The Role of Attention in Face Processing

Markus Bindemann

Department of Psychology

University of Glasgow

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

August 2004

Abstract

Selective attention is widely regarded as a crucial component of human perception. In the visual domain, attentional mechanisms have been implicated in stimulus encoding, implicit recognition, conscious perception and goal-directed behaviour. To date, however, the role of attention in face processing has been largely overlooked. This is remarkable given the social and biological importance of faces, and the wealth of psychological research that has focused on faces as stimuli. Moreover, if we are to better understand how the human brain processes faces, then this would also require an insight into the interaction between attention and face processing. The experiments in this thesis addressed the relation of attention and face processing directly by assessing the consequences of various attentional manipulations in response-competition and repetition priming tasks. The first line of enquiry examined observers' ability to attend selectively to facial expression and identity, and whether attention is required for the integration of these types of information into a multi-dimensional face percept. Subsequent experiments examined capacity limits in face processing and attention biases to faces and nonface comparisons. The main findings indicate that face processing is capacity limited, such that only a single face can be processed at a time, and that faces are particularly efficient at retaining and engaging visual attention in comparison to nonface objects. However, while face processing limits appear to proceed independent of a general capacity, attention biases to faces may reflect processing stages that are shared with other stimuli. These findings are discussed in relation to existing research on faces and attention.

Acknowledgements

I would like to thank Professor Mike Burton for his support and superb supervision throughout this project. I would also like to offer my thanks to Rob Jenkins for generously sharing his expertise on faces and attention. This research was supported by an ESRC Postgraduate Training Award, R42200134060.

To my parents, the *minderbroeders* and Jane.

Declaration

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

Markus Bindemann

Publications

Chapters 3 and 5 of this thesis have been submitted for publication.

Chapter 3

Bindemann, M., Burton, A.M., & Jenkins, R. (in press). Capacity limits for face processing. *Cognition.*

Chapter 5

Bindemann, M., Burton, A.M., Hooge, I.T.C., Jenkins, R., & De Haan, E.H.F. (submitted). Faces retain attention. *Psychonomic Bulletin & Review.*

Table of contents

Chapter 1 General Introduction

1.1 Introduction

There are a range of selective issues implicated in human perception that come under the umbrella term *attention* (Driver, 2001; Styles, 1997). Thus it is widely held that attentional mechanisms afford the selection of information from our senses for conscious perception and goal-directed behaviour (e.g. Broadbent, 1958), but also for implicit recognition (e.g. Lavie, 1995) and for the rudimentary encoding of stimuli (e.g. Treisman & Gelade, 1980). Although a great deal is now known about selective attention (e.g. Baddeley & Weiskrantz, 1993; Pashler, 1998; Styles, 1997), rather little is understood about how visual attention and face processing interact. This is remarkable as there is probably no other class of visual stimuli that can match the social and biological importance, and that has been studied as extensively as the human face (e.g. Bruce & Young, 1998; Young, 1998). This is also unfortunate as the role of attention may be imperative in understanding how the human brain processes faces. Similarly, the study of visual attention might benefit from considering stimuli of immense intrinsic significance such as faces.

This thesis explores the interaction of attention and face processing across four themes by measuring task-irrelevant processing in response-competition and priming tasks. The first theme concerns the ability to attend selectively to different types of facial information. The second theme explores the role of attention in integrating these types of facial information into multidimensional percepts during early visual processing. The third theme concerns the limit of the number of faces that can be processed simultaneously. The final theme focuses on late visual processes involved in awareness and response, by examining whether faces are particularly efficient in affecting responses to a target in comparison with other stimulus classes. I begin by outlining the relevant principles of attention. This is followed by a review of what is currently understood about the relation of attention and face processing. I end this chapter by describing the general methodological approach of the current work.

1.2 Principles of attention

1.2.1 Selectivity and capacity limitation

Two principles that have dominated the study of visual attention are *selectivity* and *capacity limitation.* Selectivity is apparent in most human behaviour and refers to the observation that we continually assign priority to only a part of our entire sensory input. For example, when reading we focus on a small set of words at a time even though a page usually contains other words that we could also read. Selectivity is closely related to the principle of capacity limitation, which refers to the notion that attention is a finite resource that can only be devoted to a subset of the total sensory input. Consequently, selectivity and capacity limitation are often portrayed as two sides of the same coin, whereby selectivity prioritizes limited attentional resources to enhance the processing of important stimuli.

