meaningfully has since emerged (see Kouider & Dehaene, 2007 for a review). For example, responding faster with a left response to an arrow prime following a prime of an arrow pointing to the left (Schlaghecken & Eimer, 2004). There is also growing evidence that prime images can influence ‘free choice’ on trials where the target has no clear response associated with it, resulting in participants being more likely to make a choice congruent with the prime e.g. in a task choosing between left and right arrows, more likely to choose ‘left’ to an arrow pointing upwards if primed with a left pointing arrow (Schlaghecken & Eimer, 2004). Since the presence of ‘subliminal’ priming effects has been shown many times to date, the shift in focus of research in recent years has been to try to understand the conditions and limitations of subliminal processing for various types of images, the precise effect of subliminal processing on behaviour, and the neurological basis of subliminal processing (and what this can suggest for visual processing and awareness in general).

Of particular interest here is subliminal priming with emotional stimuli. In the priming literature, it has been found that an emotional prime (such as a face expressing an emotion) can affect the speed with which people respond to the valence of a subsequent emotional target, with faster responses for same valence primes and targets (e.g. Aguado, Garcia-Gutierrez, Castaneda & Saugar, 2007; Houwer, Hermans, Rothermund & Wentura, 2002; Tamietto & de Gelder, 2007). Emotional primes can also influence how attractive a neutral stimulus is rated with more attractive ratings following positive primes (Murphy & Zajonc, 1993), and how highly a beverage is rated with higher ratings following positive primes (Winkielman, Berridge & Wilbarger, 2005). Emotional primes have also been found to attract attention when participants complete a non-emotional task (Mogg, Bradley, de Bono & Painter, 1997). Emotional primes, it seems, can influence people’s RTs, choices, and even preferences - at least in a similar manner to non-emotive primes.

There are still many points of debate within the broader subliminal priming literature, as covered in Chapter 1. To summarise these, controversy remains as to whether priming effects are truly due to subliminal processing, or some form of partial awareness, and whether the priming effect results from semantic processing of the prime or is due to activation of S-R mappings formed by recent encounters with the prime stimulus. If the prime was presented consciously at some point during the experiment (typically because the prime and target stimuli sets consist of the same images), it is possible for the image(s) to become associated with a response and for subliminal presentation to activate these mappings without necessarily activating meaning (or an emotional reaction). Experiments using prime stimuli that have not been shown consciously (i.e. have been encountered in the experiment only as a subliminal prime, which I will refer to as 'target independent') have found significant priming effects (i.e. Kiesel, Kunde, Pohl & Hoffman, 2006) indicating that priming can occur when the primes do not have S-R mappings associated with them. Evidence of emotion processing of emotional subliminal stimuli is evident from experiments where the subliminal image elicits physiological responses associated with the emotion e.g. altering skin conductance responses (Stone & Valentine, 2001), or sympathetic muscle activity associated with an emotion, e.g. a smile (Dimberg, Thunberg & Elmehed, 2000). Also, the apparent ability of emotional prime stimuli to capture attention outside awareness (Mogg et al., 1997) is not susceptible to an S-R mapping account, as there is no overt response.

Further questions not yet addressed in the literature concern the nature of the processing of the prime beyond the semantic/S-R debate. For example, in Chapter 5 it was found that a masked inverted face could lead to emotional advantages even when the stimuli were low visibility. This provides evidence that the emotional stimuli were being processed (or at least, could be processed) featurally, as opposed to configurally, at low levels of visibility. An interesting question, then, would be whether priming effects can be found with inverted faces, and hence whether any effect of subliminal images is due to featural or configural processing.

The following four experiments aimed to test whether a target-independent subliminal emotional face could:

a) influence people's RTs when they must classify the emotional valence of a target face (Experiments 5 and 6) or word (Experiments 7 and 8),
b) influence a choice between emotional classifications of a neutral target (Experiments 5, 6, and 8).

The experiments also asked whether any priming effect can be found from inverted faces (Experiments 5 and 6), how priming may interact with prime visibility (Experiments 7 and 8), and finally how manipulating the location of a target word relative to a prime face may influence the priming effects found (Experiment 8). The rational for using inverted faces follows the findings from Experiments 3 and 4 (Chapter 5) that suggested any bias for processing emotions may be featurally based, and relatively unaffected by inversion. By including inverted faces, Experiments 5 and 6 tested whether the same is true for a priming paradigm. In Experiment 7, the visibility of the prime was varied to assess the ‘visibility window’ in which priming effects are found, and how visibility interacts with priming in general. Finally, the location of the target words in Experiment 8 was varied as a check to see whether the point of spatial overall between the prime and target was an important factor in the priming effects found in Experiment 7.

Masking protocols were used based on those developed in Chapter 5 as they provide some level of knowledge about the masking function, so that SOAs could be chosen that are likely to result in subliminal presentation. It also allows for a degree of comparison across the two sets of experiments. The experiments were designed primarily to look for response congruence effects, that is relative facilitation of RT for categorisation congruent with the prime's emotional valence (which does not appear to have been demonstrated for subliminal primes), or bias of free choice towards the congruent category. However, I anticipated that we might also find effects attributable to emotional arousal due to the prime, in addition to effects of congruence.

Experiment 5 used quite a large set of prime images consisting of computer-generated faces displaying happy and angry facial expressions, occurring as upright and inverted faces (as used in Experiment 3). The masking parameters were the same as in Experiment 3 for the 10 ms masking condition. In “congruence priming” blocks, the target face could be happy or angry, and RT as a function of the target’s emotional congruence with the prime’s valence was the intended measure of emotion processing of the prime (and its effect on target processing). In “forced choice” blocks, the target was neutral, and the measure of emotion processing of the prime was the effect of the prime’s valence on forced choice of the neutral target. The prime images were displayed only subliminally to the participants during these two kinds of priming blocks, and the targets were from a quite different set of faces. Following the priming phase of the experiment, a visibility test was conducted in which participants were told of the presence of the prime face and asked to discriminate its emotional expression or orientation (upright or inverted).

6.2. Experiment 5

6.2.1. Methods

6.2.1.1. Participants

Twenty-eight undergraduate students participated for course credits and up to £1 as a bonus. Fourteen were male, and fourteen female with an average age of 23 (SD = 5) years old. The experiment lasted around 50 minutes. All gave informed consent before participating.

6.2.1.2. Materials

Examples of the target and prime stimuli can be seen in Figure 6.1. Forty identities from the NimStim face set (as used in Experiments 2 and 4) were used as target images. The target images occurred with happy, angry and neutral expressions, were always upright, and were colour photos. The prime faces were generated using computer software (FaceGen). Computer generated faces were used because generating the faces gave a high level of control over the properties of the faces, and because the same images had been used in Experiment 3 giving a good idea as to the visibility of the faces when masked. The prime images consisted of happy and angry facial expressions from 20 identities. All prime images were grey scale, and each of the identities and expressions occurred as both upright and inverted faces. The mask stimuli were created by splitting a prime face into 360 squares and randomly re-ordered them (using a plug-in for Adobe PhotoShop Elements). See Figure 6.1 for examples. In total there were 8 different masks, with 4 occurring as forwards masks and 4 as backwards masks.

The faces were presented centrally on the screen, and were 200 x 200 pixels in size (2.76 x 2.76°).

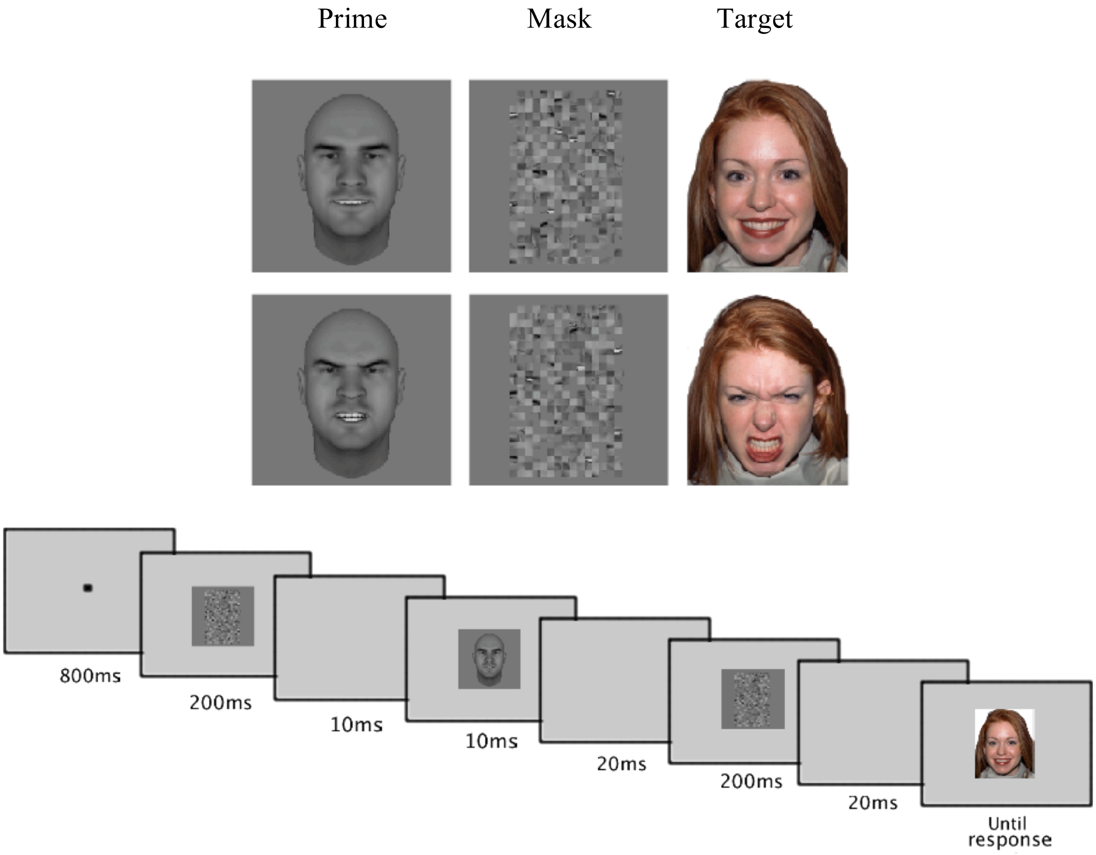

Figure 6.1: Examples of the prime, mask and target stimuli (top) and a standard trial outline for the priming phase of the experiment (bottom) in Experiment 5.

6.2.1.3. Procedure

The experiment consisted of two phases: a priming phase, and a visibility phase. The priming phase of the experiment was made up of two practice blocks and 12 experimental blocks. All blocks were 40 trials long. A standard trial is outlined in Figure 6.1. The priming phase of the experiment consisted of two types of blocks: emotional target blocks in which RTs were measured (to test for congruence priming) and neutral target blocks where un-speeded forced choice was measured (to test forced choice effects of the prime). The emotional target blocks consisted of 320 trials, 40 trials for each combination of target valence, prime congruence, and prime orientation. For all the possible trial types of the experiment, each target identity occurred once, and each prime identity occurred twice. The exact pairings of prime and target identity were randomised for each participant. The presentation of specific mask images was randomised across the experiment. The blocks where the target was neutral in valence consisted of 160 trials: 40 trials for each prime type, with each target identity occurring once per prime type, and each prime identity twice. Which block was given (emotional or neutral valence targets) alternated throughout the experiment with two blocks in a row for emotional target blocks, followed by one neutral target block.

Participants were not informed of the presence of the prime face. For the congruence priming blocks, participants were instructed to respond as quickly and accurately as they could to the emotion present in the target faces. For the forced choice blocks, participants were instructed to try and discriminate the “emotion” of the neutral faces, with no emphasis on speed, and warned that the task would be difficult. In all blocks, error feedback was given after each trial with an incorrect response in the form of a red cross presented centrally on the screen. For blocks with emotional target faces, the feedback simply reflected the correct classification of the target emotion; for blocks with neutral target faces, feedback reflected the emotion present in the prime images, essentially giving participants training for the discrimination.

Bonuses were available from each emotional target block based on a combination of speed and accuracy. On each trial, participants were awarded 0 points for an incorrect response, 1 point for a correct response after more than 500 ms, and 2 points for a correct response 500 ms or less following target onset. A 4 p bonus was given for the emotional target blocks if the score was over 49 initially, and then based on improving the score from the previous block. For the neutral target blocks, performance was not based on speed (0 points if incorrect, 1 point if correct) and a score of 25 or above gave a 4 p bonus in the first block, to set the target score for improvement over subsequent blocks. Participants were informed of their score and whether they got the bonus at the end of each block.

After the priming phase of the experiment, participants completed six blocks (40 trials each) of a visibility test. At this point they were informed about the presence of the prime face, and asked to complete one of three tasks: discriminate the emotion (happy or angry) when the face was upright, discriminate the emotion (happy or angry) when the face was inverted, or the orientation of the face (upright or inverted). The task varied between every block, with the following order run twice in succession: emotion discrimination with the faces upright, orientation discrimination, and emotion discrimination with the faces inverted. There were 80 trials for each task, with each prime identity presented twice for each emotion for the emotion discriminations (for both upright and inverted tasks), and once for the orientation blocks for each possible prime type. Bonuses were awarded based on average performance, with two 25 p bonuses available in total (one at the end of each cycle of the three tasks). A bonus was given if performance was numerically better than chance 50%. Although this resulted in some rewards when the participants were just guessing, the criterion was set this way to give participants a serious chance of wining the bonuses and to motivate them to try.

Half of the participants responded with their left index finger to angry faces, and right index finger to happy faces. The other half of participants had the reversed button mappings.

6.2.1.4. Data analysis

For the priming phase of the experiment, the mean error rate (% incorrect) and mean correct RT was computed for each condition. For the RT analysis, only responses on correct trials that were slower than 200ms but faster than 2000ms were included. The total number of RT trials that were faster than 200 ms or slower than 2000 ms, and so removed from the analysis, was < 1 %. The mean RTs and accuracy rates per condition were then subtracted to give congruency effects – calculated by subtracting the congruent trials for each target emotion from the incongruent trials for each target emotion. As a result, a positive congruency effect, as reported in this thesis, reflected slower responses on incongruent than congruent trials. ANOVAs were run separately for the accuracy and RT data with the factors of target emotion, prime congruency, and prime orientation.

Appendix 3 reports the effect of repeating the analysis with median correct RTs. There are many strategies for dealing with RT outliers, and no agreement on the optimum one. But removing as outliers RTs in each cell that are some multiple of the standard deviation of the RTs from the mean does not work well, as the standard deviation is even less robust to outliers than the mean, and the RT distribution is also asymmetric. Outlier-exclusion strategies have been proposed using more robust and asymmetrical non-parametric measures of spread such as multiples of the upper and lower hinges (e.g. Ratcliff, 1993). The approach taken here is to ask whether the outcome is markedly changed if cell medians are used as the measure of central tendency. Although medians are less arithmetically tractable a measure of central tendency than the mean, and can be biased by sample size, they are robust to outliers. If the pattern of results is much the same with means and with medians, we can be confident that no strategy for removing outliers is likely to make much difference, and removes the need to justify that strategy. Of course if the pattern of means and medians does differ, that is a signal that the distributions need to be examined to see whether the difference is due to outliers or, instead, to the variable changing the shape of the distribution (see, for example, Wheeldon and Monsell, 1994).

For the visibility data, d' values were calculated for each visibility task using the formula: $d' = z(H) - z(FA)$. For the inversion task, H is p(correct) for upright faces, and FA is p(error) for inverted faces; for the emotion tasks, H is p(correct) for angry faces, FA is p(error) for happy faces.

For the forced choice neutral blocks, d' values were calculated using the same formula and format as the visibility data.

6.2.2. Results

6.2.2.1. Visibility

Performance on the three tasks given in the visibility test were computed as d' (shown in Figure 6.2) and analysed with one-way t-tests to check for departure from chance. Performance at each of the three visibility tasks was very low, and non were reliably different from chance; emotion for upright faces t (28) = -1.82, emotion for inverted faces t (28) = -0.96, orientation task t (28) = 0.49. This indicated that the visibility of the prime faces was very low, and arguably subliminal. None of the participants reported noticing the presence of the prime face during the priming phase of the experiment; indeed, when informed of the prime's

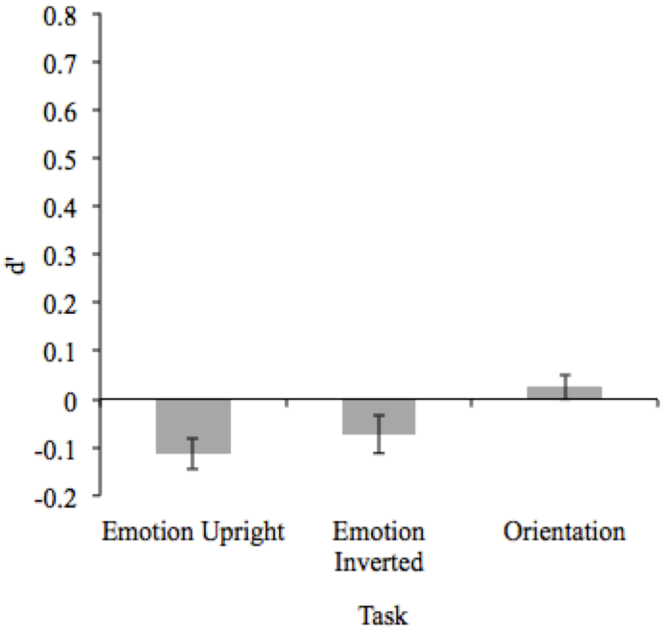

Figure 6.2: Mean discrimination performance (d') for the three visibility tasks in Experiment 5. Error bars represent SEM.

presence, all participants showed surprise.

Prime orientation	Prime emotion	Angry Target		Happy Target	
		RT	Error	RT	Error
Upright	Angry	569	3.67	558	3.29
	Happy	579	4.21	555	4.14
Inverted	Angry	568	3.84	566	2.99
	Happy	585	3.89	556	4.09

Table 6.1: Mean RT (ms) and error rate (%) for each condition in Experiment 5.

6.2.2.2. Congruence priming

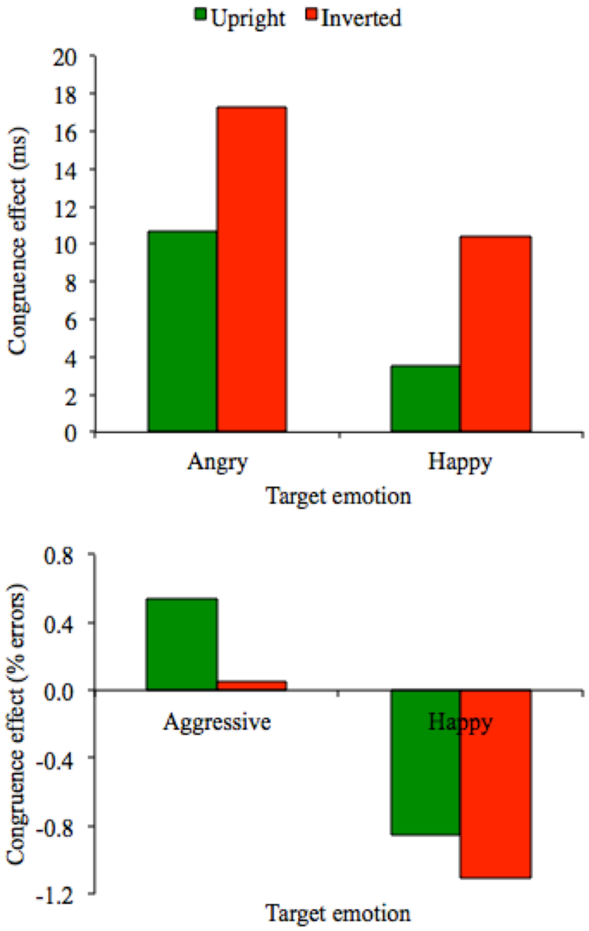

Figure 6.3: Congruency effects in RTs (top) and error rates (bottom) for the upright and inverted prime conditions in Experiment 5.

The mean RTs and error rates for individual trial types are shown in Table 6.1, and the congruence effects can be seen in Figure 6.3.

There was a significant effect of target valence, F (1, 26) = 7.18, $p < 0.05$, with the angry faces being responded to more slowly than the happy faces (mean difference = 17 ms, SEM = 6). This was not matched with a significant difference between target emotions in the error rates, F < 1, although there were slightly more errors to angry than to happy face targets (mean difference = 0.3, SEM = 0.6). There was also a main effect of congruence of the prime for the RTs, F (1, 26) = 7.52, $p < 0.05$, indicating that there was a small but significant facilitation of RT when the prime and target had the same valence. For the error data, there was no significant main effect of prime congruence, F < 1.

Figure 6.3 shows a trend for larger congruence effects for inverted faces than upright faces (mean difference = 7 ms, SEM = 7 ms), however congruence did not reliably interact with inversion in the ANOVA for the RTs, F < 1, or for the error rates, F < 1. It also appears that there are larger congruence effects for the angry targets than the happy targets in the RTs (mean difference = 7, SEM = 6), and a trend for a positive congruence effect in the error data for angry target trials, and a negative congruence effect for happy target trials (mean difference = 1, SEM = 1). There was a general trend for an effect of prime emotion on the error rates, where there are fewer errors following angry than happy primes (mean difference = 0.6, SEM = 0.5), and a very small difference in RTs with angry primes leading to slightly faster responses than happy primes (mean difference = 4 ms, SEM = 3). The interaction between target emotion and prime congruence was not significant for the RT data, F (1, 26) = 1.52, or for the error data, F (1, 26) = 1.56, indicating the differences between congruence and RTs for the different prime and target emotions were not significant.

6.2.2.3. *Forced choice neutral trials*

For the blocks where the target faces were neutral, performance was measured with d' for accuracy. Performance is shown in Figure 6.4. Performance for upright (mean d' = 0.17, SEM = 0.08) and inverted (mean d' = 0.05, SEM = 0.09) faces was tested for departure from chance with a two-tailed t-test. For the upright faces, performance was reliably different from chance, t (28) = 2.15, p < 0.05. For the inverted faces, performance was not reliably different from chance, t (28) = 0.59. But the difference between performance on the two orientations was not significant, t (28) = 1.54. The emotion of the prime, then, did

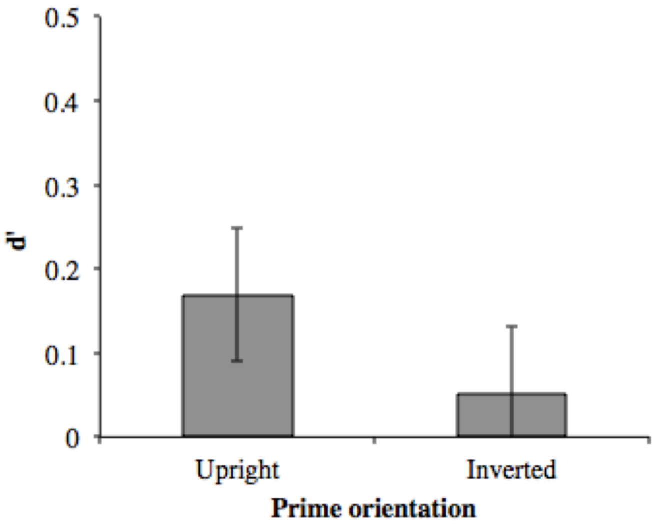

Figure 6.4: Group mean performance at the forced choice neutral trials for Experiment 5. Error bars represent SEM for performance at each prime orientation.

appear to have a detectable but small effect on forced choices to neutral targets, but it was significant only for upright face primes.

6.2.3. Discussion

Experiment 5 found that the prime faces did indeed influence subsequent responses for both the RT task and the forced choice task. This was true even though the prime stimuli were from a different set of images and were presented only subliminally, so it cannot be due to response learning. However, the effects were small, and not consistent across target types and measures. Classification RTs were faster when the prime had the same emotional valence as the target. This congruence effect was in the region of 12 - 14 ms for responses to the angry target faces. The effects for the happy target face responses were considerably smaller (2 - 6 ms), indicating the priming effect was predominantly found for responses to negative target faces although the difference in congruence for positive and negative target faces was not significant. Moreover, for error rates there was a negative congruence effect for positive target faces.

Somewhat surprisingly, the congruence effect was at least as strong when the prime face was inverted. The typical face inversion effect (reduced performance for inverted, compared with upright, faces) is thought to reflect the disruption of configural information (Yin, 1969). Equivalent congruence priming for upright and inverted faces suggests that the priming is due to a match in emotion-conveying features between prime and probe face, rather than configural processing of the prime's facial emotion. For example, the mouth shape is quite salient in the computer-generated images used (see Figure 6.1).

An alternative account of the priming effects found in this experiment is the effects could arise from structural similarity between the targets and the primes and are nothing specifically to do with emotion. In other words, the priming could be a form of repetition priming. Against this point is the finding of priming effects with inverted faces – the targets were always upright so features and basic properties did not overlap. Also, the use of different face stimuli sets for the primes and targets should help to minimise this possibility.

Another point to consider here is the role of repetition blindness. This is a phenomenon whereby the repetition of a stimulus reduces participants’ ability to report the repetition. This is not an issue in the visibility test participants completed as only the prime face was presented on those trials. For the priming trials, however, both prime and a target face were shown in rapid succession. This could have actually reduced the visibility of the prime face relative to the visibility test, and could in part account for the small priming effects found. The visibility was very low as measured in the formal visibility test; if another factor reduced this further during the priming it could have disrupted processing to a too greater extent to find substantial priming. The inversion of the prime face on some trials would have minimised this, and could in fact explain in part why priming effects were if anything larger on the inverted prime trials: visibility may have been higher due to no repetition blindness. This point is difficult to avoid with the design used in this experiment, but will be addressed in Experiments 7 and 8.

The effect on the forced choice of emotion for a neutral face was also very small, but reliably above chance. As the participants received feedback in the response bias blocks, the effect was likely to be due to some degree of learning to use the products of subliminal processing of the image. Only the upright faces produced an effect on discrimination reliably above chance, although the difference between upright and inverted prime faces was not itself significant.

Explicit discrimination of both the emotion and orientation of the prime faces was no better than chance, even though visibility was measured after a lot of practice at the forced choice task. Moreover, no participant noticed the prime face during the priming phase of the experiment, and all expressed surprise when they were told. Hence the face primes in this experiment were subliminal by the toughest criteria. Thus the conclusion so far is that emotional attributes in at least angry faces can be extracted when the faces are presented subliminally, as indicated by their impact on RT for classifying facial emotion in a following target, or forced choices about the emotion in an actually neutral target. However, it may be that this reflects detection of emotion-conveying features rather than configural processing of the subliminal face's emotion.

6.3. Experiment 6

Although the results of Experiment 5 were promising, the conclusions that could be drawn from Experiment 5 were limited. The effects of the prime were small and a clear effect was limited to angry target faces. The computer-generated face primes are by definition artificial and neither they nor their emotional expression may have been sufficiently visible to give substantial priming effects; they may also have been discriminable largely at a featural level. Hence in Experiment 6 I used a different set of prime faces, Ekman's classic set of photos of actors displaying exaggerated expressions, to try and increase the emotional effects, on RTs, error rates, and decision bias. Scrambling Ekman faces in the same way as Experiment 5 generated the mask stimuli. In addition, an easier task was given as a visibility test (presence or absence) in case the orientation task was not a sufficient test of visibility. Presence/absence is the easiest discrimination possible and hence provides an even stricter measure of visibility.

6.3.1. Methods

6.3.1.1. Participants

Twenty-four undergraduate students participated for course credits and up to £1 as a bonus. Twelve were male, and twelve female with an average age of 20 years (SD = 4 years). The experiment lasted for around 50 minutes. All participants gave informed consent before participating.

6.3.1.2. Materials

The target faces were the same as Experiment 5. For the primes, faces were taken from the Ekman set. The prime images consisted of happy and angry facial expressions from all 10 identities available in the Ekman set, and all were used at both upright and inverted orientations (giving a total of 40 prime images). The mask stimuli were created by splitting a prime face into 864 squares and randomly re-ordered them (using a plug-in for Adobe PhotoShop Elements). There were eight mask images in total, four used as forwards masks and four as backwards masks. Faces were presented centrally on the screen, and were 120 x 180 pixels in size (1.66º x 2.49º).

6.3.1.3. Design and Procedure

The design was essentially that of Experiment 5, as was the timing of events of a trial. For the priming phase of the experiment, the feedback and bonus system were the same as Experiment 5. For the emotional target block types, there were 320 trials with 40 trials per target and prime emotion and orientation combination. Each prime identity occurred four times for each type of trial, and the target faces once for each identity, per trial type. The specific pairing of identities was randomised for each participant. For the neutral target blocks, there were 160 trials, with 40 trials per prime type, again consisting of one presentation of each target identity, and four presentations of each prime identity.

For the visibility phase, there were three tasks: emotion judgement with upright faces, emotion judgement with inverted faces, and present/absent judgement. Participants cycled through the tasks (emotion upright, present/absent, emotion inverted) twice. For the present/absent block, half the time a face was presented (at any of the possible orientations and emotions), the other half nothing was presented. There were 240 trials for this phase of the experiment, 80 per task. For the emotion tasks (both upright and inverted) each prime identity occurred four times per condition. For the present/absent task, each prime identity occurred once in each emotion and each orientation.

6.3.1.3. Data analysis

The data were analysed in the same manner as Experiment 5. The total number of RT trials that were faster than 200 ms or slower than 2000 ms was < 1 %..

As in Experiment 5, a median analysis of the RTs in included in Appendix 3. An items analysis was also run and is shown in Appendix 2.

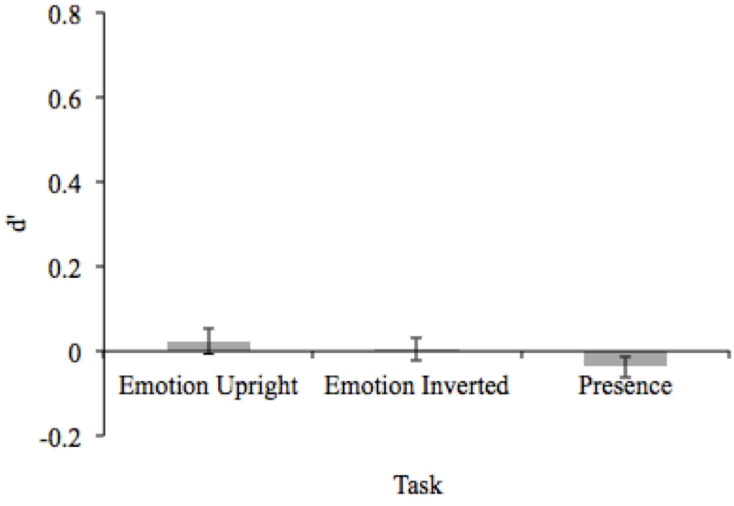

Figure 6.5: Mean group discrimination performance (d') for the three visibility tasks in Experiment 6.

6.3.2. Results

6.3.2.1. Visibility

Performance on the three tasks given in the visibility phase of the experiment were computed as d' (see Figure 6.5). A one-way t-test was run on each task to test for departure from chance performance. None of the three tasks, emotion

Prime orientation	Prime emotion	Angry Target		Happy Target	
		RT	Error	RT	Error
Upright	Angry	535	5.14	539	6.53
	Happy	539	4.7	538	5.48
Inverted	Angry	533	4.59	545	4.84
	Happy	539	4.19	538	5.73

Table 6.2: Mean RT (ms) and error rates (%) for each trial type for Experiment 6.

upright faces, emotion inverted faces, presence of face, were performed reliably above chance with the respective t-values of: t (23) = 0.39, t (23) = 0.08, t (23) = -0.73. None of the participants reported noticing the presence of the prime face, and all expressed surprise when told of the primes' existence.

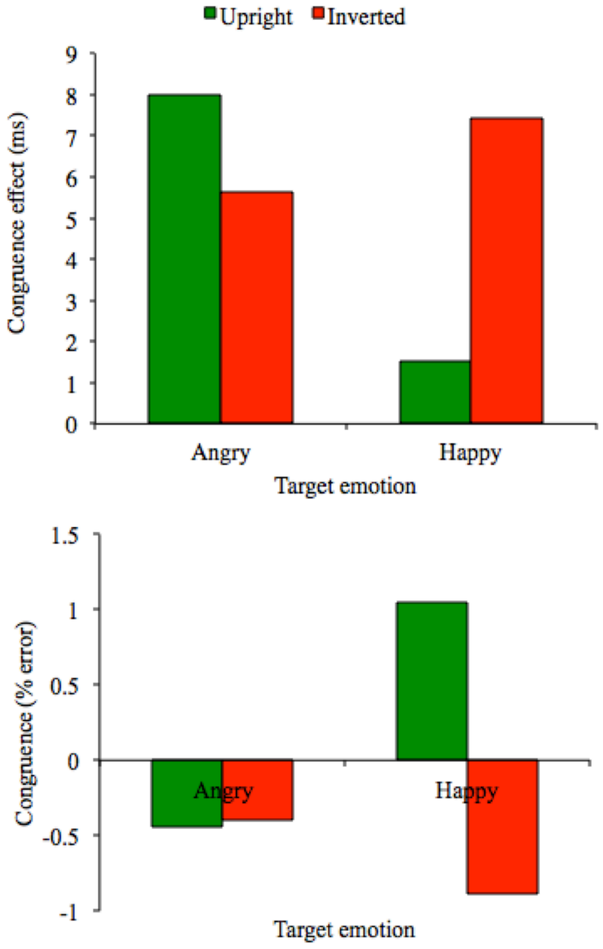

Figure 6.6: Mean congruence effects in RTs (top) and error rates (bottom) for each prime-target stimulus type in Experiment 6.

6.3.2.2. Congruence priming

The mean RTs and error rates per condition are displayed in Table 6.2, and the congruence effects are shown in Figure 6.6.

ANOVAs were run with the within-subjects factors of target valence, prime congruence and inversion and the between subjects factor of response assignment on the mean RTs and mean error rates. The within-subjects factors of target emotion, prime congruence and prime orientation, and the between-subjects factor of response assignment were included. There were no significant main effects or interactions in the RT or error data, F <= 1 for all.

6.3.2.3. Forced choice neutral trials

For the blocks with neutral target faces, there was no significant effect of either upright, t (23) = -0.09 (M = -0.01), or inverted, t (23) = -0.7 (M = 0.05), prime faces. There is no evidence that the prime reliably influenced participant's choice when the target faces were neutral.

6.3.3. Discussion

Sadly, Experiment 6 provided no evidence of significant effects of congruence between the prime and target emotion, either for classification RT or for forced choice about a neutral target. There was a hint of faster responses following an angry prime compared with happy, but this was not significant, and the happy responses were very slightly more accurate overall.

Three possible reasons why priming effects in Experiments 5 and 6 have been very small, and congruence unclear or entirely absent, are that the prime stimuli were ineffective, visibility was too low for any meaningful processing, or the prime faces were not sufficiently emotive/life like to lead to substantial emotional effects. Experiment 5 used computerised faces as primes, and Experiment 6 used a relatively small set of real life faces that showed exaggerated expressions and are, arguably, not very realistic. It is possible the emotions in these faces were not sufficiently powerful or salient to lead to congruence effects in the data (or only led to very weak ones). Alternatively, if priming in Experiment 5 was based on facial features, the exaggerated expressions in the prime faces from Experiment 6 may have been too dissimilar to the target features to induce priming effects. If so, this would argue against the priming effects reflecting semantic processing of expression, but hint to more basic perceptual and featural processing effects.

To summarise the findings so far it appeared from Experiment 5, using computer-generated face primes that subliminal primes never seen as targets could affect how quickly participants classified target images, and influenced free choice about the emotion in a neutral face, although the effect was small, significant only for angry primes, and may largely have reflected featural priming. I was, however, unable to replicate this effect using Ekman faces as primes in Experiment 6.

6.4. Experiment 7

In the two experiments so far, I picked a single SOA on the masking function (from Experiment 3) where discrimination was very poor and used it to ensure that presentation was subliminal. In one sense, this was successful: visibility tests in both experiments showed that participants could not explicitly discriminate the emotion in the prime face, or even its orientation (Experiment 5) or whether there was a face prime (Experiment 6). But perhaps the conditions I adopted to meet this demanding objective criterion for non-awareness — zero d' for discrimination of any property of the prime — were so stringent that they abolished all (Experiment 6) or almost all (Experiment 5) processing of the prime. Even if there can be subliminal processing of emotion, to demonstrate this in the masking paradigm requires locating the right point in what may be a narrow window between adequately preventing awareness and preventing any processing at all. In the remaining experiments, therefore, I varied the masking parameters of the prime to investigate the link between visibility and priming effects, and to attempt to ensure that priming effects were not small or absent merely because visibility of the stimuli was just too low. A similar principle was used by Schwarz & Mecklinger in 1995. They manipulated visibility of flankers in a flanker task and generated a flanker compatibility effect that increased with visibility, but were able to find a mask SOA at which flanker discrimination was still at chance whilst there was nevertheless a significant flanker compatibility effect

Another issue raised by Experiment 5 was whether any priming effect obtained could be due merely to priming of the detection of specific features, such as a smiley mouth, rather than to processing the prime's emotional valence per se. To avoid this possibility, the target stimuli were changed to emotional words, so that the target stimuli had no overlap in perceptual form with the prime stimuli, only a shared or conflicting positive or negative emotional valence. This carries an advantage of avoid possible repetition blindness effects as discussed in relation to Experiment 5. The disadvantage with this change is that the overlap of emotions between the prime and target stimuli is less precise, being at a generic level of valence rather than specific overlap in emotion (e.g. anger). It is not clear whether it is possible to find measurable priming effects when the prime and target stimuli are different types of images (e.g. words and faces). When prime stimuli are visible it is possible to find measurable priming effects across modalities (Spruyt, Hermans, De Houwer & Eelen, 2002), although with subliminal stimuli the findings are not so clear with both failures (Andrews, Lipp, Mallen & Konig, 2011) and successes (Hermans, Spruyt, De Houwer & Eelen, 2003) evident. I also used a large number (96) of different emotional word targets, to minimise the possibility that participants would learn S-R mappings to the target stimuli.

A different face set was used as prime stimuli (the set used as targets in Experiment 4), in case the small/non-existing priming effects found so far were due to unrealistic faces, or small stimulus sets. Finally, prime faces with neutral expressions were included to provide a neutral 'baseline' measure of performance with no emotional prime.

6.4.1. Methods

6.4.1.1. Participants

Twenty-four participants took phase in exchange for payment (up to £6 depending on performance). Six were male and eighteen were female, with an average age of 22 (SD = 4) years. All gave informed consent before participating.

6.4.1.2. Materials

The experiment used faces taken from the Nim-Stim database. In total, sixteen male and sixteen female identities were used, converted to grey-scale and cropped to remove some of the hair/outer image. All identities occurred with a happy, angry and neutral expression giving 96 images in total. These were then saved at two orientations, upright and inverted, for the visibility phase of the experiment. (Inverted faces were not used in the priming phase of the experiment). The mask stimuli were created by splitting a prime face into 864 squares and randomly re-ordered them (using a plug-in for Adobe PhotoShop Elements). There were eight mask images in total, four occurred as forwards masks and four as backwards masks. Examples of the stimuli can be seen in Figure 6.7 (a and b).

The target stimuli were positive and negative words taken from the ANEW database (Bradley & Lang, 1999). The ANEW database was chosen as it gives a comprehensive list of words rated for valence. In total 96 emotional words were used – 48 positive and 48 negative. The negative and positive words were matched for word

c)

Negative	Positive	Neutral
Abuse	Loyal	Grove
Starving	Adorable	Tracking
Misery	Trophy	Corner
Thief	Peace	Front
Rejected	Paradise	Chairman

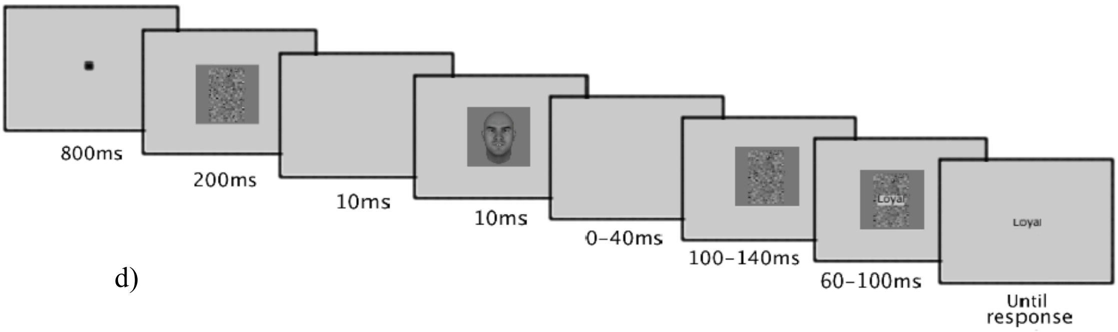

Figure 6.7: Examples of the stimuli used for the prime faces (a), showing neutral (left), happy (middle) and angry (right) expressions, examples of the mask stimuli (b), examples of the target word stimuli (c), and trial outline (d) in Experiment 7.

length and frequency. An additional set of 48 neutral words was generated using the English Lexicon Project (Balota, Yap, Cortese et al., 2007) to display as (irrelevant) targets in the visibility phase to make the visibility test as similar to the priming phase of the experiment as possible, but without distracting emotional content. The neutral words were matched individually to each of the emotional words for word length and frequency (HAL). Examples of the positive, negative and neutral words can be seen in Figure 6.7 (c). For the experiment, all words were written in upper case.