A long-standing issue within the attention domain has been the locus of selectivity. Early selection theorists suggest that some basic physical attributes of all sensory inputs are analyzed independent of their attentional status at an early processing stage, but only attended stimuli are processed to identification and beyond (e.g. Broadbent, 1958). Late selection theorists, on the other hand, propose that selective processing only begins after the analysis of all stimuli is completed. Thus, the identity of attended and unattended stimuli is computed alike in a capacity-unlimited analysis and selection only occurs after full perception to gain access to systems required for awareness, memory, and response (e.g. Deutsch & Deutsch, 1963; Duncan, 1980; Norman, 1968). The issue of early versus late selection is therefore a debate about the extent to which unattended stimuli are processed.

In the visual domain, attention researchers have employed many variations of a rather modest number of tasks to investigate the locus of selection, such as *response-competition* and *distractor-priming* tasks. In response-competition, subjects are typically instructed to make two-alternative speeded responses to a target item while ignoring distractor items in the display. Distractor processing can then be assessed via their congruency on target response times (RTs). To the extent that the distractors are processed, RTs are slowed when the distractors and the target belong to opposite response categories (the incongruent condition) relative to when distractors belong to the same response category as the target (the congruent condition). Using this type of paradigm, Eriksen & Eriksen (1974) asked subjects to classify a central letter while ignoring flanking letter distractors. Each response category consisted of two letters (e.g. A and U to be responded to with one hand, and M & K with the other hand), and the distractors could belong to the same or the opposite response category (e.g. AAUAA or MMUMM). They found that the distractor letters influenced target responses, with slower classification times in the incongruent condition. This is consistent with the central claim of late selection that attended and unattended stimuli are processed to identification (see also e.g. Logan, 1980; Eriksen & Schultz, 1979; Miller, 1987).

However, Eriksen & Hoffman (1972, 1973) also obtained results with a response-competition task that are difficult to integrate into strict late selection accounts. They presented circular displays subtending 2° of visual angle in diameter and containing 12 letters, one of which was marked as a target by a short arrow-cue. As in Eriksen & Eriksen's (1974) study, target RTs were slowed on incongruent relative to congruent trials, but these congruency effects were eliminated when the distractors appeared more than 1° from the target. In a similar way, Yantis & Johnston (1990) found that letter distractors only produced interference when they were presented next to the target or separated by one response-neutral item, but not for more remote distances.

In another twist of the selection debate, Tipper and associates first produced another measure for the processing of unattended visual stimuli and showed later that distractors can still be processed in the absence of response-competition with a target. Tipper (1985) presented subjects with superimposed, different-coloured line drawings of two objects and asked them to name the object in a specified colour, but to ignore the other. Importantly, the relationship between ignored and attended objects was manipulated, so that the ignored object on one trial was occasionally presented as the target object on a subsequent trial. In these instances, Tipper (1985) obtained slower target naming times (a phenomenon that Tipper termed 'negative priming'), even though subjects were incapable of reporting the identity of the unattended stimuli. Moreover, this effect occurred between letters of the same identity but a different shape (Tipper & Cranston, 1985) and pictures and names of objects sharing a semantic category (e.g. CAT-DOG, Tipper, 1985; Tipper & Driver, 1988), showing that the unattended stimuli were subject to considerable processing.

Subsequently, Driver & Tipper (1989) examined whether distractors give rise to negative priming even when they do not interfere with target classification. Subjects were asked to count the number of red items in a display, while ignoring black distractor digits that could be congruent or incongruent with the correct counting response. The black digits did not appear to interfere with task-relevant processing. However, when this was followed by an interference display in which the red target items were congruent with the numerical value of the preceding distractors, negative priming was found. It appears then that even non-interfering distractors can be processed to identification.

Nevertheless, some researchers urge caution in interpreting negative priming as an unequivocal measure of late selection (Lavie & Tsal, 1994). One possibility, for example, is that the distractor may only be subject to 'raw' processing on trial *n*, but may then be fully primed by the related target on trial *n +1*. It is also uncertain what negative priming theorists would predict in situations in which only a proportion of distractor stimuli produce interference (e.g. Eriksen & Hoffman, 1972, 1973; Yantis & Johnston, 1990). Indeed, Driver (2001) argues that a position that is often overlooked in this debate is that many effects do not fit into strict early or late selection accounts, but may be driven by only partial processing along a continuum between early and late selection.