The faces and masks were presented centrally on the screen, and were 120 x 168 pixels in size (1.66 x 2.23°). The words were presented centrally in font size 18 (the letters were around 0.2° tall).

6.4.1.3. Procedure

On each trial of the experiment the prime face was presented between a sandwich mask (with a forwards mask before the face and backwards mask after the face) and then a word was presented. The word was embedded in the backwards mask for a variable interval depending on the mask SOA. That is, a constant interval separated the prime offset and the target onset, and the mask had a constant duration of 200 ms. All that varied was the interval between the prime onset and the mask onset (See Figure 6.7, d). The experiment consisted of two halves: a priming phase, and a visibility phase. The order in which the two were completed was controlled across participants. Each phase lasted for 480 trials split into 15 blocks.

For the priming phase, there were three possible prime types (happy, angry and neutral), two target word types (positive and negative), and five different mask SOAs (10, 20, 30, 40 and 50 ms). There were 16 trials per trial type, 480 trials in total. The prime identities were split into two groups, and one group was used for each trial type. For example, at SOA of 10, 30 and 50 ms, one group would be used for incongruent trials, the second group for congruent (and vice versa for 20 and 40 ms SOA). The specific pairings of SOA and identities for congruent and incongruent conditions was controlled across participants. I.e. half the participants would have one set of identities for congruent trials at 10, 30 and 50 ms and another set of identities for the incongruent trials at the same SOAs, the other half of the participants would have this pairing the opposite way round. Each of the 48 positive and 48 negative words occurred once per SOA, and for each participant they were allocated randomly to prime types.

During the priming phase of the experiment, participants were told to ignore the faces presented at the start of each trial and focus on the word. They responded based on the valence of the word – positive or negative, and feedback was given (a red cross) if an error was made. Participants were encouraged to respond as quickly as they could, and received bonuses based on a combination of their speed and accuracy. The prime faces were always upright, and occurred with happy, angry and neutral expressions. Up to 15 bonus payments of 3p could be earned, based on a scoring system the same as the previous two experiments. A bonus was awarded for a score above 34 on each block, and the score was given at the end of each block.

For the visibility test, participants were asked to ignore the words (which were now neutral in valence) and to decide whether the prime was upright or inverted. No trial to trial feedback was given, but bonuses could be earned based on accuracy. In addition to face orientation, participants made a second response on each trial to indicate, on a scale of 1-7, how well they saw the face. This gave a measure of the participant's subjective experience, in addition to the objective discrimination of face orientation. For each block performance at the orientation discrimination of 63% or higher (20 or more correct responses out of 32) gave a bonus of 3p, and participants were informed of their score at the end of each block.

6.4.1.4. Data analysis

Data were analysed in the same way as previously described. The total number of RT trials that were faster than 200 ms or slower than 2000 ms was < 1 %. Appendix 3 gives the median RTs, and Appendix 2 reports an items analysis.

6.4.2. Results

6.4.2.1. Visibility

The visibility data for the objective task (discriminating face orientation) were computed as d's (shown in Figure 6.8), and entered into an ANOVA with the within-subjects factor of SOA. There was a significant effect of SOA on performance, $F(1,23) = 5.88$, $p < 0.01$. Figure 6.8 shows the effect of SOA, and suggests that the significant effect of SOA reflected (largely) increasing performance with increasing SOA.

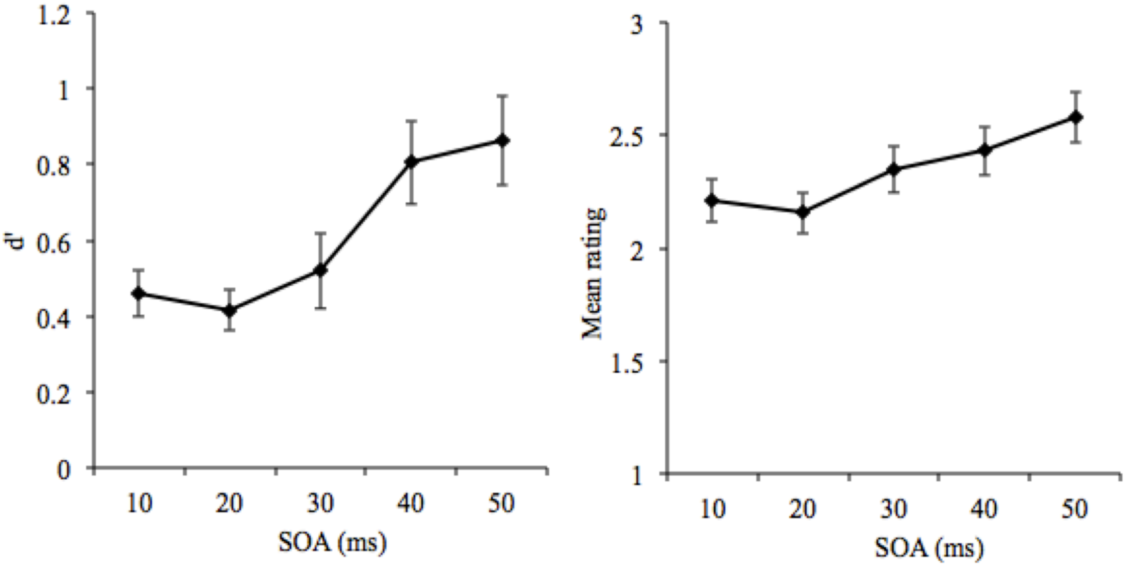

Figure 6.8: Group mean discrimination performance (d') at the visibility task (left) and the group mean subjective ratings (scale of 1-7) of visibility (right) from Experiment 7.

For the subjective measure of awareness, there was also a significant effect of SOA on the subjective rating of visibility, F (1,23) = 12.87, $p < 0.001$. Figure 6.8 shows the mean ratings given for the subjective measure of awareness and suggests that ratings generally increase with increasing SOA – as for the objective discrimination measure.

Performance for all five SOAs for the objective measure was reliably above chance: t (23) = 3.79, 3.82, 2.63, 3.69 & 3.64 respectively from 10 – 50 ms SOA ($p < 0.05$ for all). Subjective visibilities for all five SOAs were reliably above ‘0,’ t (23) = 6.47, 6.23, 6.48, 6.69, 6.86 respectively ($p < 0.05$ for all). The visibility of the stimuli, then, was above chance, although performance was certainly low, particularly at the short SOAs.

6.4.2.2. Priming effects

The mean RT and error rate for each condition are plotted in Figure 6.9.

There was a significant interaction between target emotion and prime emotion for the RTs, F (2, 54) = 3.44, $p < 0.05$, indicating that the prime emotion effected performance and interacted with target valence in the RTs. The error rates did not show an equivalent significant effect, $F<1$. The interaction between target emotion and prime emotion in the RTs was also found to reliably interact with SOA, F (8, 216) = 2.31, $p < 0.05$. From Figure 6.9, it is apparent that there is a considerable difference between conditions particularly at 10 ms, where the angry prime, negative target condition is slower than the other conditions (mean difference = 50, SEM = 9). The interaction between prime emotion, target emotion and SOA was not significant in the error rates, F < 1. In the errors there was a trend for fewer errors for emotional primes than neutral at 10ms for the responses to negative words. The errors, then, do not support the findings in the RTs but they also do not contradict the RTs.

There was a significant main effect of target valence on RTs, F (1, 27) = 14.07, p < 0.01, with positive words being responded to faster than negative words (mean difference = 20, SEM = 5). This was supported by a small trend in the error data for more errors for negative targets compared with positive targets (mean difference = 0.34, SEM = 0.38), although this was not close to being significant, F < 1.

Prime emotion reliably effected RTs, F (2, 54) = 3.3, p<0.05, with faster responses on trials with a happy prime (M = 479, SEM = 15) than with a neutral prime

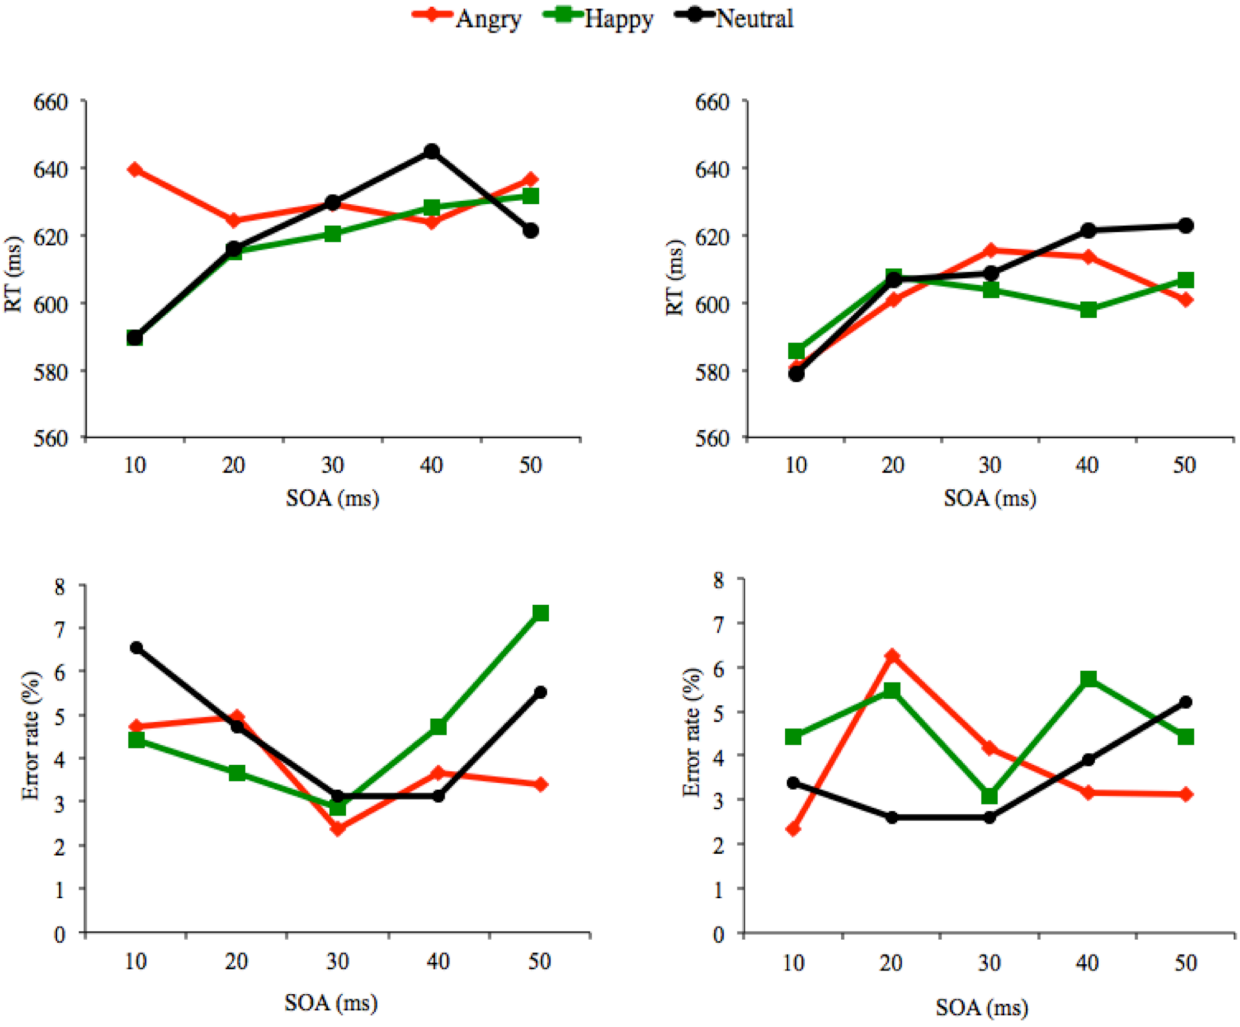

Figure 6.9: Mean RTs (top) and error rates (bottom) for the negative (left) and positive (right) target words, shown for each prime type (angry, happy and neutral) from Experiment 7.

(M = 484, SEM=16) or an angry prime (M=487, SEM=16). The error data did not show a significant effect of prime emotion, F (2, 54) = 2.12, although there was a trend for more errors on trials showing a happy prime (M = 4.99, SEM = 0.46) than a neutral prime (M = 4.28, SEM = 0.41) than an angry prime (M = 4.13, SEM = 0.51).

The RTs, as shown in Figure 6.9, appeared to show a roughly linear trend with the fastest performance at the shortest SOA, with RT generally increasing with SOA. The overall effect of SOA was significant for the RTs, F (4, 108) = 12.56, $p < 0.001$, and for the error rates, F (4, 108) = 2.95, $p < 0.05$. For the error rates, there was more of a quadratic trend with the fewest errors being made at a mid SOA time (30 ms). SOA also interacted reliably with the effect of prime emotion in RTs, F (8, 216) = 2.7, $p < 0.05$, and for the error rates, F (8, 216) = 2.19, $p < 0.05$. The effect of the prime, then, was affected by the SOA and so by stimulus visibility.

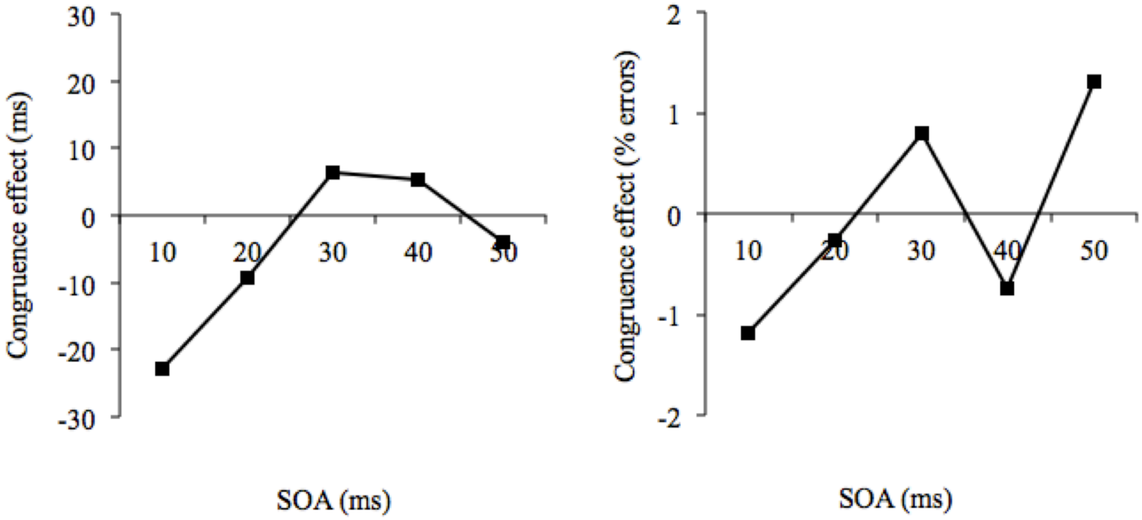

Figure 6.10: Mean congruence effect (incongruent trials – congruent trials) between prime and target at each of the SOAs for the RTs (left) and the error rates (right).

6.4.2.3. Priming effects: emotional primes only

The above analyses included all three prime types: neutral, happy and angry. The differences could reflect a mixture of the effects of response congruence and of emotional versus neutral primes. To isolate effects of response congruence per se, further analyses was run on the RTs and error rates including only the emotional primes.

The data are plotted as congruence effects in Figure 6.10.

The congruence effect – reflected in this analysis as the interaction between target emotion and prime emotion – was significant for the RTs, F (1, 27) = 4.98, $p < 0.05$ (M = 9, SEM = 4), but not for the error rates, $F < 1$ (M = -0.28, SEM = 0.51). The interaction between target emotion and prime emotion interacted reliably with SOA for the RTs, F (4, 108) = 4.63, $p < 0.01$. In the error data, the interaction between prime emotion, target emotion and SOA was not significant, $F < 1$. In the RTs it appeared that the significant result from the main analysis were not driven purely by differences between the neutral and emotional primes - there were significant differences between the happy and angry prime faces that interacted with both target emotion and SOA.

In the RTs there was a significant main effect of prime emotion, F (1, 27) = 5.44, $p < 0.05$, reflecting slower responses following angry prime faces than happy (mean difference = 8, SEM = 4). The main effect of prime emotion was not significant in the error rates, $F < 1$, and there was very little difference in accuracy between the two prime emotion conditions (mean difference = 0.07, SEM = 0.51). The main effect of target emotion was significant in the RTs, F (1, 27) = 13.03, $p < 0.01$, with generally faster responses to positive than negative words (mean difference = 20, SEM = 6). In the error rates, the main effect of target emotion was marginally significant, F (1, 27) = 3.57, $p = 0.07$, with generally fewer errors for the angry prime trials than the happy prime trials (mean difference = 0.73, SEM = 0.39).

6.4.3. Discussion

To summarise the findings of Experiment 7, there were significant effects of the primes valence on RTs as a main effect, and interactions with SOA. The effect of the prime valence did interact with the effect of the valence of the target word, in that responses to negative words were more affected by the prime emotion than positive words. In particular, responses to negative words were slower following a negative prime, and this interacted (reliably) with SOA (and so visibility). This ‘double negative’ effect was largest at 10 ms SOA. Responses to positive words showed smaller, less clear, effects of the prime. The visibility measures indicate that performance increases with increasing SOA, and so the largest effect of prime emotion on responses to negative words was when visibility was lowest.

However, the effect of the prime did not appear to be primarily due to congruence between prime and target valence, as was expected. The only point at which a large, convincing effect of the prime was present was highly specific, and would be better described as a ‘double negative’ slowing than as a congruence effect. That is, only when a negative prime precedes a negative target does it slow RTs. This is quite a complex finding. It could be a form of arousal effect: negative prime stimuli have been shown to capture attention (Mogg, et al., 1997; Algom, Chajut & Lev, 2004). In the present experiment that seemed to particularly be the case when a visible negative target was also used, somehow magnifying this capturing effect and slowing responses. This finding was not predicted (although it is at the SOA where I hoped to find evidence of emotion processing), but it is a considerable size (50 ms) and the relevant interactions are significant. The effect of the angry prime on overall RTs (independent of the target valence) was also significant, indicating an overall trend was present that could reflect arousal.

The findings of Experiment 7 gave some evidence of emotion processing of the prime: priming was found when the prime was a face and the target was a word, with the only overlap being the semantic, affective content. This argues against the possibility that priming was based on basic feature processing of the stimuli, as there was no featural overlap between the two.

The patterns in the error data were dissimilar to those in the RTs. There were very few significant effects in the errors rates, and those present did not reliably interact with target valence. Given this lack of reliability, and the absence of a clear pattern in general, not much of a conclusion can be drawn. There was a significant interaction between SOA and prime emotion, so a hint of some effect of the prime on error rates existed in the data. However, the error rates did not directly contradict the RTs – effects were mostly smaller, non-significant, or simply not present.

The visibility measures used both showed a general linear increase with SOA, and SOA reliably affected both measures of visibility. Overall, the two visibility measures give an equivalent picture of visibility, at least in how visibility interacts with SOA, and both measures appear to be equally valid for measuring visibility. However, participants did generally find giving two responses (one 2AFC and one a rating of 1-7) cumbersome so taking both (at the same time) may not be entirely practical.

Overall, visibility was reliably above chance. Higher visibility may, in the experiment, have been due to some extent to the change in masking conditions – in this experiment, the target words were embedded into the masks (and at variable intervals) to keep the time between prime and target constant. The presence of the word in the mask could have reduced the masks effectiveness somewhat. Given the higher levels of visibility here than in Experiment 5 or 6, the larger effect of prime emotion on RTs could be explained by the greater visibility, though the "double negative" effect is quite different to the congruence effect seen for angry primes in Experiment 5, and it is unclear why it should be maximal at the lowest levels of visibility.

To conclude, in Experiment 7 the emotional valence of a low visibility prime face did affect RT for classifying the emotional valence of a target word by as much as 40ms, but only at the two shortest SOAs. However, there was little sign of the emotional congruence effect seen in Experiment 5. Rather the effect was that classification of the target word was slowed by an angry face prime, but only when the target valence was also negative. Hence the data suggest that low visibility emotional primes can influence behaviour, if not in the way expected. Prime visibility overall was clearly higher than in the previous experiments (though the primes were still hard to see), though it is not clear why this should act to prevent the expected congruence effect.

6.5. Experiment 8

Experiments 5 and 7 provided evidence that emotional face primes can affect responses to subsequent stimuli. This was true even when visibility was low, the primes never occurred as targets and (in Experiment 7) the targets (words) were perceptually very different to the primes. However only in Experiment 7 was a substantial effect obtained. Nor was the effect at the shortest SOA the expected congruence effect; indeed there was little sign of an emotional congruence effect at any SOA. Instead the effect of prime emotion at the shortest SOA was a slowing of responses to an emotionally negative target word following an angry face prime: what I am calling a "double negative" interference effect.

As the double negative effect was both striking and novel, Experiment 8 was designed to see if could be replicated. Some changes were made. A smaller and lower range of SOAs was used to reduce the overall visibility of the prime faces, nor were participants told about the presence of the prime faces. Both of these measures should increase confidence that any priming effects are due to subliminal primes, and cannot be sensibly explained by learnt S-R mappings or partial awareness.

The location of the target word on screen also varied between one of four locations. This was manipulated in case any attention effects were present that could increase the effect of the mask, or draw attention to, or push attention away from, specific parts of the screen where specific features of a face had previously been presented. The location of the word was constant in Experiment 7, and the words were considerably smaller than the prime faces, so it is possible that the findings were at least partially due to the drawing of attention to or away from the centre of the prime face. Whilst the target was presented after the prime, the location of the target may still direct attention to where a mouth, a nose, or an eye was very recently. If processing of the faces was featural (as is likely given the congruence effect with inverted faces in Experiment 5), then it is plausible that performance would vary with the location of the target word relative to these features. In Experiment 7 the location of the target was specifically over the centre of the figure (the nose), so changing this to being over other features or towards the edge of the face may influence whether priming is found. This issue was not relevant in Experiments 5 and 6 as the target stimuli were faces of a similar size to the primes and so fully overlapped.

As another change, blocks were included requiring forced choice to emotionally neutral targets, similar to those in Experiments 5 and 6: participants were asked to classify the valence of an emotionally neutral target word following an emotional prime face. In Experiment 8, there was an emphasis on speed of response allowing some comparison of the RTs in these response bias blocks to those in the emotional target blocks. The subjective measure of awareness used in Experiment 7 was omitted, as it is hard to interpret and awkward for participants to give two responses on every trial. In addition, there was, as in the prior experiments, another subjective measure available – whether they noticed the faces during the first phase of the experiment.

6.5.1. Method

6.5.1.1. Participants

Twenty-four participants took phase in exchange for credit and up £1 as a bonus (depending on performance). Seven were male and seventeen were female, with an average age of 22 (SD = 6) years. All gave informed consent before participating.

6.5.1.2. Materials

The stimuli used (prime faces, masks, and target words) were identical to Experiment 7.

6.5.1.3. Procedure

The experiment was split into two phases: priming (including target priming and response bias blocks), and a visibility test. The priming phase was completed first, with one response bias block for every 3 blocks of target priming. Before the start of the experiment, two short blocks of practice were given (24 trials each), one for the target priming task, and one for the response bias task. The visibility test was given at the end, after participants had been told about the prime faces.

The trial outline for the priming phase of the experiment was similar to Experiment 7, with the following changes. Only 3 SOAs between the prime onset and the backwards mask onset were used: 10, 20 and 30ms. The location of the target word varied randomly between 4 locations: 40 or 80 pixels above or below the fixation point (0.55 and 1.1^{o} above and below fixation). This corresponded with 4 locations on the prime faces: the top of the forehead, the eyes, the mouth, and the bottom of the chin.

In the target priming blocks, the target words were either positive or negative, and participants were asked to classify the words' valence as quick and accurately as they could. Prime faces could be happy, angry or neutral. Feedback was given (a red cross appeared on screen) if a wrong response was made. Each block lasted for 48 trials, and there were 12 blocks in total. Each trial type (each combination of target location, SOA, target emotion and prime emotion) occurred 8 times. Each prime identity (of which there were 32) occurred 18 times, 6 times with a happy expression, 6 times with an angry expression, and 6 times with a neutral expression. The SOA, location, and target type each identity was paired with was controlled across participants. I.e. for one quarter of the participants each prime identity would occur once per target valence and

SOA, for only one of the four locations. For the other participants, the same would apply but the specific location would be different. Each target word occurred 6 times.

For the forced choice blocks, each trial was identical to the classification RT phase of the experiment except the target word was emotionally neutral, and prime faces either happy or angry (not neutral). Participants were informed hat there would be a cue present to indicate whether the face was positive or negative, but this was very subtle and would be difficult to detect. They were instructed to try and respond with their ‘gut instinct’ as quickly as they could. Feedback was given based on the prime face emotion. Each block lasted for 48 trials, and there were 4 blocks in total. Each trial type (target location, SOA, prime emotion) consisted of 8 trials. Each of the 96 possible target words were presented twice, and controlled so that 48 occurred in the first block, a different 48 in the second block, then repeated for the 3rd and 4th blocks. The prime identities (of which there were 32) occurred a total of 6 times each, 3 times with a happy expression and 3 times with an angry expression, with each identity and emotion presented once per SOA. The pairing of identity and location was controlled across participants in the same way as the priming phase of the experiment (described above).

Once the priming and response bias blocks were completed, the participants were asked if they had noticed anything in addition to the ‘grey squares’ (the masks) and the target words. They were then informed of the presence of the prime face, and instructed to respond as accurately as they could to whether the face was upright or inverted. No feedback was given for this phase of the experiment, which lasted for 6 blocks of 48 trials. Each prime face identity occurred 9 times, and the pairing of identity to SOA and prime emotion was controlled (i.e. each identity occurred at each SOA and for each emotion) and the pairing of target location and identity was again controlled across participants in the same way as the priming and response bias phase of the experiment. Each of the neutral words occurred 3 times and was randomly selected across the experiment (i.e. not linked to the trial type).

At the end of each block in the experiment a bonus was awarded based on a combination of speed and accuracy for the priming blocks, and for accuracy in response bias and visibility blocks. For the classification RT blocks, a score was calculated with 0 points for wrong responses, 1 point for correct responses with RT longer than 500ms, and 2 points for correct responses with RT less than 500ms. A score of 64 or higher (out of 96) yielded a bonus. For the forced choice blocks, a score of 28 or higher yielded a bonus. For the visibility test a score of 30 or higher gave a bonus. Each bonus was worth 5p.

The response mappings (which button and finger was used to respond to positive and negative) was balanced across participants with 12 using the left index finger to respond to negative and right index finger to positive, and the other 12 the opposite mapping.

6.5.1.4. Data analysis

Data were analysed in the same manner as the Experiments 5-7. The total number of RT trials that were faster than 200 ms or slower than 2000 ms was < 1 %. Appendix 3 gives the median RTs, and Appendix 2 reports an items analysis.

6.5.2. Results

Five participants commented that they noticed the face during the priming phase of the experiment. These participants were removed from the analysis to minimise any chance of awareness of the stimuli leading to any of the effects of interest. The sample size for the following analyses is 19 as a result, with 10 participants with response allocation one, and 9 for response allocation two. All 19 not only failed to notice the presence of the prime face, but also expressed surprise when told of its existence.

Based on a preliminary analysis, the four locations factor was collapsed to "inner" and "outer" locations, as this appeared to capture its effect, and reduces the complexity of the analysis.

6.5.2.1. Visibility

The overall accuracy for discriminating the orientation of the face during the visibility test was d' = 0.18 (SEM = 0.1). The difference from zero d' was only marginally significant, $t(18) = 1.72$, $p = 0.1$, and is so low that, coupled with the fact that none of the included participants noticed the presence of the faces during the priming phase of the experiment, makes it likely that any priming effects were due to subliminal processing of the emotional expression.

Figure 6.11 shows the results for the two location types (inner and outer), for the three SOAs. ANOVA on the results, breaking performance down for each of the trial types, found no significant main effect of location, F<1, prime emotion, F<1, or SOA, F(1,18)=2.18. There was no significant interaction between location and SOA, F<1, or SOA and emotion, F<1. The interaction between location and emotion was marginally significant, F(1,18)=2.89, p=0.07. The interaction between location, prime emotion and SOA was not significant, F<1. Figure 6.11 shows that performance was low for both locations and all SOAs, with a slight trend for lower performance when the neutral word (a mere distractor in the visibility test) appeared at an inner rather than an outer location (mean difference d'=0.04, SEM=0.08).

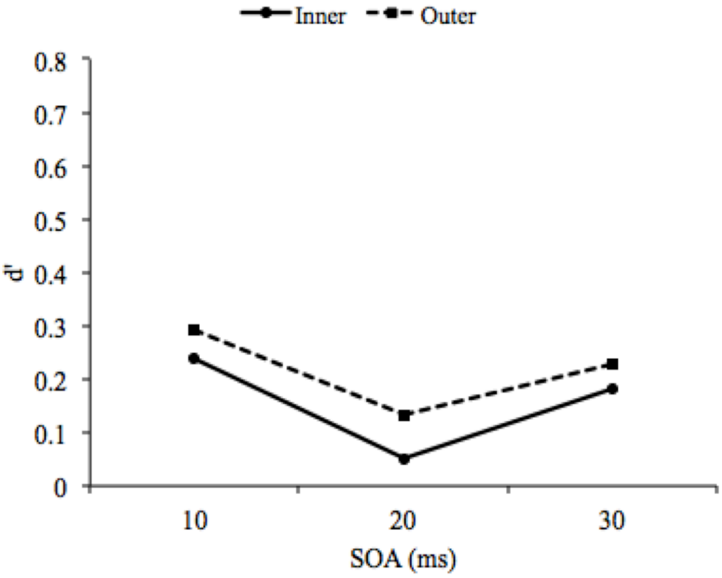

Figure 6.11: Group mean discrimination performance (d') of prime face orientation for the inner and outer target word location conditions from the visibility test in Experiment 8.

6.5.2.2. Target priming trials

Figure 6.12 displays the RTs and error rates for each type of target-prime pair (collapsed across SOA and location). The prime emotion by target valence interaction was not significant in the RTs, F<1, but it was in the error data, F (2, 46) = 8.26, p<0.001, indicating a congruence effect was present for the errors. From Figure 6.12, it appears there is a trend for neutral (in comparison to emotional) primes being responded to with less accuracy than the emotional primes for the positive target words. Overall, responses on trials with an angry prime are generally more accurate, regardless of target valence. Responses to negative words show fewer errors for neutral and angry primes than happy – an effect that appears to represent a congruence effect. There is very little difference in the RTs across prime-target pairs.

The interaction between target valence, prime emotion and SOA was not significant in the RTs, F < 1, or the error rates, F < 1, and there was no significant main effect of the prime in the RTs, F < 1, or the error rates, F (2, 46) = 1.28. There was no significant interaction between prime emotion and SOA in the RTs, F (4, 72) = 2.02, or the error rates, F (4, 92) = 1.72. The main effect of target valence was not significant for the RTs, F < 1, and there was only a very slight trend for faster responses to positive than negative words (mean difference = 4, SEM = 6). In the error rates, the effect of target emotion was not significant, F < 1. There was a marginally significant main effect of SOA in the RTs, F (2, 36) = 2.98, p = 0.06, with faster responses at 30 ms (mean = 558, SEM = 11)

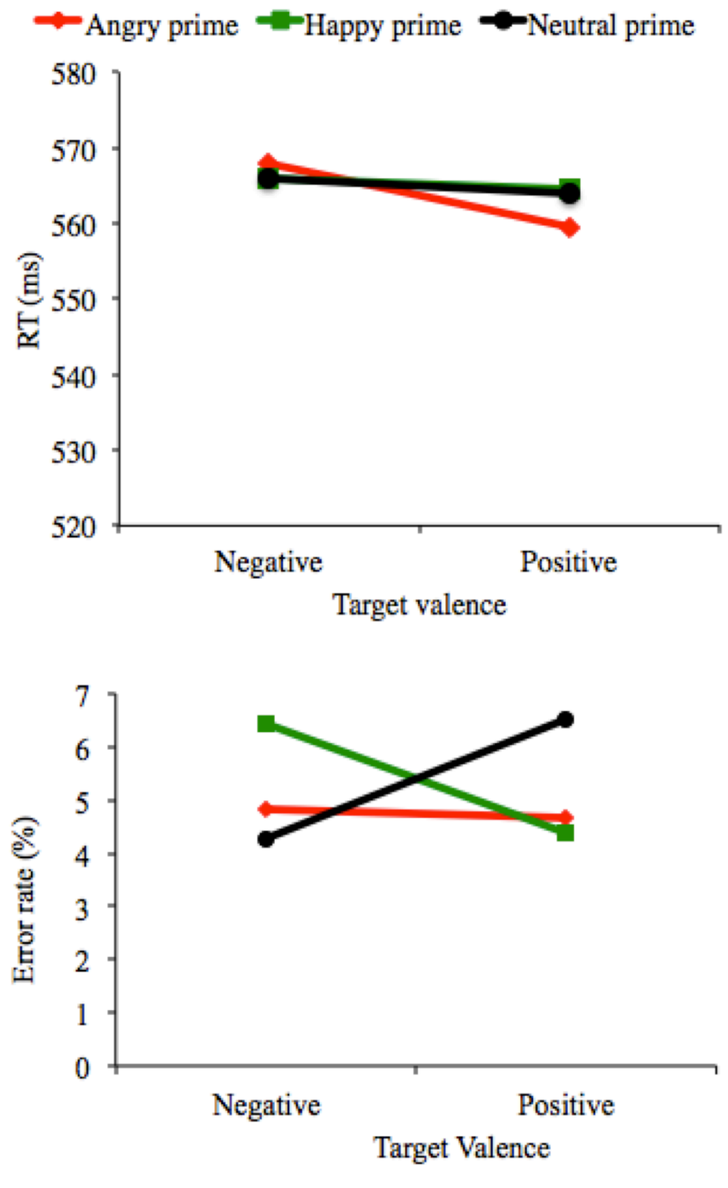

Figure 6.12: The interaction between target valence and prime emotion for the RTs (top) and the error rates (bottom) from Experiment 8.

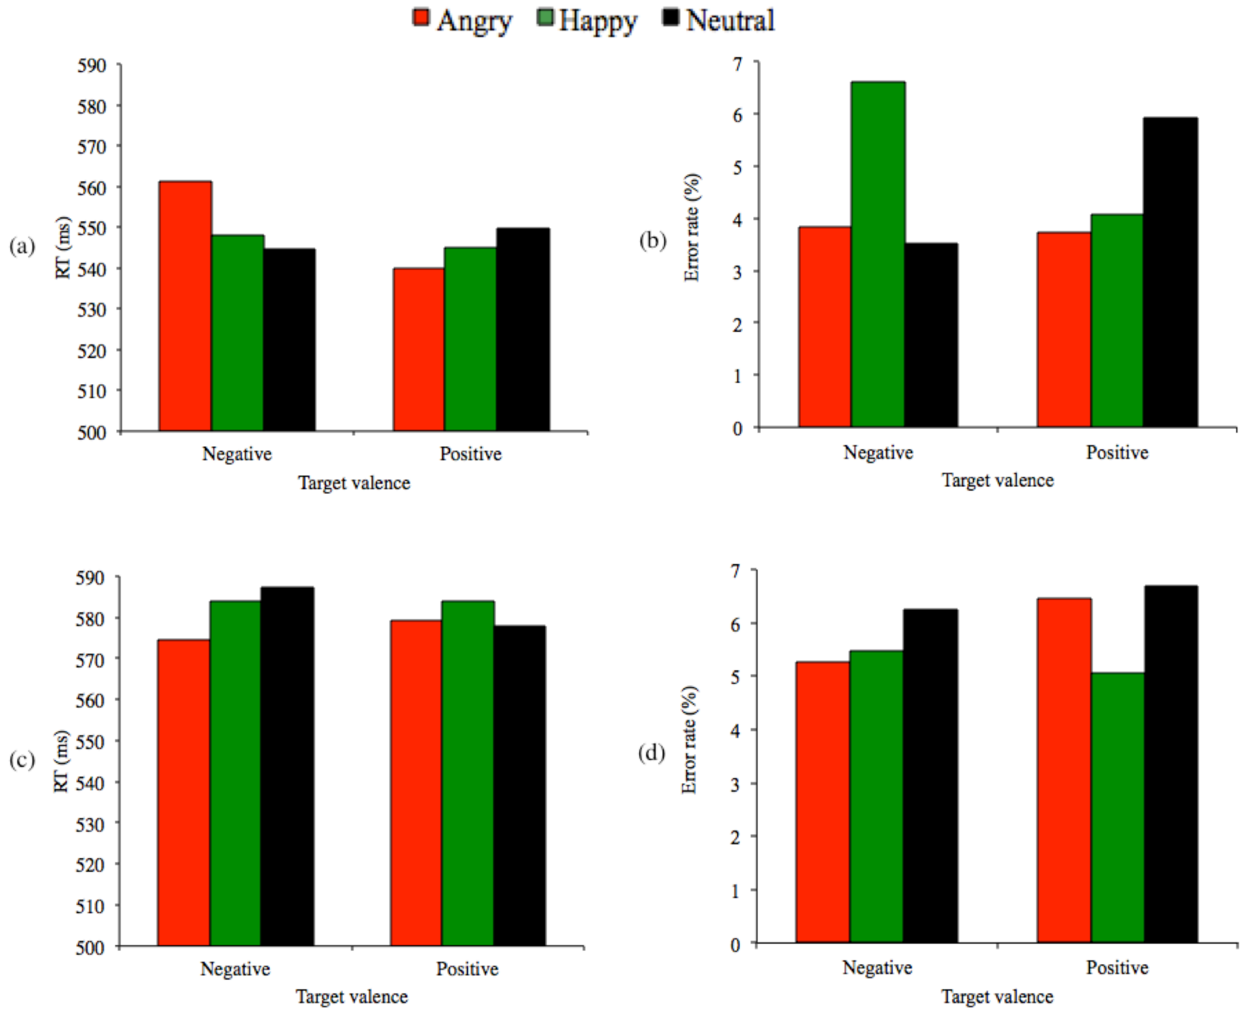

Figure 6.13: The interaction between target valence and prime emotion for the inner (a & b) and outer (c & d) locations for the RTs (a & c) and error rates (b & d) from Experiment 8.

than 10 ms (mean = 568, SEM = 12) or 20 ms (mean = 567, SEM = 12). There was no significant main effect of SOA for the errors, F<1.

So far, the results have indicated that there is essentially no significant effect of the prime emotion on RTs and very little on error rates. Inasmuch as there was an effect on error rate, it was in the direction of lower error rates with an affectively congruent prime for a negative target: i.e. there is no sign of the "double negative" effect identified in Experiment 7. However, the location of the target may have affected the results. The results are shown in Figure 6.13 broken down into the ‘inner’ (locations 2 and 3) and ‘outer’ (locations 1 and 4) locations. In an ANOVA run with the two location types (inner and outer), the interaction between target valence, prime emotion, and target location was significant in the RTs, F (2, 36) = 4.08, $p < 0.05$, although not for the error rates, F < 1. The interaction between target location, SOA, target valence and prime emotion was not significant in the RTs, F (4, 72) = 1.6, or the error rates, F (4, 72) =

1.25. The location of the target word on screen, then, appeared to affect the interaction between prime emotion and target valence for the RTs. There appeared to be no significant effects on the error rates. Figure 6.13 indicates that the interaction may be sourced from slower responses on 'double negative' trials than the happy or neutral trials for the inner locations (mean difference = 15, SEM = 6), but if anything the opposite trend for the outer locations where trials showing an angry prime generally led to faster responses than happy or neutral primes (mean difference = 11, SEM = 4). For the errors, there is no clear pattern apparent although for the inner location there appears to be evidence for a congruence effect for negative targets (more accurate responses on neutral and angry prime trials than happy prime trials – mean difference = 2.9, SEM=0.8), with the same pattern largely absent for the outer locations (mean difference = 0.8, SEM=0.6).