1.2.2 PERCEPTUAL LOAD

Lavie (1995) offered an explanation of how early and late selection theories might combine along a continuum. According to her perceptual load theory of selective attention, the most important determinant whether stimuli within the visual field are processed is the *perceptual load* of attended and unattended items. A major assumption of this theory is that processing resources cannot be voluntarily withheld, but that visual analysis of relevant and irrelevant stimuli proceeds automatically until available capacity is exhausted. Consequently, to-be-ignored distractor stimuli are only excluded from analysis when the perceptual load of the relevant task requires all available capacity. If relevant stimuli do not exhaust this capacity, excess processing resources automatically spill over to irrelevant stimuli, thus enabling their processing.

In a review of past findings, Lavie & Tsal (1994) integrate a wealth of existing data in support of a perceptual load account, and since then considerable evidence has also been accumulated (Lavie, 1995, 2000, 2001; Lavie & Cox, 1997; Lavie & Fox, 2000; Rees, Frith & Lavie, 1997). For example, Lavie (1995) showed subjects displays consisting of one of two possible target letters and an additional task-irrelevant distractor letter, which could be congruent (i.e. the same letter as the displayed target) or incongruent (the same letter as the alternative target) with the target response. Perceptual load was manipulated by presenting only the target and distractor (the low load condition), or by embedding the target in a horizontal string of response neutral letters (the high load condition). Lavie (1995) predicted that distractor interference would depend on task-relevant load, so that it would be reduced in the high-load compared to the low load condition. This was confirmed by the results, which showed that distractor congruency effects were only found under low load.

1.2.3 Feature integration

Although Lavie's perceptual load theory suggests a continuum between early and late selection, it also implies that perception occurs in an all-or-nothing manner such that task-irrelevant distractor stimuli either are or are not processed. In contrast, Treisman's feature integration theory (FIT) suggests that task-irrelevant stimuli always undergo some processing even if they are not identified by the perceptual system (e.g. Treisman, 1988, 1993; Treisman & Gelade, 1980; Treisman & Schmidt, 1982). According to FIT, different stimulus attributes, such as colour, size, or orientation are registered separately by the perceptual system without any attentional effort, but are integrated into multidimensional percepts by visual attention. Attention is thus portrayed as a type of glue that binds different features together.

One source of support for FIT comes from visual search tasks, in which subjects are required to detect a target in an array of distractor items. Target search times are seemingly unaffected by increasing the number of distractors in a display, provided that all distractors are identical and differ from the target in terms of a single dimension such as shape or colour (e.g. Smith, 1962; Treisman & Gelade, 1982). This pattern of results is often referred to as perceptual 'pop-out' and is interpreted as capacity-free parallel search. However, search times generally increase with display size when targets and distractors consist of feature conjunctions, such as colour *and* shape (e.g. Heathcote & Mewhort, 1993; Treisman & Gelade, 1980), a serial search pattern that implies that target and distractor processing requires focused attention. Thus it seems that feature integration, unlike feature perception, is attentionally demanding.

This interpretation has not gone unchallenged as visual search for conjunction targets seems relatively easy when the distractors can be grouped into a common shape (e.g. Duncan & Humphreys, 1989; Humphreys, Quinlan & Riddoch, 1989; Wolfe, 1994). One can also question the extent to which feature search proceeds without attention as subjects are always looking deliberately at search displays to locate the target. However, FIT also receives support from reports that conjunction information appears sometimes unavailable in unattended stimuli. For example, Lavie (1997) showed that response-competition effects with a colour-shape conjunction target (e.g. a purple cross) are equivalent for conjunctive distractors, in which critical colour and shape information is combined within one of two distractors (e.g. a purple cross and a blue triangle, where 'blue' and 'triangle' are response-neutral features), and disjunctive distractors (e.g. a purple triangle and a blue cross). In accordance with FIT, this suggests that correct conjunction information was unavailable under conditions of inattention. In addition, FIT receives some support from neuropsychological patients who can describe the colours and shapes of objects quite accurately, but have difficulty in reporting their correct conjunctions (Friedman-Hill, Robertson & Treisman, 1995; Humphreys, Cinel, Wolfe, Olson & Klempen, 2000).