Not surprisingly, there was a highly significant main effect of target location on the RTs, F (1, 18) = 81.56, p < 0.001, and this was also significant in the errors, F (1,18) = 4.81, p<0.05. As can be seen in Figure 6.13, responses were slower and less accurate

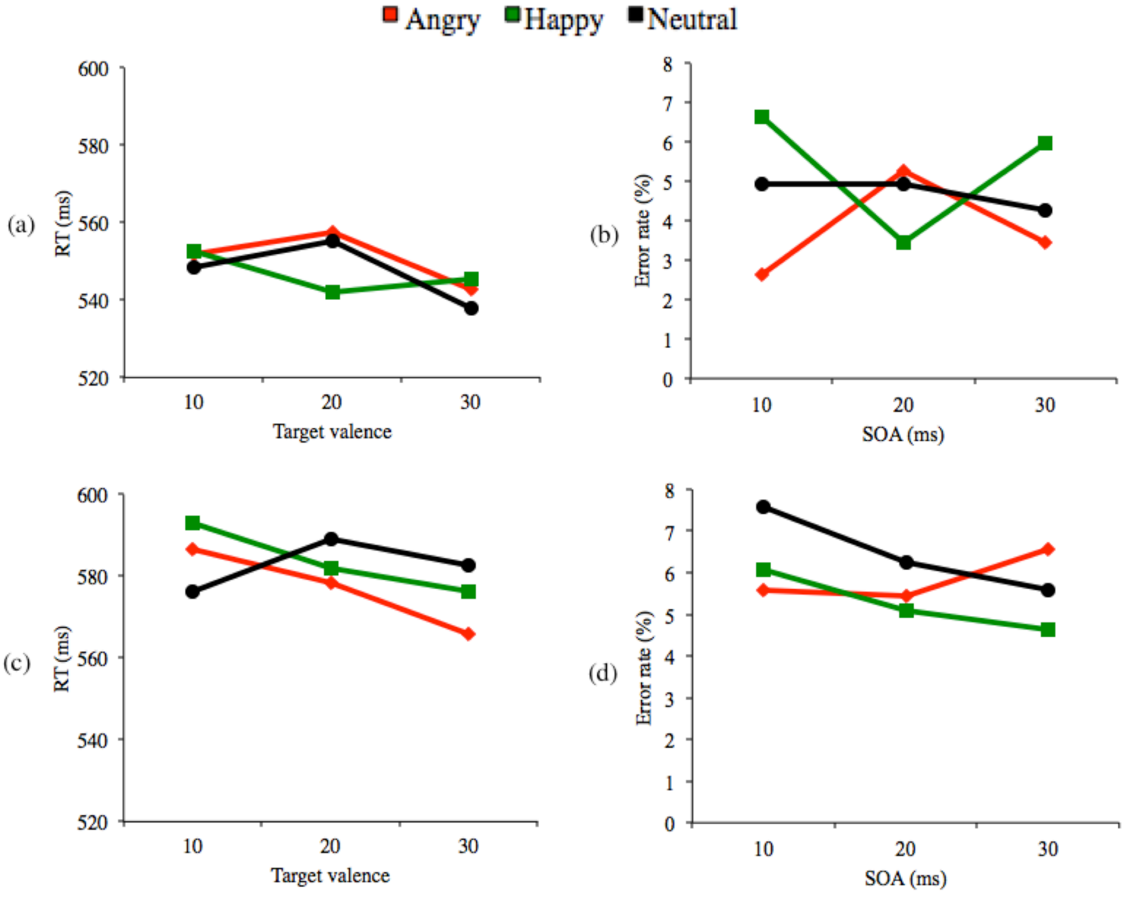

Figure 6.14: The interaction between prime emotion and SOA for the inner (a & b) and outer (c & d) locations for the RTs (a & c) and error rates (b & d) from Experiment 8.

responses for the outer target locations than the inner target locations (mean difference between the two pairs for the RTs =33, SEM=3, and the error rates = 1.13, SEM=0.51).

Figure 6.14 shows the results collapsed across target valence, for each of the two location types and for each SOA. There was no significant interaction between prime emotion and target location in the RTs, F (2, 36) = 1.14, and a marginally significant interaction in the error rates, F (6,138) = 2.52, p = 0.1. The interaction between prime emotion, target location and SOA was not significant for the RTs, F (4, 72) = 1.12, but was significant for the errors, F (4,92) = 3.31, p < 0.05. There was, then, very little evidence that the prime emotion affected responses independent of the target valence. From Figure 6.14 it appears that responses were slower and less accurate for the angry (than happy) prime trials particularly at 20ms SOA for the inner locations (RT mean difference =15, SEM=6; error rates mean difference =1.8, SEM=0.8) whereas the angry prime trials for the outer locations were generally faster and with similar numbers of errors than the happy prime trials (RT mean difference =4, SEM=7; error rates mean difference =0.49, SEM=1.29). These effects were very small, and not significant in the ANOVA.

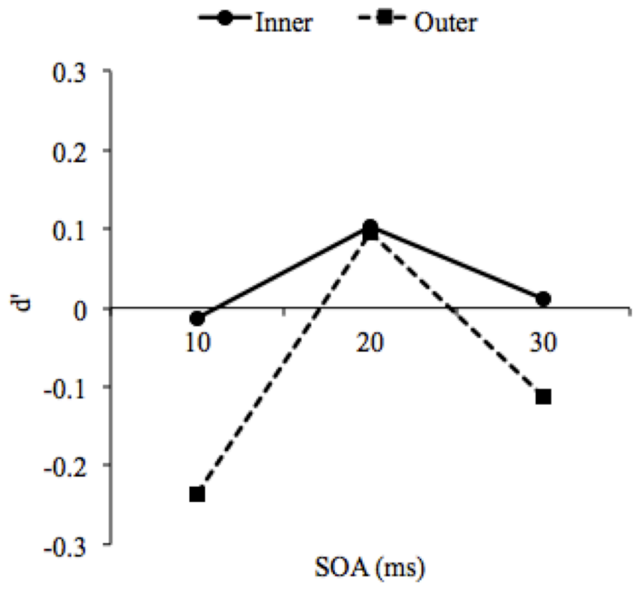

Figure 6.15: Group mean discrimination performance (d') for the decision bias blocks from Experiment 8.

6.5.2.3. Forced choice neutral trials

For accuracy, the overall performance was at a d' of 0.03 proportion correct (SEM=0.04), and this was not reliably different from chance, t (18) = -0.4. Figure 6.15 shows performance for each of the two location types (inner and outer). An ANOVA was run on the accuracy data with the factors of target location and SOA. There was no significant main effect of target location, F (3, 69) = 1.41, or SOA, F (2, 46) = 2.16, and no significant interaction between location and SOA, F<1. There was no hint, then, that performance was reliably above chance or was reliably affected by target location or SOA. No evidence of any bias for the response bias blocks was evident.

The effect of target location, prime emotion and SOA on RTs is shown in Figure 6.16. An ANOVA was run on the RT data with the factors of target location, prime emotion and SOA. There was no significant main effect of SOA, F < 1, or prime emotion, F < 1, although there was a slight trend for faster responses to angry primes (mean difference = 10, SEM = 13.36). The main effect of location was marginally significant, F (3, 69) = 4.28, p = 0.05, reflecting generally slower responses to the outer locations (M = 26, SEM = 18). There was no significant interaction between SOA and prime emotion, F < 1. The interaction between prime emotion and target location was not significant, F < 1, and did not reliably interact with SOA, F < 1. The interaction between location and SOA was very marginally significant, F (6, 138) = 2.47, p = 0.1. From Figure 6.16 it would appear that the two locations show opposite RT patterns with SOA: the inner location shows fastest responses at 20 ms (mean difference = 18, SEM = 23), the outer the slowest responses at 20 ms (mean difference = 28, SEM = 24). Other than this, there is very little evidence of any effect of the prime.

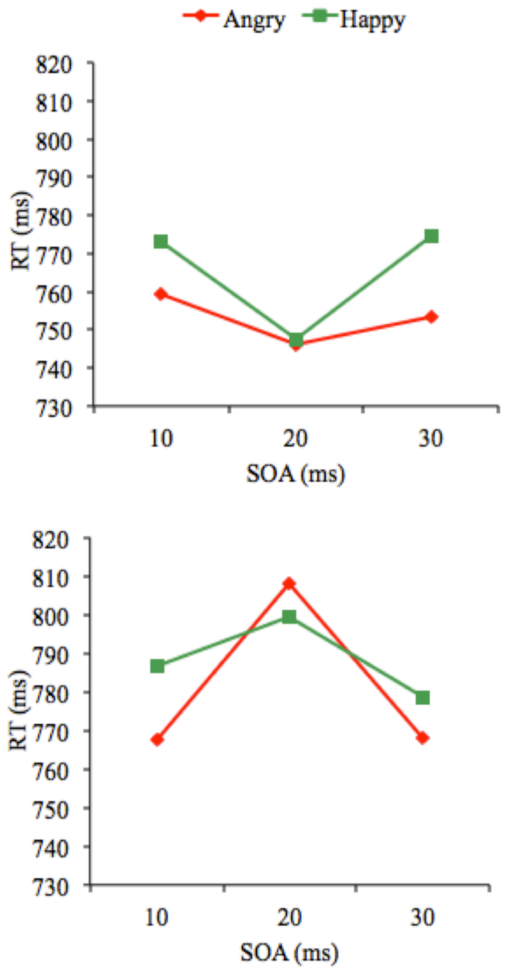

Figure 6.16: Group mean RTs for the inner (top) and outer (bottom) locations for the decision bias blocks from Experiment 8.

6.5.3. Discussion

The main aim of Experiment 8 was to replicate the double negative effect found in Experiment 7. This was rather limited in its success: the double negative trials were the slowest, but only for the 'inner' target word locations and the size of the effect was smaller. Moreover, for the outer locations, if anything the opposite pattern was found, questioning the reliability of the effect for the inner locations. In Experiment 8, unlike the previous experiments, there were substantial sized effects of the primes in the error rates, and these appeared to be congruence effects for the negative targets only.

The findings from Experiment 8 showed a complex interaction between target emotion, prime emotion, and the location of the target word. An effect of the prime is evident, despite none of the participants included in the analysis noticing the presence of the prime face. This, along with low visibility levels in the visibility test, make it likely that any effect of the prime was subliminal. In addition, visibility did not reliably vary with location of the target word, or the SOA.

Comparing the two target word locations (inner and outer), there were substantial differences in the priming effects found. One possible interpretation of the interaction of priming effects with target word location is that priming requires a degree of spatial overlap between the prime and target, and in the ‘outer’ location condition in this experiment the words were sufficiently removed from the emotional phase of the face that this did not occur. The inner locations, by contrast, corresponded with the location of the eyes and the mouth. If overlap is needed, then it may be that we can extend this is to meaning overlap with the informative phase of the stimulus, not any phase of the stimulus.

The error rates in this experiment are the only in this series of experiments to give significant interactions with the primes. In particular for the double slowing effect for the inner location, there were slower but also more accurate responses for the angry-negative trials than for the happy-negative condition. This finding hints at a speed accuracy trade off – leading to faster and more accurate responses for negative prime and negative target trials. However, in Experiment 7, the ‘double negative’ effect was specific to the RTs – there was no hint of an increase in accuracy.

The neutral target trials failed to demonstrate any detectable effect of the face primes on forced classification of the emotional valence of a neutral word. With the design of this experiment, then, it appears that there was essentially no effect of the prime on choice. The effect of the prime, then, was limited to priming responses to emotional target words.

To summarise the main conclusion from Experiment 8, the double negative effect was replicated, but only to a limited extent. There were significant effects of the error rates whereby for negative target words responses were more accurate if the prime was angry or neutral, compared with happy. There was very little evidence of any effect of the primes for the response bias blocks, in terms of RTs or choice.

6.6. General Discussion

Of the four experiments reported here, three show (at least some) significant effect of a low-visibility emotional face prime and these will be the focus of this discussion. In Experiment 5, congruence effects in the RTs were significant, but quite small and predominantly sourced from responses to angry target faces, with no consistent effects for happy face responses (although the difference in congruence effect was not significant). No significant effects were found in the error rates. An effect on decision bias was found when the targets were neutral faces, with participants more likely to choose angry following an angry prime, or happy following a happy prime. For inverted faces, the priming effects were comparable to those of upright faces, but the decision bias effect tended to be smaller and not significant for inverted primes. In Experiment 7, with word targets, a significant and substantial priming effect was found but it appeared to be different in nature to Experiment 5. There was a particularly large difference between happy and angry prime trials when participants responded to negative words, which was overall 41ms. There was no equivalent effect on error rates for the same condition. Experiment 8 found again different results, with the effects of the prime varying depending on the location of a target word. The error rates showed relatively large effects for the responses to negative target words, with generally more accurate responses when the prime was angry and neutral, compared with happy for the inner location (2.9%), but less so for the outer location (0.8%).

There are, in short, substantial differences between the experiments. One possible argument here is that there were no real priming effects, and the experiments simply picked up on experimental noise. Given that the priming effects were at least partially significant in Experiments 5, 7 and 8, the relative consistency for larger/more significant effects for responses to negative targets, and the differences in design between experiments, it is also possible that the effects were real with the differences due mostly to a combination of congruence and arousal effects, and the effect of varying visibility/knowledge of the prime. I will start with discussing the latter, and then describe how congruence and arousal may have led to the findings.

In Experiments 5 and 8, participants were not told about the presence of the prime face, and only participants who did not notice the presence of the primes were included in the analysis. Both of these experiments showed small, but significant (at times only marginally significant), effects of the primes on RTs and (for Experiment 8) on error rates. Experiment 7 used a range of visibilities and participants were explicitly told of the presence of the primes. It is certainly plausible that larger priming effects were found at least partly as a result of this difference in visibility and knowledge of the primes. Prior research has shown that lowering visibility could lead to smaller (or even non-existent) priming effects (Andrews, Lipp, Mallen & Konig, 2011). In addition, low visibility primes that are not used as target stimuli and are not shown consciously at any point have also been shown to give smaller, but significant, priming effects (Finkbeiner & Friedman, 2011). The differences in priming magnitude between the experiments reported in this chapter, then, are not necessarily that surprising. This explanation implies that whilst low visibility, target-independent priming can be found, it tends to be smaller in magnitude. Given the potentially small effects, designs may need to be highly sensitive to these differences, and failures in the literature to find significant effects (and, indeed, the failure of Experiment 6 in this chapter) are not that surprising. A sensible conclusion here may be that low visibility, target-independent emotional primes can influence behaviour, but only to a rather small extent.

A second point that could explain the differences between experiments would be that there are two ways emotional primes could lead to effects on target classification: through response congruence, and through arousal. I would argue that a combination of arousal and congruence led to the effects found in this chapter. Experiments 5, 7 and 8 all showed larger RT effects of prime emotion for the responses to negative targets, regardless of whether the target was a face (Experiment 5) or a word (Experiments 7 and 8). Even in Experiment 5 where significant congruence effects were found, the congruence effect was largest for the responses to negative target (although the difference between negative and positive responses was not significant), again indicating negative target responses may lead to larger effects. The highly specific double negative effect in Experiment 7 may be due to a magnified threat alert when two negative stimuli are shown in succession, which is more prominent than when one is shown in isolation. The threat alert could lead to more cautious processing, and so produce the double negative slowing in Experiment 7. If there is some combination of congruence and arousal effects, then given the differences in methodology, the extent either one influences behaviour could certainly vary.

The differences in the effect of the primes between experiments can also be described as due to the differences in methodology. Most notably, Experiment 5 used faces as primes and targets, where as Experiments 7 and 8 used words as targets and faces as primes. This manipulation of the relationship between the prime and the target could have altered the type of priming found: a combination of congruence and arousal with similar primes and targets, but shifting to more arousal than congruence when the relationship is restricted to an overlap in general valence. Alternatively, the prime faces in Experiments 7 and 8 could have been more arousing. Given this fairly major difference, some change in the nature of priming would not be that surprising. Again, some differences were also present between all experiments in the visibility of the stimuli, and there were other differences in design as well (e.g. varying the location of the target word in Experiment 8).

One implication of the results of the reported experiments is that low visibility emotional priming is not a consistent phenomenon. Changes in design did change the results, and it is not the case that there is a robust, replicable, substantial effect from low visibility emotional faces. This may seem counter intuitive given that one of the main theories of low visibility processing of emotions postulates a dual route model, and is based around emotions being processed rapidly to aid threat detection (e.g. LeDoux, 1992). However, given the nature of what the experiments were trying to achieve – to measure some effect of the prime emotion on a subsequent, unrelated response – finding relatively small effects may not be that counter-intuitive.

We can safely assume that where emotional priming effects have been seen in these experiments, they reflect processing of the emotion rather than S-R learning, because the primes were not used as targets, or clearly seen at any point (especially in Experiments 5 and 8). In Experiments 7 and 8 the prime stimuli had the added virtue of being perceptually completely different to the target word stimuli, thus ruling out priming at the level of perceptual features.

Given the variation in findings, the results from the experiments reported in this chapter cannot give definitive support for clear, consistent, behavioural effects of a low visibility prime. However, they do indicate meaningful (semantic) processing of low visibility primes can occur, under some conditions at least, and this can lead to effects on behaviour. The precise nature of this effect appears to be moderated by the visibility of the stimulus, knowledge of the primes presence, location of the target stimulus, and the relationship between the prime and target stimulus. This is complex, and more research is certainly required to establish beyond doubt the nature of low visibility emotional primes processing and its influence on behaviour.

CHAPTER 7

Summary, methodological reflections, and future directions

7.1. Introduction

This thesis aimed to investigate processing of the emotional content of low visibility faces and body posture images organised around three themes: V1's role in visual awareness and emotion processing, the effect of visual masking on decision making about emotional stimuli, and subliminal emotional priming. The findings were generally mixed, with some evidence for semantic emotion processing of subliminal (or at least very low visibility) stimuli, and for emotion processing not being dependent on V1. However, having found these effects I then generally failed to find equivalent effects with different designs. This chapter aims to briefly summarise the main findings of this thesis, and discuss the interpretations, limitations and future directions. The emphasis of this brief final chapter is on explaining the findings that were in line with the predictions, and addressing the issues surrounding the failed extensions.

The results from the experiments reported in this thesis will first be summarised for each theme, then the following will be discussed: the inconsistencies found in the experiments reported, alternative explanations of the evidence for privileged emotion processing, the measures of awareness used, and future research directions.

7.2. Theme summaries

7.2.1. V1's role in visual awareness and emotion processing

Experiment 1 used TMS to disrupt processing in V1 during a critical interval (70-140ms) following a body posture stimulus. Participants learnt to classify eight body posture images as belonging to one of two apparently arbitrary categories. In fact, half the time there was a redundant cue of emotion to the decision. The stimuli were presented either at the location on screen of a phosphene, or a control location in the opposite hemifield. TMS selectively disrupted processing of neutral body posture images at the phosphene location, and this was true over the whole range of pulse onsets (although the disruption did diminish at 120-140ms). In contrast, categorisation of the emotional body posture images (for which the emotion provided an additional cue to response category) was relatively robust to disruption.

Experiment 2 used a similar design to Experiment 1, but with face stimuli, offline cTBS to V1, and instead of relatively arbitrary categories, participants performed a gender discrimination task, with (as in Experiment 1) emotion providing a redundant cue for half the decisions. The gender task was chosen because it had the advantage of not requiring a learning phase, and also allowed the use of a larger set of stimuli. The risk was, however, that because facial features associated with gender are already well learned as a cue for this decision, learning of the redundant emotion cue might be blocked. Indeed there was no evidence that the emotion cue assisted the decision. A greater problem, however, was that there was no detectable effect of the cTBS stimulation on the gender discrimination at all. The most likely reason is that the cTBS stimulation protocol was ineffective in disrupting processing locally in V1.

7.2.2. The effect of visual masking on decision making about emotional stimuli

Experiments 3 and 4 both used visual masking to disrupt the processing of faces, and measured the effect of that disruption on several type of explicit discrimination. The strength of the masking was manipulated by varying the SOA between the target face and backwards mask (between 10 and 130ms). In both experiments, the expected monotonic masking function was obtained: ability to discriminate properties of the faces was at its lowest point for the shortest SOAs and increased with increasing SOA. In Experiment 3, using computer generated faces, there was an advantage for emotion discriminations (happy or angry) over orientation (upright or inverted) discriminations at the shortest SOA, despite the orientation discrimination being performed much better at higher levels of visibility. The effect was strongest for the upright faces. An attempt to exploit the computer-generated faces to match an age discrimination task to the emotion discrimination task in difficulty, using on-line calibration, was thwarted by the age discrimination being too difficult even for the extremes of the age range.

Experiment 4 used a similar design to Experiment 2 but using masking rather than TMS. Participants performed a gender discrimination task, with emotion providing a redundant cue to the decision for half of the stimuli. The hope was to find an advantage in classifying the emotional subset of stimuli at short SOAs, due to sparing of emotion processing at short SOAs. An appropriate masking function was obtained, but there was no evidence of any advantage for the emotional stimuli. It seems likely that this was because there was no learning of the S-R association for the emotional cue in this experiment (and in Experiment 2) because it was ‘blocked’ by the 100% predictive gender cue. (Unfortunately, the two experiments were run concurrently, so the outcome of neither could forestall the running of the other.)

7.2.3. Subliminal emotional priming

Experiments 5-8 used a response-congruence masked priming paradigm with a low visibility emotional face prime preceding classification of the emotional valence of a face or word target. All the priming experiments used prime images that were never used as targets, nor clearly seen at any point, to rule out S-R learning as an account of any response congruence.

Experiment 5 used faces as targets, and found a small but significant effect of congruence between the primes and targets, for angry target faces only. Experiment 6 used a similar design but with different prime stimuli and failed to show any substantial or significant priming effect. Experiments 7 and 8 both manipulated the visibility of the prime by varying the SOA between the prime and backward mask, and used word targets. Experiment 7 found little sign of a congruence effect, but a marked slowing of RT at the shortest SOAs when an angry face prime preceded a negative word target - a “double negative effect.” Experiment 8, in contrast, found little evidence this of effect on RT (except for a hint of such an effect for central word target locations only) but lower error rates for double negative – which could be interpreted as prime congruence facilitating accuracy for negative target words only or perhaps double negatives triggering more cautious responding.

Experiments 5, 6 and 8 included forced choice blocks where the target stimuli were neutral, but participants were still asked to classify them as happy/angry or positive/negative. Experiment 5 found a small but significant effect of the prime emotion on the choice made when the prime was upright. Experiment 6 found no evidence of any effect. Experiment 8 time pressured the choice, and found no significant effects on choice or RTs.

7.3. Why where the results inconsistent?

7.3.1. Methodological difficulties

The effects this thesis set out to find – relative sparing of emotion processing when processing in V1 was disrupted by TMS, or when masking was used to disrupt visual processing, were always likely to be small. Indeed, the effects found in Experiments 1, 3, 5, 7 and 8 that were significant were relatively small in size. Seeking effects that are likely to be relatively subtle poses difficulties for the methodology: the measurements need to be sufficiently sensitive, the noise minimal, and the general design effective and sufficiently persuasive to elicit a convincing effect. This latter point is perhaps the most difficult, and the one the rest of this section will focus upon.

To demonstrate that emotion processing is relatively spared the effect of some form of disruption (TMS, masking) requires comparison to some form of non-emotional control task and/or images. In this thesis I used several techniques to do this. In Experiments 1, 2 and 4, I used emotion as a redundant cue to the decision to be made (a newly learned arbitrary discrimination in Experiment 1 or a gender discrimination in Experiments 2 and 4). This was successful in Experiment 1, where performance for the emotional and neutral body posture images was almost identical for the control condition (stimuli presented in the opposite hemifield to the V1 disruption). Emotional body posture images in Experiment 1 were also relatively robust to the disruption effect of TMS to V1, suggesting that the emotions themselves were processed meaningfully and used for the discrimination.

Redundant cues in Experiments 2 and 4 were not as successful as in Experiment 1. When a discrimination of gender was made on difficult to see faces, there was no evidence of any benefit for the emotional faces. Null results are difficult to interpret, but one possibility is that an already learned 100% predictive cue – facial features associated with gender – blocked the learning of the second cue that was only predictive 50% of the time. Other possibilities worth consideration include the possibility that the emotions in the faces were not sufficiently strong, although the same stimuli did lead to emotional priming effects in Experiments 7 and 8. And these same stimuli have produced substantial effects of emotional valence in other studies (e.g. Hadjikhani, Hoge, Snyder & de Gelder, 2008; Perez-Edgar, Bar-Haim, McDermott et al., 2010; Mitchell, Nakic, Fridberg et al., 2007).

Another technique I used was to compare emotion discrimination with discrimination of non-emotional properties of the same stimuli. In Experiment 3, the task of orientation discrimination was used. Although this is clearly easier than the emotion discrimination task at supra-liminal SOAs, it was harder at the shortest SOA, and this provided one positive piece of evidence for emotion processing being spared the effect of masking. However, because the orientation task was not matched in difficulty to the emotion discrimination task, this effect may have been weaker and less significant than it could be. The other non-emotion task used, as requiring similar configural processing to emotion – age discrimination – was entirely unsuccessful as it was far too difficult, even with online calibration of difficulty.

Matching the emotion and non-emotion discrimination tasks for difficulty as attempted in Experiment 3, seems almost impossible, at least with currently available artificial face generation or real face sets. The relative success of matching task difficulty in Experiment 3 may be as close as good as that type of design can practically get. If 2AFC based on low visibility stimuli is required, then matching the difficulty of tasks may not be the best approach. Having said that, the finding of an advantage for the emotion discrimination over the orientation discrimination at low levels of visibility in Experiment 3 was really quite convincing given how much easier the orientation discrimination was at longer SOAs.

The final technique used in this thesis was a priming paradigm, where the measure is the effect of the prime emotion on a subsequent emotional classification of another stimulus. In many respects this is a neat way to investigate low visibility processing of emotions– it certainly gets round issues of task matching. All that is needed is a basic measure at the end of the experiment[6] to confirm the visibility of the stimuli as being low enough to convince the sceptic that the primes were subliminal. However, it may in practice be very difficult to locate the masking conditions within the narrow window between suppressing awareness, and suppressing processing altogether. This design was used in Experiments 5-8 with mixed results, and it certainly did not lead to consistent findings. Using priming as the measure for processing is also indirect – it does not directly measure the participants ability to discriminate properties of the

---

[6] Measuring visibility separate to the priming is not ideal, as one could argue fatigue or practice effects can lead to a different visibility measure compared to that experienced during the priming part itself. However, if participants are not to be told about the primes, it is somewhat unavoidable. The added measure of asking participants if they noticed the primes gives some added rigor.

prime, and there are many issues within the priming literature (e.g. possibility of S-R learning, reversal of response priming effects with delay) that may complicate interpretation, and several possibilities as to what is being primed (e.g. perceptual features, emotional semantics, responses). Moreover the emotional content of the prime may have multiple effects (category congruence priming, automatic alerting), which make it hard to know exactly what pattern of effect to look for.

Given the difficulties inherent in any one method, this thesis has used the converging operations approach of using two ways of disrupting visual processing, and several ways of assessing the relative immunity of processing of emotion and other attributes. Most techniques used gave some evidence that processing the emotional properties of low visibility social stimuli was relatively spared, but all three also gave evidence of failed replications or inconsistencies. The issue of matching tasks and difficulty, or how to get around these issues with techniques like priming, remains one for further consideration.

7.3.2. Failures to replicate

As noted previously, one problem found consistently in this thesis was the failures to replicate findings with different designs. Experiment 1 showed evidence of privileged emotion processing, but a different design for Experiment 2 failed to show equivalent findings. Experiment 3 found evidence using visual masking of privileged emotion processing, and the Experiment 4 failed to replicate this. Experiment 5 found evidence of congruence priming effects, but then Experiment 6 did not. Experiments 5, 7 and 8 all showed effects of priming, but these effects differed considerably across experiments. This consistent difficulty to find similar results with different designs is a problem, and it certainly questions the strength of any conclusion that emotion processing is relatively robust to disruption from TMS to V1 or visual masking.

As discussed above, null results are difficult to interpret but it is possible the failures in Experiments 2 and 4, that both used redundant cues of emotion during a gender discrimination task, were at least partly responsible. In retrospect, using this design in the two experiments without checking that the redundant emotion cue would be learned and used in this case was a mistake, although the experiments did tell us something. Experiment 2 appeared to show no significant effect of the cTBS on the discriminations at all (regardless of the processing of emotions), suggesting further work is needed to understand effects of cTBS on V1 before using it for present purposes. Experiment 4 then clarified that the redundant cue design, if coupled with a familiar rather than a newly learned discrimination (as in Experiment 1) was also flawed. To apply the blocking argument to the experiments in a broader sense, it is possible that the familiar gender discrimination used in Experiments 2 and 4 blocked the learning of the redundant cue of emotion. In Experiment 1 the discrimination was not gender, but a new and apparently arbitrary discrimination that was learnt at the start of the experiment.

The failure to find priming in Experiment 6 could be explained as the experiment used different stimuli to the other experiments and also showed really low visibility – perhaps too low for any meaningful processing to occur. The priming experiments present a second replication failure in that the priming effects differed between Experiments 5, 7 and 8. Initially, a congruence effect was predicted for Experiment 7 similar to that found in Experiment 5. However, post-hoc it is plausible to interpret the findings as arousal effects. As reviewed in Chapter 1, there is evidence that low visibility emotional stimuli can capture attention and effect responses to subsequent, non-emotive, stimuli (e.g. Sutton & Altarriba, 2011). In Experiment 5, the target stimuli were very similar to the prime stimuli – they were both sets of emotional faces. In Experiments 7 and 8 the target stimuli were words, and so different to the prime stimuli of faces. This distinction between the two experiment designs could have led to some of the differences in the findings.

The failures to replicate or extend findings in this thesis, then, can be explained. However, future research is needed to clarify whether the processing of emotions is relatively robust to disruption with TMS to V1 and visual masking. I discuss potential research directions to achieve this below.

7.4. Privileged emotion processing?

The experiments reported in this thesis did find some evidence of privileged processing of emotions. This section considers possible alternative account of these effects (assuming them to be genuine), and the possible anatomical basis of privileged processing of emotions.

7.4.1. Alternative accounts

The previous chapters mentioned several possible accounts for the findings that do not involve emotions, per se. There is the possibility that the effects were not due to the *emotional* content of the stimuli, but to their meaningfulness. This argument has two variants: firstly that there was an alternative, meaning cue present and the emotions were not involved at all, or secondly that the emotions led to the effect because they are meaningful, and other kinds of meaningfulness could have similar effects. The first argument carries little weight as it is difficult to imagine, across the different sets of stimuli, what cue to "meaning" could exist that happens to highly correlate with the emotional categories distinguishing the responses. The second is certainly a possibility, although it has to be noted that gender is a meaningful attribute of considerable adaptive significance, and no advantage appeared to be present for the gender discrimination task. The argument itself is also perhaps to some extent redundant – emotions are very meaningful and important social cues. It is difficult to imagine another cue that could carry as significant a meaning in order to directly test this effect.

Another kind of explanation would be that emotional expressions may be processed configurally (or holistically) and it was this type of processing that gave the discrimination of emotion its advantage, rather than affective content per se. This argument is implausible because of several findings in the reported experiments. The task used to test awareness in Experiments 3, 5, 7 and 8 was the orientation task (discriminate whether a face was upright or inverted). This task can, presumably, be done with configural processing if this is advantaged at low levels of visibility. There were several findings that indicated that emotional effects could be found with inverted faces (Experiment 3, 4, and 5) and inversion presumably disrupts configural processing. The inverted faces did show some differences to the upright faces, i.e. larger priming effects in Experiment 5, but the presence of effects at all certainly indicates featural low visibility emotion processing is possible, and making it unlikely that configural processing, alone at least, led to the emotional advantages found.

A final possibility is that there could be some templates or specific properties of the emotional faces that affected the discriminations somehow. For example, it could be the shape of the whites shown by bared teeth in a smile leads to different effects than the whites from an angry snarl. I.e., it was not the emotions themselves that led to effects, but some sort of template of basic perceptual features. This is essentially an argument against meaningful, semantic processing of the low visibility stimuli. This argument is unlikely to be true, at least for the priming experiments, as in Experiments 7 and 8 the emotional prime faces affected response times when the targets and primes were perceptually very different (faces and words). Also, in Experiment 3, the same stimuli, showing the same 'templates,' were only advantaged for the emotion discrimination, and not for the other measures (orientation).

In summary it seems probable that it was the emotional properties of the faces and body posture images, and most likely featural processing of those emotions, that led to the advantages found.

7.4.2. Anatomical basis

The processing of emotions appeared to be relatively immune to TMS to V1. This, along with the research reviewed in Chapter 1, provides support for at least basic processing of emotional content (e.g. threat versus trust) being possible in spite of processing being disrupted in striate cortex. This independence could be due to emotion discrimination being based on features detected in V1 that are robust to the disruption, or to the contribution of an alternative route to emotion processing bypassing V1.
One possibility is that the disruption methods used mostly targeted recurrent feedback processing in V1, leaving feedforward processing intact, and that feedforward processing is sufficient for emotion processing, but not for processing of other features. Alternatively, emotion-conveying features may simply be more robust to disruption of feedforward processing methods used. But it is unclear why either preserved feedforward processing, or compromised feedforward processing would be sufficient to allow discrimination of emotional stimuli, not other properties such as identity and orientation, especially as the cues to identity and emotion in Experiment 1 were so similar.

Research reviewed in Chapter 1 has found that when emotional stimuli are presented under low visibility/subliminal conditions, there was a shift in processing to sub-cortical structures in the brain (e.g. the amygdala – Morris, de Gelder, Weiskrantz et al, 2001; Williams, Das, Liddle, et al., 2005; Dannlowski, Ohrmann, Bauer et al., 2007; Williams, Moss & Bradshaw, 2004). With this literature in mind, it seems at least possible that a subcortical route to emotion processing (as proposed by LeDoux, 1986), operating largely outside of awareness, could have led to the advantages for the processing of emotions at minimal visibility levels as found in this thesis.

Of course there remains debate in the literature to what extent the whole emotion processing system (rather than just its inputs) can be characterised in terms of a two-route dichotomy. Indeed, there is a considerable amount of evidence that many cortical areas associated with emotional and face processing are involved in low visibility/subliminal processing of emotional faces (Pessoa & Adolphs, 2011). These areas (e.g. temporal-occipital cortex) are also associated with visible emotion processing. My findings are certainly compatible with one dynamic, highly interconnected system, possibly with two input routes – one via striate cortex, and one via subcortical structures. Such a dynamic system would also give some rationale as to why the processing of emotions was so difficult to disentangle from the processing of other facial and body expressions – if they are all processed by the same system, dissociations could be harder to make.

With the detailed critique Pessoa & Adolphs recently made (2010), and notable failures in the literature to find subliminal emotion effects (e.g. Andrews et al., 2011), the dual route model as outlined by LeDoux (1986) could certainly seen to be in some trouble. More research will be needed to tease apart the various factors and components involved in low visibility and subliminal emotional processing, however this thesis did find some support for a route independent of V1 processing emotional stimuli. TMS of cortical areas, real time measures of brain activity with EEG, and fMRI on cortical and subcortical areas, may be well placed for answering these questions. However it may also be that more sophisticated methods, including combinations of the above and more sensitive behavioural/physiological measures, are needed to sort the noise from the emotion.

7.5. Measures of awareness: choice of measurement in this thesis

In Experiment 1 (and 2), there was no systematic measurement of awareness. The point of the experiments (successful in Experiment 1) was to dissociate the effects of TMS on the discrimination of emotional and non-emotional attributes. Nevertheless, participants consistently reported that they found the body posture images very difficult to see, and the pattern of results was not correlated with overall discrimination performance.

In the remaining experiments, an explicit discrimination was used as an objective measure of awareness. For Experiments 3, 5, 7, and 8 the measure of awareness was discrimination between upright and inverted faces. This task was used as it was expected to be very easy – if awareness of the face were present the task should be simple to complete. For Experiment 3 this worked very well, as the task was overall easier than an emotional discrimination by a considerable way, showing that it was indeed a relatively easy task. Despite this, performance for the orientation task reached a point (at a 10ms SOA) at which discriminations were not reliably different from chance. For Experiments 5, 7 and 8 the measure also appeared to work well – performance at this easy task was essentially at chance in Experiment 5, and in Experiment 8, whilst being fairly low in Experiment 7. Whilst an easier task could have been used – i.e. presence/absence of a stimulus, this might have been too easy a task. After all, the point of the visibility/awareness measure is to test whether participants can "see" enough of the face to discriminate its properties, and specifically the visibility of the information relating to emotional expressions. Choosing a task that is comparable in difficulty, or easier, to discriminating and seeing emotions is sufficient to assess explicit discrimination of the stimuli. As it was, at longer SOAs in Experiment 5 the orientation task was much easier than the measures of emotional discriminations providing particularly strong evidence that this measure of visibility was actually fairly strict.

The issues surrounding whether objective or subjective measures should be used were discussed in Chapter 1. Objective measures were used in most experiments as they (in theory at least) have fewer issues with how participants interpret them. A subjective rating measure was used, in addition to an objective measure, in Experiment 7, though it did not really add much, except to confirm that subjective visibility increased, as did explicit discrimination, with SOA. More important, Experiments 5, 6, and 8 all included a less formal, but nevertheless powerful assessment of subjective awareness: whether the prime faces were noticed during the priming phase of the experiment. If participants did not report noticing the presence of any prime stimulus, it seems highly unlikely that they had any phenomenal awareness of the stimuli.

Overall, it seems fair to conclude that in Experiments 3 and 4 there were points (at short SOAs) where awareness was truly absent. For Experiments 5, 6 and 8 the combination of participants not noticing the prime faces, and performing poorly the visibility test at around chance level, were good evidence that the prime faces were indeed subliminal. For Experiment 7 the visibility was a bit higher, but the faces were still very difficult to see and visibility was fairly low. I have nevertheless tended to used the term “low visibility” rather than “subliminal” in recognition of the fact that this is a continuum for which the meaningful threshold is uncertain. For Experiments 1 and 2, I make no claim about awareness (except that the stimuli were hard to see).

7.6. Future research directions

7.6.1. Replications

Given the uncertainty of the findings from this thesis, the first direction for future research would be to replicate the findings. If the processing of emotions is privileged, then there should be evidence for this with a range of different designs. This section will outline some ideas as to how this could be achieved, building on my findings so far.

Clearly the TMS stimulation protocol in Experiment 2 was ineffective. There was no significant evidence of changes to the phosphene threshold following cTBS, and very little evidence of any disruptive effect of the cTBS. There are also inconsistent findings in the limited literature on disrupting processing in V1 with cTBS. Experiment 2 may thus have ‘jumped the gun’ it its use of developing TMS methods. Perhaps the next step would be to give a simpler visual task (e.g. a visual perimetry test), and initially aim to demonstrate impaired detection or discrimination in a selective region of the visual field before moving on to complex stimuli of faces, and testing the processing of emotional stimuli.

An alternative direction would be to apply the online single pulse design that was successful in Experiment 1 to faces, or, in order to increase the size of the TMS effect, switch to using a small train of pulses on each trial to maximise the disruption. A stronger disruption effect would be a considerable benefit as it may lead to larger differences between emotional and neutral discriminations and allow for stronger conclusions to be made.

The failures of Experiments 2 and 4 may have been, at least in part, due to the emotional cues to the required gender discrimination not being learned. One could perhaps repeat the design but with a discrimination less likely to be so overlearned that learning of the emotion cue is blocked. Gender is highly socially relevant discrimination we make everyday. In contrast, the body posture stimuli in Experiment 1, which did show some evidence of privileged emotion processing, were classified on an entirely arbitrary basis (apart from the emotional postures). The stimuli in Experiment 1 were also shown clearly to the participants before the start of the experiment, which have aided the formation of the association, whether consciously or not. The question here would be what task to use that fulfils this requirement. One option would be to generate faces that vary by basic physical features (e.g. small features vs. larger features). Another possibility would be to use a particularly difficult discrimination on a familiar dimension – e.g. morphing male and female, or famous and non-famous faces, on a dimension that could vary from relatively easy to very difficult and ambiguous. Under those conditions, redundant cues to emotion may be more likely to be harnessed. Both of these have the convenience of not requiring participants to learn an arbitrary category as was necessary in Experiment 1, which allows for larger sets of stimuli.

For the priming experiments, the distinction between congruence and arousal priming – i.e. when each occurs, needs to be clarified. This could be done through simply increasing the data and running more experiments similar to those of Experiment 5, 7 and 8 to ascertain when each occurs. Another option would be to use additional measures of arousal or congruence. For example, ERPs could be measured and used to analyses differences in activity for trials where there is a semantic mismatch, or a semantic match **(e.g.** Mecklinger, Kramer & Strayer, 1992**)**, to give an additional, objective, indication of congruence. A measure of participants' physiological responses (e.g. skin conductance, sympathetic facial muscle activity) would give some measure of emotional reaction and arousal. By taking these measures it may be possible to ascertain how the priming effects come about under different designs, and further understand the differences found in Experiments 5-8 of this thesis.

7.6.2. Extensions

Although the findings of this thesis are not definitive, it is worth some brief discussion of where this research area could go if replications, like those suggested above, are successful and the evidence for privileged emotion processing increases. One direction would be to image where in the brain the processing of emotion in subliminal stimuli occurs relative to that observed with supraliminal stimuli. Is it limited to subcortical structures, or does cortex get activated via the "back door" in much the same regions as via an undisrupted visual cortex? This sort of question could be addressed using fMRI during a masking paradigm similar to those used in Chapter 5, or a priming paradigm similar to those used in Chapter 6. If offline stimulation (cTBS) can be used successful to disrupt processing in V1, this would also allow for the combination of TMS and fMRI to image where in the brain the relatively spared processing of emotions occurs. This could also be done using online TMS as techniques for the combination of TMS and fMRI progress, although it is technically more difficult.

Another possibility would be to use TMS to disrupt other areas of the visual processing network. If the low visibility emotion processing is (at least partly) based in the visual cortex, disrupting these areas should lead to a reduction in the emotion advantage. For example, if stimuli of faces are used in a masking or priming paradigm TMS could be applied to the occipital face area (OFA). If the OFA is involved in low visibility emotional face processing, the emotional advantage should reduce. If not, then the interpretation would be more difficult, but it would suggest that the OFA is not involved in the processing of emotional stimuli.

7.7. Conclusions

To conclude, the mixed findings of this thesis indicate that processing of the emotion in faces and body posture images may be relatively robust to disruptions of visual processing resulting in low levels of visibility and even absence of awareness. Moreover, the TMS evidence suggests this is should be attributed to a sub-cortical “hot line” to emotion processing centres (whether sub-cortical or cortical). However, the marginal reliability of some of the effects found, the small effect sizes, and the difficulty in replicating the effects with converging methods and across different stimulus sets suggests that any processing advantage for emotional content is relatively subtle. There is much work yet to do to clarify these issues

References

Abrams, R.L., & Grinspan, J. (2006). Unconscious semantic priming in the absence of partial awareness. *Consciousness and Cognition, 16,* 942–953.

Adolphs, R., Gosselin, F., Buchanan, T. W., Tranel, D., Schyns, P., & Damasio, A. R. (2005). A mechanism for impaired fear recognition after amygdala damage. *Nature, 433,* 68-72.

Adolphs, R., Russel, J. A., & Tranel, D. (1999). A role of the human amygdala in recognising emotional arousal from unpleasant stimuli. *Psychological science, 10(2),* 167-171.

Adolphs, R., & Tranel, D. (2003). Amygdala damage impairs emotion recognition from scenes only when they contain facial expressions. *Neuropsychologia, 41*, 1281-1289.

Adolphs, R., Tranel, D., Damasio, H., Damasio, A. (1994). Impaired recognition of emotion in facial expressions following bilateral damage to the human amygdala. *Nature, 372,* 669-672.

Aguado, L., Garcia-Gutierrez, A., Castaneda, E., & Saugar, C. (2007). Effects of prime task on affective priming by facial expressions of emotion. *The Spanish Journal of Psychology, 10 (2),* 209-217.

Algom, D., Chajut, E. & Lev, S. (2004). A rational look at the emotional Stroop phenomenon: A generic slowdown, not a Stroop effect. *Journal of Experimental Psychology: General, 133*, 323-338.

Amaral, D. G., Behniea, H., Kelly, J. L. (2003). Topographic organization of projections from the amygdala to the visual cortex in the macaque monkey. *Neuroscience, 118(4),* 1099-1120.

Amaral, D. G., Price, J. L., Pitkanen, A., Carmichael, S. T. (1992). *Anatomical organization of the primate amygdaloid complex*. In: The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction (Aggleton J, ed), pp 1-66. New York: Wiley-Liss.

Amassian, V.E., Cracco, R.Q., Maccabee, P. J., Cracco, J. B., Rudell, A., & Eberle, L. (1989). Suppression of visual perception by magnetic coil stimulation of human occipital cortex. *Electroencephalography and Clinical Neurophysiology, 74*, 458-462.

Anders, S., Eippert, F., Weiskopf, N., & Veit, R. (2008). The human amygdala is sensitive to the valence of pictures and sounds irrespective of arousal: an fMRI study. *Social cognitive and affective neuroscience, 3(3),* 233-243.

Andrews, V., Lipp, O. V., Mallen, K. M., & Konig, S. (2011). No evidence for subliminal affective priming with emotional facial expression primes. *Motivation and Emotion, 35,* 33-43.

Ashley, V., Vuilleumier, P., & Swick, D. (2003). Time course and specificity of event-related potentials to emotional expressions. *Neuroreport, 15,* 211–216.

Balota, D.A., Yap, M.J., Cortese, M.J., Hutchison, K.A., Kessler, B., Loftis, B., Neely, J.H., Nelson, D.L., Simpson, G.B., & Treiman, R. (2007). The English Lexicon Project. Behavior Research Methods, 39, 445-459.

Banse, R. (1999). Automatic evaluation of self and significant others: affective priming in close relationships. *Journal of social and personal relationships, 16,* 803-821.

Banse, R. (2001). Affective priming with liked and disliked persons: prime visibility determines congruency and incongruency effects. *Cogntiion and emotion, 15 (4),* 501-520.

Barbur, J.L., Watson, J. D., Frackowiak, R. S., & Zeki, S. (1993). Conscious visual perception without V1. *Brain, 116,* 1293–302.

Bard, P. (1928). A diencephalic mechanism for the expression of rage with special reference to the sympathetic nervous system. *American journal of physiology, 84,* 490–515.

Baron-Cohen, S., Ring, H. A., Bullmore, E. T., Wheelwright, S., Ashwin, C., & Williams, S. C. R. (2000). The amygdala theory of autism. *Neuroscience and behavioural reviews, 24,* 355-364.

Bateson, M., Nettle, D., & Roberts, G. (2006). Cues of being watched enhance cooperation in real-world setting. *Biology letters, 2,* 412-414.

Bernston, G. G., Bechara, A., Damasio, H., Tranel, D., & Cacioppo, J. T. (2007). Amygdala contribution to selective dimensions of emotion. *Social cognitive and affective neuroscience, 2(2),* 123-129.

Block, N. (1990). Consciousness and accessibility. *Behavioural and brain sciences, 13 (4),* 596-598.

Boy, F., & Sumner, P. (2009). Tight coupling between positive and reversed priming in the masked prime paradigm. *Journal of Experimental Psychology: Human Perception and Performance, 36,* 892-905.

Bradley, M.M., & Lang, P.J. (1999). *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings.* Technical report C-1, Gainesville, FL. The Center for Research in Psychophysiology, University of Florida.

Breitmeyer, B. (2007). Visual masking: past accomplishments, present status, future developments, *Advances in Cognitive Psychology, 3,* 9-20.

Breitmeyer, B. G., Kafaligönül, H., Öğmen, H., Mardon, L., Todd, S., & Ziegler, R. (2006). Meta- and paracontrast reveal differences between contour- and brightness-producing mechanisms. *Vision research, 46 (17),* 2645-2658.

Breitmeyer, B., & Ogman, H. (2006). *Visual Masking: Time slices through conscious and unconscious vision.* OUP, Oxford.

Corthout, E., Uttl, B., Ziemann, U., Cowey, A., & Hallett, M. (1999). Two periods of processing in the (circum)striate visual cortex as revealed by transcranial magnetic stimulation. *Neuropsychologia, 37,* 137-145.

Corthout, E., Uttl, B., Walsh, V., Hallett, M., & Cowey, A. (1999). Timing of activity in early visual cortex as revealed by transcranial magnetic stimulation. *Neuroreport, 10,* 2631-2634.

Corthout, E., Uttl, B., Juan, C. H., Hallett, M., & Cowey, A. (2000). Suppression of vision by Transcranial magnetic stimulation: a third mechanism. *Neuroreport, 11*, 2345-2349.

Costafreda, S. G., Brammer, M. J., David, A. S., & Fu, C. H. Y. (2008). Predictions of amygdala activation during the processing of emotional stimuli: a meta analysis of 385 PET and fMRI studies. *Brain research review, 58,* 57-70.

Croner, L. J., & Kaplan, E. (1995). Receptive fields of P and M ganglion cells across the primate retina. *Vision research, 35(1),* 7-24.

Dalgleish, T. (2004). The emotional brain. *Nature perspectives, 5,* 582-589.

Dalgleish, T., Dunn, B. D., Mobbs, D. (2009). Affective neuroscience: past, present and future. *Emotion review, 2 (4),* 355-368.

Darwin, C. (1872). *The Expression of Emotion in Man and Animals. University of Chicago Press*. Chicago.

de Gelder, B., & Hadjikhani, N. (2006). Non-conscious recognition of emotional body language. *Neuroreport, 17,* 583-586.

de Gelder, B., van Honk, & Tamietto. (2011). Emotion in the brain: of low roads, high roads and roads less travelled. *Nature Reviews Neuroscience, 12,* 425-426.

de Gelder, B., Vroomen, J., Pourtois, G., & Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of the striate cortex. *Neuroreport, 10,* 3759-3763.

De Houwer, J., Hermans, D. (2010). Do feelings have a mind of their own? In J. De Houwer & D. Hermans (Eds.), *Cognition & Emotion: Reviews of current research and theories* (pp. 38-65). Hove, UK: Psychology Press.

Dannlowski, U., Ohrmann, P., Bauer, J., Kugel, H., Arolt, V., Heindel, W., & Suslow, T. (2007). Amygdala reactivity predicts automatic negative evaluations for facial emotions. *Psychiarty research: neuroimaging, 154,* 13-20.

Day-Brown, J. D., Wei, H., Chomsung, R. D., Petry, H. M., & Bickford, M. E. (2010). Pulvinar projections to the striatum and amygdala in the tree shrew. *Frontiers in neuroanatamy, 4,* 143.

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Science, 10(5),* 204-11.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, 79,* 1-37.

Dimberg, U., Thunberg, M., & Elmehed, K. (2000). Unconscious facial reactions to emotional facial expression. *Psychological science, 11(1),* 86-89.

Duncan, S., Barrett, L. F. (2007). Affect is a form of cognition: a neurobiological analysis. *Cogn Emot., 21 (6),* 1184-1211.

Ekman, P. (1972). Universals and cultural differences in facial expressions of emotion. In J. Cole (Ed.), *Nebraska Symposium on Motivation 1971*, 19, 207-283.

Ekman, P. (1999). Basic Emotions. In Dalgleish, T., Power, M. *Handbook of Cognition and Emotion*, Sussex, UK: John Wiley & Sons.

Ekman, P., Friesen, W. V. (1971). Constants across cultures in the face and emotions. *Journal of personality and social psychology, 17,* 124-129.

Eckstein, D., & Henson, R. (2012). Stimulus/Response learning in masked congruency priming of faces: Evidence for covert mental classifications? *Quarterly Journal of Experimental Psychology, 65,* 92-120.

Eckstein, D., & Perrig, W. J. (2007). The influence of intention on masked priming: a study with semantic classification of words. *Cognition, 104,* 345-376.

Eimer, M., & Holmes, A. (2002). An ERP study on the time course of emotional face processing. *NeuroReport, 13,* 427-431.

FaceGen software: http://www.facegen.com/index.html 15th April, 2011

Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2007). Masking disrupts re-entrant processing in human visual cortex. *Journal of cognitive neuroscience, 19,* 1488-1497.

Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). The spatiotemporal profile of cortical processing leading up to visual perception. *Journal of Vision, 8*, 1-12.

Felleman, D. J., & van Essen, D. C. (1991). Distributed hierarchical processing in primate visual cortex, *Cerebral Cortex, 1,* 1-47.

Fernandez, E., Alfro, A., Tormos, J. M., Climent, R., Martinez, M., Vilanova, H., Walsh, V., & Pascual-Leone, A. (2002). Mapping of the human visual cortex using image-guided transcranial magnetic stimulation. *Brian Research Protocols, 10,* 115-124.

Finkbeiner, M., & Friedman, L. (2011). The flexibility of nonconsciously deployed cognitive processes: evidence from masked congruency priming. *PlosOne, 6 (2),* 1-13.

Fisk, G. D., & Haase, S. J. (2005). Unconscious perception or not? An evaluation of detection and discrimination as indicators of awareness. *American journal of psychology, 118 (2),* 183-212.

Franca, M., Koch, G., Mochizuki, H., Huang, Y., & Rothwell, J. C. (2006). Effects of theta burst stimulation protocols on phopshene thresholds. *Clinical Neurophysiology, 117,* 1808-1813.

Freese, J. L., Amaral, D. G. (2005). The organisation of projections from the amygdala to visual cortex areas TE and V1 in Macaque Monkey. *The journal of comparative neurology, 486,* 295-317.

Frings, C., & Wentura, D. (2003). Who is watching big brother? TV consumption predicted by masked affective priming. *European journal of social psychology, 33,* 779-791.

Gaillard, R., Cul, A. D., Naccache, L., Vinckier, F., Cohen, L., & Dehaene, S. (2006). Nonconscious semantic processing of emotional words modulates conscious access. *PNAS, 103 (19),* 7524-7529.

Gaillard, R., Dehaene, S., Adam, A., Clemenceau, S., Hasboun, D., Baulac, M., Cohen, L., & Naccache. (2009). Converging intracranial markers of conscious access. *PLoS Biology, 7 (3),* 472-492.

Gallagher, M., Holland, P. C. (1994). The amygdala complex: multiple roles in associative learning and attention. *PNAS, 91,* 11771-11776.

Garavan, H., Pendergrass, J. C., Ross, T. J., Stein, E. A., & Risinger, R. C. (2001). Amygdala response to both positively and negatively valenced stimuli. *Neuroreport, 12 (12),* 1-5.

Glascher, J., & Adolphs, R. (2003). Processing of the arousal of subliminal and supraliminal emotional stimuli by the human amygdala. *The Journal of Neuroscience, 23(22),* 10274-10282.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action, *Trends in neurosciences, 15(1),* 20-25.

Gupta, R., & Raymond, J. E. (2012). Emotional distraction unbalances visual processing. *Psychological Bulletin Review, 19(2),* 184-189.

Hamm, A. O., Weike, A. I., Schupp, H. T., Treig, T., Dressel, A., & Kessler, C. (2003). Affective blindsight: intact fear conditioning to a visual cue in a cortically blind patient. *Brain, 126,* 267-275.

Hinojosa, J. A., Carretie, L., Mendez-Bertolo, C., Miguez, A., & Pozo, M. A. (2009). Arousal contributions to affective priming: electrophysiological correlates. *Emotion, 9(2),* 164-171.

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences, 9,* 1–23.

Houwer, J. D., Hermans, D., Rothermund, K., & Wentura, D. (2002). Affective priming of semantic categorisation responses. *Cognition and emotion, 16 (5),* 643-666.

Huang, Y-M., Baddeley, A., Young, A. W. (2008). Attentional capture by emotional stimuli is modulated by semantic processing. *Journal of experimental psychology: human, perception and performance, 34 (2),* 328-339.

Huang, Y., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta burst stimulation of the human motor cortex. *Neuron, 45,* 201-206.

Humphrey, N. (1974). Vision in a monkey without striate cortex: a case study. *Perception, 3,* 241-255.

James, T. W., Culham, J., Humphrey, G. K., Milner, A. D., & Goodale, M. A. (2003). Ventral occipital lesions impair object recognition but not object-directed grasping: A fMRI study. *Brain, 126,* 2463–2475.

Jiang, Y., Costello, Fang, Huang & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *PNAS, 103(45),* 17048-17052.

Jiang, Y., & He, S. (2006). Cortical responses to invisible faces: dissociating subsystems for facial-information processing. *Current biology, 16,* 2023-2029.

Jolij, J., & Lamme, V. A. F. (2005). Repression of unconscious information by conscious processing: Evidence from affective blindsight induced by transcranial magnetic stimulation. *Proceedings of the National Academy of Sciences, 102,* 10747-10751.

Juan, C. H., & Walsh, V. (2003). Feedback to V1: a reverse hierarchy in vision. *Experimental brain research, 150 (2),* 259-263.

Kamin, L.J. (1969). Predictability, surprise, attention and conditioning. In Campbell, B. A., & Church, R. M. (eds.), *Punishment and aversive behavior*, 279–96, New York: Appleton-Century-Crofts.

Kammer, T. (2007). Visual masking by transcranial magnetic stimulation in the first 80 milliseconds. *Advances in Cognitive Psychology, 3,* 177-179.

Kiesel, A., Kunde, W., & Hoffman, J. (2007). Unconscious priming acording to multiple S-R rules. *Cognition, 104,* 89-105.

Kiesel, A., Kunde, W., Pohl, C., & Hoffmann, J. (2006). Priming from novel masked stimuli depends on target set size. *Advances in cognitive psychology, 2 (1),* 37-45.

Kiesel, A., Wagener, A., Kunde, W., Hoffmann, J., Fallgatter, A. J., & Stocker, C. (2006). Unconscious manipulation of free choice in humans. *Consciousness and cognition, 15,* 397-408.

Kiesel, A., Kunde, W., Pohl, C., Berner, M. P., & Hoffmann, J. (2009). Playing chess unconsciously. *Journal of Experimental Psychology: Learning, Memory and Cognition, 35 (1),* 292-298.

Killcross. S., Robbins, T. W., Everitt, B. J. (1997). Different types of fear-conditioned behaviour mediated by separate nuclei within amygdala. *Nature, 388 (6640),* 377-380.

Kober, H., Barrett, L. F., Joseph, J., Bliss-Moreau, E., Lindquist, K., & Wager, T. D. (2008). Functional grouping and cortical–subcortical interactions in emotion: A meta-analysis of neuroimaging studies. *Neuroimage, 42(2),* 998-1031.

Kolb, F. C., & Braun, J. (1995). Blindsight in normal observers. *Nature, 377(6547),* 336-338.

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of The Royal Society Biological Sciences, 362,* 857-875.

Kouider, S., & Dupoux, E. (2004). Partial awareness creates the "illusion" of subliminal semantic priming. *Psychological science, 15 (2),* 75-81.

Kouider, S., & Dupoux, E. (2007). How "semantic" is response priming restricted to practiced items? A reply to Abrams & Grinspan (2007). *Consciousness and cognition, 16,* 954-956.

Kousta, S-T., Vinson, D. P., & Vigliocco, G. (2009). Emotion words, regardless of polarity, have a processing advantage over neutral words. *Cognition IN PRESS.*

Kawasaki, H., Adolphs, R., Kaufman, O., Damasio, H., Damasio, A., Granner, M., Bakken, H., Hori, T., & Howard III, M. A. (2001). Single-neuron responses to emotional visual stimuli recorded in human ventral prefrontal cortex. *Nature neuroscience, 4(1),* 15-16.

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

Krolak-Salmon, P., Henaff, M. A., Vighetto, A., Bertrand, O., & Mauguiere, F. (2004). Early amygdala reaction to fear spreading in occipital, temporal, and frontal cortex: a depth electrode ERP study in human. *Neuron, 42(4),* 665-676.

Lachter, J., Forster, K. I., & Ruthruff, E. (2004). Forty-Five Years After Broadbent (1958): Still no Identification Without Attention. *Psychological Review, 3(4),* 880-913.

Lamme, V. A. F. (2001). Blindsight: the role of feedforward and feedback corticocortical connections. *Acta Psychologica, 107,* 209-228.

Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in cognitive sciences, 10 (11),* 494-501.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Cognitive Neuroscience, 23,* 571-579.

Lamme, V. A. F., Super, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current opinion in neurobiology, 8,* 529-535.

Lang, P. J., Bradley, M. M., Fitzsimmons, J. R., Cuthbert, B. N., Scott, J. D., Moulder, B., & Nangia, V. (1998). Emotional arousal and activation of the visual cortex: an fMRI analysis. *Psychophysiology 35,* 199 – 210.

Lazarus, R., Averill, J., and Opton, E. (1970). Towards a cognitive theory of emotion. In M. Arnold (Ed.), *Feelings and emotions*. New York: Academic Press.

LeDoux, J. E. (1996). *The Emotional Brain.* New York: Simon & Schuster.

Liddle, B. J., Brown, K. J., Kemp, A. H., Barton, M. J., Das, P., Peduto, A., Gordon, E., & Williams, L. M. (2005). A direct brainstem-amygdala-cortical 'alarm' system for subliminal signals of fear. *Neuroimage, 24,* 235-243.

Lovheim, H. (2011). A new three-dimensional model for emotions and monoamine neurotransmitters. *Med Hypotheses, 78,* 341-348.

Luo, Q., Holroyd, T., Jones, M., Hednler, T., & Blair, J. (2007). Neural dynamics for facial threat processing as revealed by gamma band synchronization using MEG. *Neuroimage, 34,* 839-847.

Marcel, A. (1983). Conscious and unconscious perception: experiments on visual masking and word recognition. *Cognitive psychology, 15,* 197-237.

MacLean, P. D. (1949). Psychosomatic disease and the "visceral brain." Recent developments bearing on the Papez theory of emotion. *Psychosomatic Medicine, 11,* 338-353.

Macmillan, N. A., & Kaplan, H. L. (1985). Detection theory analysis of group data: estimating sensitivity from average hit and false alarm rates. *Psychological bulletin, 98,* 185-199.

McGaugh, J. L. (2004). The amygdala modulates the consolidation of memories of emotionally arousing experiences. *Annual review in neuroscience, 27 (1),* 1-28.

Mizuno, N., Uchida, K., Nomura, S., Nakamura, Y., Sugimoto, T., Uemura-Sumi, M. (1981). Extrageniculate projections to the visual cortex in the macaque monkey: an HRP study. *Brain Res 212,* 454–459.

Mogg, K., Bradley, B. P., de Bono, J., & Painter, M. (1997). Time course of attentional bias for threat information in non-clinical anxiety. *Behaviour Research and Therapy, 35,* 297–303.

Morland, A. B., Lê, S., Carroll, E., Hoffmann, M. B., & Pambakian, A. (2004). The role of spared calcarine cortex and lateral occipital cortex in the responses of human hemianopes to visual motion. *Journal of Cognitive Neuroscience, 16,* 204-18.

Morris, J. S., Öhman, A., & Dolan, R. J. (1999). A subcortical pathway to the right amygdala mediating "unseen" fear. *Proceedings of the National Academy of Sciences, 96,* 1680-1685.

Morris, J. S., de Gelder, B., Weiskrantz, L., & Dolan, R. J. (2001). Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain, 124,* 1241–1252.

Moutoussis, K., & Zeki, S. (2002). The relationship between cortical activation and perception investigated with invisible stimuli. *PNAS, 99 (14),* 9527-9532.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: affective priming with optimal and suboptimal stimulus exposures. *Journal of personality and social psychology, 64 (5),* 723-739.

Nagarajan, S. S., Durand, D. M., & Warman, E. N. (1993). Effects of induced electric fields on finite neuronal structures: a simulation study. *IEEE Trans Biomed Eng., 40 (11),* 1175-1188.

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

Öhman, A., Carlsson, K., Lundqvist, D., **&** Ingvar, M. (2007). On the unconscious subcortical origin of human fear. *Physiology & Behavior, 92,* 180-185.

Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of personality and social psychology, 80,* 381–396.

Paus, T., Jech, R., Thompson, C.J., Comeau, R., Peters, T., and Evans, A.C. (1997). Transcranial magnetic stimulation during positron emission tomography: a new method for studying connectivity of the human cerebral cortex. *J Neurosci. 17(9),* 3178-3184.

Pessoa, L., & Adolphs, R. (2011). Emotion processing and the amygdala: from a ‘low road’ to ‘many roads’ of evaluating biological significance. *Nature Reviews Neuroscience, 11,* 773-782.

Pessoa, L., McKenna, M., Gutierrez, E., Ungerleider, L. G. (2002). Neural processing of emotional faces requires attention. *PNAS, 99 (17),* 11458-11463.

Phan, K.L., Wager, T.D., Taylor, S.F., & Liberzon, I. (2002). Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI. *Neuroimage, 16,* 331–348.

Phelps, E. A., LeDoux, J. (2005). Contributions of the amygdala to emotional processing: from animal models to human behaviour. *Neuron, 48,* 175-187.

Phelps, E. A., O'Connor, K. J., Cunningham, W. A., Funayama, E. S., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2000). Performance on indirect measures of race evaluation predicts amygdala activation. *Journal of cognitive neuroscience, 12 (5),* 729-738.

Pissiota, A., Frans, O., Michelgard, A., Appel, L., Langstrom, B., Flaten, M. A., & Fredrikson, M. (2003). Amygdala and anterior cingulate cortex activation during affective startle modulation: a PET study of fear. *European journal of neuroscience, 18,* 1325-1331.

Pitcher, D., Charles, L., Devlin, J., Walsh, V., & Duchaine, B. (2009). Triple dissociation between faces, bodies and objects in extrastriate cortex. *Current Biology, 19,* 319-324.

Pitcher, D., Garrido, L., Walsh, V., & Duchain, B. C. (2008). Transcranial magnetic stimulation disrupts the percetion and embodiment of facial expressions. *The journal of neuroscience, 28(36),* 8929-8933.

Pitcher, D., Walsh, V., & Duchaine, B. (2011). The role of the occipital face area in the cortical face perception network. *Experimental Brain Research, 209,* 481-493.

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, *114*, 510-532.

Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. Journal of Experimental Psychology. *Human Perception and Performance,36 (3),* 596-618.

Robinson, J. L., Laird, A. R., Glahn, D. C., Lovallo, W. R., & Fox, P. T. (2010). Metaanalytic connectivity modelling: delineating the functional connectivity of the human amygdala. *Human brain mapping, 31 (2),* 173-184.

Rossi, S., Hallett, M., Rossini, P. M., & Pascual-Leone, A. (2009). Safety, ethical considerations, and application guidelines for the use of transcranial magnetic stimulation in clinical practice and research. *Clinical neurophysiology, 120,* 2008-2039.

Rozenkrants, B., Olofsson, J. K., & Polich, J. (2008). Affective visual event-related potentials: arousal, valence, and repetition effects. *International Journal of Psychophysiology, 67,* 114-123.

Sabatini, E., Penna, S. D., Franciotti, R., Ferretti, A., Zoccolotti, P., Rossini, P. M., Romani, G. L., & Gainotti, G. (2009). Brain structures activated by overt and covert emotional visual stimuli. *Brain research bulletin, 79,* 258-264.

Sahraie, A., Weiskrantz, L., Barber, J. L., Simmons, A., Williams, S. C. R., & Brammer, M. J. (1997). Pattern of neuronal activity associated with conscious and unconscious processing of visual signals. *PNAS, 94,* 9406-9411.

Salin, P, A., & Bullier, J. (1995). Corticocortical connections in the visual system : Structure and function. *Physiol. Reviews, 75*, 107-154.

Schlaghecken, F., & Eimer, M. (2004). Masked prime stimuli can bias "free" choices between response alternatives. *Psychonomic bulletin & review, 11 (3),* 463-468.

Schmidt, T. (2007). Measuring unconscious cognition: Beyond the zero-awareness criterion. *Advances in Cognitive Psychology, 3,* 275-287.

Schmidt, T., & Vorberg, D. (2006). Criteria for unconscious cognition: three types of dissociation. *Perception & psychophysics, 68 (3),* 489-504.

Schneider, G. E. (1969). Two visual systems. Science, 163 (3870), 895-902.

Schoenfeld, M. A., Noesselt, T., Poggel, D., Tempelmann, C., Hopf, J-M., Woldorff, M. G., Heinze, H-J., & Hillyard, S. A. (2002). Analysis of pathways mediating preserved vision after striate cortex lesions. *Ann Neurol, 52,* 814-824.

Sherman, S. M., & Guillery, R. W. (1996). The functional organization of thalamocortical relays. *J. Neurophysiol. 76,* 1367–1395.

Shipp, S. (2004). The brain circuitry of attention, *Trends in Cognitive Sciences, 8 (5),* 223–230.

Silvert, L., Lepsien, J., Fragopanagos, N., Goolsby, B., Kiss, M., Taylor, J. G., Raymond, J. E., Shapiro, K. L., Eimer, M., & Nobre, A. C. (2007). Influence of attentional demands on the processing of emotional facial expressions in the amygdala. *Neuroimage, 2(1),* 357-366.

Sincich, L. C., Park, K. F., Wohlgemuth, M. J., & Horton, J. C. (2005). Bypassing V1: a direct geniculate input to area MT. *Nature Neuroscience, 7 (10),* 1123-1128.

Smith, F. W., Schyns, P. G. (2009). Smile through your fear and sadness: transmitting and identifying facial expression signals over a range of viewing distances. *Psychological science, 20(10),* 1202-1208.

Stepniewska, I. (2004). The pulvinar complex. In: *The primate visual system*. CRC Press LLC. (Eds.), Kaas J. H., & Collins, C.E.

Stewart, P. A., & Schubert, J. N. (2006). Taking the "low road" with subliminal advertisements: a study testing the effect of precognitive prime "RATS" in a 2000 presidential advertisement. *The Harvard International Journal of Press/Politics, 11 (4),* 103-114.

Stone, A., & Valentine, T., & Davis, R. (2001). Face recognition and emotional valence: processing without awareness by neurologically intact participants does not simulate covert recognition in prosopagnosia. *Cognitive, affective and behavioural neurosience, 1 (2),* 183-191.

Strahan, E. J., Spencer, S. J., & Zanna, M. P. (2002). Subliminal priming and persuasion: striking while the iron is hot. *Journal of Experimental Social Psychology, 38,* 556-568.

Sumner, P. (2008) Mask-induced priming and the negative compatibility effect. *Experimental Psychology, 55*, 133-141.

Sumner, P., Tsai, P., Yu, K., & Nachev, P. (2006). Attentiona modulation of sensorimotor processes in the absence of perceptual awareness. *PNAS, 103(27),* 10520-10525.

Suslow, T., Ohrmann, P., Bauer, J., Rauch, A. V., Schwindt, W., Arolt, V., Heindel, W., & Kugel, H. (2006). Amygdala activity during masked presentation of emotional faces predicts conscious detecion of threat-related faces. *Brain and Cognition, 61,* 243-248.

Sutton, T. M., & Altarriba, J. (2011). The automatic activation and perception of emotion in word processing: evidence from a modified dot probe paradigm. *Journal of cognitive psychology, 23(6),* 736-747.

Tamietto, M., & de Gelder, B. (2007). Affective blindsight in the intact brain: neural interhemispheric summation for unseen fearful expressions. *Neuropsychologia,*

Tamietto, M., & de Gelder, B. (2010). Neural bases of the non-conscious perception of emotional signals. *Nature Neuroscience Reviews, 11,* 697-709.

Tapia, E., & Breitmeyer, B. G. (2011). Visual consciousness revisited: magnocellular and parvocellular contributions to conscious and nonconscious vision. *Psychological Science, 22(7),* 934-942.

Tapia, E., Breitmeyer, B. G., & Shooner, C. R. (2010). Role of task-directed attention in nonconscious and conscious response priming by form and color. *Journal of experimental psychology: human perception and performance, 36(1),* 74-87.

Tsuchiya, N., Moradi, F., Felsen, C., Yamazaki, M., Adolphs, R. (2009). Intact rapid detection of fearful faces in the absence of the amygdala. *Nature neuroscience, 12 (10),* 1224-1225.

Ungerleider, L. G., & Mishkin, M. (1982). Object vision and spatial vision: two cortical pathways. In Ingle, D. J., Goodale, M. A., Mansfield, R. J. W. *Analysis of visual behaviour,* Boston: MIT press.

Van den Bussche, E., Van den Noortgate, W., & Reynvoet, B. (2009). Mechanisms of masked priming: a meta-analysis. *Psychological bulletin, 135(3),* 452-477.

VanRullen, R., & Koch, C. (2003). Visual selective behaviour can be triggered by a feed-forward process. Journal of cognitive neurosience, 15 (2), 209-217.

VanRullen, R., & Thorpe, S. L. (2001). The time course of visual processing: from early perception to decision-making. Journal of cognitive neuroscience, 13 (4), 454-461.

Vorberg, D., Mattler, U., Heinecke, A., Schmidt. T., & Schwarzbach, J. (2003). Different time courses for visual perception and action priming. *PNAS, 100* (10), 6275-6280.

Vuilleumier, P., Armony, J. L., Clarke, K., Husain, M., Driver, J., & Dolan, R. J. (2002). Neural response to emotional faces with and without awareness: event-related fMRI in a parietal patient with visual extinction and spatial neglect. *Neuropsychologia, 40,* 2156-2166.

Vuilleumier, P., Armony, J. L., Driver, J., Dolan, R. J. (2001). Effects of attention and emotion on face processing in the human brain: an event-related fMRI study. *Neuron, 30,* 829-841.

Vuilleumier, P., Armony, J. L., Driver, J., Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature neuroscience, 6,* 624-631.

Vuilleumier, P., Richardson, M. P., Armony, J. L., Driver, J., & Dolan, R. J. (2004). Distant influences of amygdala lesion on visual cortical activation during emotional face processing. *Nature neuroscience, 7 (11),* 1271-1278.

Walsh, V. & Pascual-Leone, A. (2003). *Transcranial magnetic stimulation: A neurochronometics of mind.* MIT Press: Cambridge, MA.

Waterston, M. L., & Pack, C. C. (2010). Improved discrimination of visual stimuli following repetitive transcranial magnetic stimulation. *PLOS one, 5(4),* e10354.

Weinberger, J., & Western, D. (2008). RATS, we should have used Clinton: subliminal priming in political campaigns. *Political psychology, 29 (5),* 631-651.

Weiskrantz, L., Warrington, E. K., Sanders, M. D., & Marshall, J. (1974). Visual capacity in the hemianopic field following a restricted occipital ablation. *Brain, 97,* 709-728.

Wheeldon, L. R., & Monsell, S. (1994). Inhibition of spoken word production by priming a semantic competitor. *Journal of memory and language, 33,* 332–356.

Wiens, S. (2006). Current concerns in visual masking. *Emotion, 6,* 675-680.

Williams, M. A., McGlone, F., Abbott, D. F., & Mattingley, J. B. (2005). Differential amygdala responses to happy and fearful facial expressions depend on selective attention. *Neuroimage, 24,* 417-425.

Williams, M. A., Moss, S. A., & Bradshaw, J. L. (2004). A unique look at face processing: the impact of masked faces on the processing of facial features. *Cognition, 91,* 155-172.

Williams, L. M., Das, P., Liddell, B. J., Kemp, A. H., Rennie, C. J., & Gordon, E. (2006). Mode of functional connectivity in amygdala pathways dissociates level of awareness for signals of fear. *The Journal of Neuroscience, 26 (36),* 9264-9271.

Winkielman, Berridge & Wilbarger (2005). Unconscious Affective Reactions to Masked Happy Versus Angry Faces Influence Consumption Behavior and Judgments of Value. *Personality and social psychology bulletin, 31 (1),* 121-135.

Yang, T. T., Menon, V., Eliez, S., Blasey, C., White, C. D., Reid, A. J., Gotlib, I. H., & Reiss, A. L. (2002). Amygdala activation associated with positive and negative facial expressions. *Neuroreport, 13 (14),* 1737-1741.

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

Yovel, & Kanwisher, N. (2005). The Neural Basis of the Behavioral Face-Inversion Effect. *Current Biology, 15,* 2256-62.

Zajonc, R. B. (1980). Feeling and thinking. Preferences need no inferences. *American Psychologist, 35,* 151-175.

APPENDIX 1

Items analysis for Experiment 2

Average performance for the 96 stimuli used in Experiment 2 are shown in Table A1.1 below. Overall, stimuli were discriminated at 73% accuracy (SD = 12%). Of the ninety-six images, a total of six were 'outliers' in that performance was less than 2SD below the mean, a very strict criteria. An analysis on the mean performance for each condition of the experiment was re-run with the outlier stimuli removed. This analysis, and the performance without the removals, is shown in Tables A1.2 and A1.3 respectively. As can be seen, the removals had no effect of note on the pattern of results.

	Male			Female		
	Neutral	Happy	Angry	Neutral	Happy	Angry
1	80	70	85	86	85	80
2	89	80	82	65	66	**36***
3	76	74	83	86	83	80
4	78	68	85	80	71	87
5	79	64	76	69	72	57
6	79	68	88	83	81	86
7	74	68	85	83	81	71
8	77	52	81	78	82	78
9	84	80	80	85	92	91
10	91	91	93	48	65	**29***
11	77	76	77	48	64	**26***
12	66	57	81	77	84	**36***
13	73	56	73	**40***	60	**29***
14	59	46	81	65	68	61
15	76	77	71	82	84	72
16	60	51	76	88	78	88

Table A1.1: Mean accuracy (% correct) for each of the ninety-six stimuli used in Experiment 2. * Stimuli were outliers in that performance was less than 2SD below the overall mean.

		Time		
		1	2	3
Neutral	Phosphene Location	73	75	76
	Control Location	72	75	75
Emotional	Phosphene Location	70	73	73
	Control Location	69	72	74

Table A1.2: Group mean categorisation performance (% correct) for each condition of Experiment 2 with no item exclusions.

		Time		
		1	2	3
Neutral	Phosphene Location	73	74	74
	Control Location	71	75	75
Emotional	Phosphene Location	70	73	73
	Control Location	70	73	73

Table A1.3: Group mean categorisation performance (% correct) for each condition of Experiment 2 excluding 6 ‘outlier” items.

APPENDIX 2

Items analysis for Experiments 6-8

Average accuracy and RTs for the prime and target stimuli used in Experiments 6-8 are shown in Table A2.1 below. Overall, none of the prime stimuli from any of the experiments give any cause for concern – they all fall comfortably within 3SD of the mean RT and accuracy. In Experiment 6, all the target stimuli were also within 3SD of the mean for the accuracy and RT measures. Both Experiments 7 and 8 had three stimuli outside 3SD from the mean – each with two stimuli for accuracy and one for RT. These stimuli were removed from the data and the analyses re-run. Table A2.2 shows the new means for Experiment 7, and Table A2.3 the new means for Experiment 8. Removing the small number of stimuli from the analysis changed the results very little. Overall, there is no evidence that any of the individual stimuli used Experiments 6-8 caused any excessive noise in responses.

		Experiment					
		6		7		8	
		Target	Prime	Target	Prime	Target	Prime
RT	Mean	577	574	481	486	558	560
	SD	23	16	40	13	40	9
	Minimum	536	545	252	462	336	544
	Maximum	634	605	598	514	683	578
Error	Mean	5.2	5.4	4.3	3.9	4.8	4.9
	SD	2.6	1.7	4.3	1.6	4.6	1.2
	Minimum	1.1	1.6	0.0	0.7	0.7	2.8
	Maximum	13.0	9.8	9.8	7.0	38.2	7.8

Table A2.1: Mean, standard deviation, minimum and maximum RT (ms) and error rate (% incorrect) for the stimuli used in Experiments 6-8.

	Target valence	Prime emotion	SOA				
			10	20	30	40	50
Full dataset	Negative	Angry	640	624	629	624	636
		Happy	590	615	620	628	632
		Neutral	589	616	630	645	621
	Positive	Angry	581	601	616	614	601
		Happy	586	608	604	598	607
		Neutral	579	607	609	621	623
After exclusions	Negative	Angry	638	624	629	623	635
		Happy	589	614	620	627	631
		Neutral	590	615	628	644	622
	Positive	Angry	579	601	615	612	600
		Happy	586	607	603	597	606
		Neutral	576	607	605	618	623

Table A2.2: Summary of the mean RTs (ms) with and without the exclusion of outlier stimuli for Experiment 7.

	Target valence	Prime emotion	SOA		
			10	20	30
Full dataset	Negative	Angry	580	566	557
		Happy	578	559	560
		Neutral	564	575	559
	Positive	Angry	558	569	551
		Happy	567	565	561
		Neutral	561	570	561
After exclusions	Negative	Angry	580	566	557
		Happy	578	559	560
		Neutral	564	575	559
	Positive	Angry	554	567	552
		Happy	563	563	559
		Neutral	559	567	558

Table A2.3: Summary of the mean RTs (ms) with and without the exclusion of outlier stimuli for Experiment 8.

APPENDIX 3

Median analysis Experiments 5-8

Means of median RTs are below for Experiments 5 – 8, in Tables A3.1 – A3.4 respectfully. The median analysis essentially shows the same pattern as reported in the main text of the thesis for each experiment, indicating the effects (and lack of effects) reported in Chapter 6 are unlikely to have resulted from outliers in RTs.

		Target	
		Angry	Happy
Upright	Congruent	518	510
	Incongruent	524	509
Inverted	Congruent	527	513
	Incongruent	529	521

Table A3.1: Mean of median correct RTs (ms) for each prime orientation, prime congruency, and target emotion from Experiment 5.

		Target	
		Angry	Happy
Upright	Congruent	521	529
	Incongruent	530	530
Inverted	Congruent	524	521
	Incongruent	523	528

Table A3.2: Mean of median correct RTs (ms) for each prime orientation, prime congruency, and target emotion from Experiment 6.

Target valence	Prime emotion	SOA 10	20	30	40	50
Positive	Angry	551	571	584	586	569
	Happy	553	581	579	570	581
	Neutral	558	584	577	586	582
Negative	Angry	602	608	603	592	605
	Happy	563	590	598	592	598
	Neutral	555	596	602	606	595

Table A3.3: Median RT (ms) for each Target valence, prime emotion, and prime – target SOA from Experiment 7.

Target valence	Prime emotion	SOA 10	20	30
Positive	Angry	530	534	527
	Happy	540	542	539
	Neutral	542	540	531
Negative	Angry	545	536	534
	Happy	548	541	538
	Neutral	545	542	533

Table A3.4: Median RT (ms) for each target valence, prime emotion, and prime – target SOA from Experiment 8. The data in this table have been collapsed across location.