1.2.4 Attention biases

The feature integration theory and the perceptual load theory meet on the principle that the allocation of attention is vital for full perception. A crucial question then concerns the extent to which this allocation can be controlled by an observer. What we see often depends on where we choose to attend in our environment, such as a task-relevant target in an experiment. Selection of the *focus of attention* can thus undoubtedly occur in a voluntary, goal-directed manner, even if spare attentional capacity may inevitably spill over to other task-irrelevant stimuli (e.g. Lavie, 1995, 2000).

Focused attention may also be driven in an involuntary manner so that stimuli receive priority even when they oppose an observer's intentions. For example, Posner, Snyder & Davidson (1980) showed that responses to a peripheral target could be cued by an immediately preceding illumination in one of the possible target locations, resulting in faster responses to validly compared to invalidly cued targets. Intriguingly, these effects were observed even when the cues were only valid on a minor proportion of all trials, so that it would have been advantageous to ignore them as most of the time they would have been misleading. This suggests that the cues captured attention in an involuntary bottom-up manner, independent of the subjects' intentions.

In addition to abrupt visual onsets (e.g. Posner et al, 1980; Jonides, 1981; Remington, Johnston & Yantis, 1992), attentional capture is also invoked by the salience of a stimulus. Thus, visual search for a uniquely shaped target is slowed if one of the distractors is printed in a salient colour, for example a red item in an otherwise green array. Similarly, search for a colour target is disrupted by the presence of a differently shaped distractor (e.g. Theeuwes, 1991, 1992, 1994). Consequently, bottom-up capture is held to operate on early visual processing, most probably at the level of feature perception in Treisman's feature integration theory (Styles, 1997). However, although impossible to eliminate, capture effects are reduced when observers are aware of the exact target shape and distractor colour prior to the experiment (e.g. Theeuwes, de Vries & Godijn, 2003). This indicates that there are also top-down processes that determine the extent to which attention may be retained by a distractor stimulus.

Comparable attentional effects have also been observed with more meaningful stimuli than with colours and simple shapes. Wolford & Morrison (1980) showed that responses regarding the parity of two peripheral digits (e.g. both odd or even versus one of each) were slowed more when a subject's own name was presented as a distractor than by control words. Shapiro, Caldwell & Sorensen (1997) also found that own names are detected more often in a rapid stream of visual stimuli than words that are of no particular significance to an observer. In addition, Mack & Rock (1998) report that visual search times for own names in 1, 6, or 12 word displays do not increase with display size. In contrast, search slopes for control names increased at an average rate of 51 ms/item. Moreover, these search slopes rose to 81 ms/item when subjects' own names rather than words were used as distractors. Thus, these studies show that task performance is facilitated when own names are used as visual targets and impaired when they are presented as distractors, indicating that attention is drawn to this stimulus class.

Note that some reports also challenge this interpretation. For example, Bundesen, Kyllingsbaek, Houmann & Jensen (1997) found that own names are no more potent as distractors during a matching task than other names, but are reported more accurately, which implies that names might not capture attention but are simply more recognizable. This receives support from claims that own names are reported more quickly in visual search than other targets but not in the capacity-free fashion reported by Mack & Rock (1998), and are not particular potent distractors (Harris, Pashler & Coburn, 2004; Harris & Pashler, 2004). Nonetheless, attentional biases have also been observed with negatively charged emotional words, particularly in anxiety-prone individuals (e.g. Broadbent & Broadbent, 1988; MacLeod, Mathews & Tata, 1986; Williams, Mathews & MacLeod, 1996). For example, MacLeod et al (1986) found that anxious individuals detect a dot probe faster when its location is validly cued by an emotionally threatening word (e.g. PANIC) than a non-emotional word (e.g. FLUTE). Since these studies employed a variety of emotionally charged and neutral words, it is improbable that these biases can be explained in terms of recognizability. Indeed, similar biases have also been observed for substance-related cues such as cigarettes and bottles of alcohol in smokers and heavy drinkers (e.g. Jones, Jones, Smith & Copley, 2003; Waters, Shiffman, Bradley & Mogg, 2003), which are items that are generally effortlessly recognizable. This suggests that attentional capture does not depend on recognizability but on the meaning that particular stimuli hold for an individual.

A number of studies suggest that some meaningful stimuli are also effective at retaining attention, particularly in anxious individuals. For example, Amir, Elias, Klump & Przeworski (2003) used threatening or neutral words to cue the locations in which a target could appear. Subjects were generally slower to respond to invalidly cued than validly cued targets. However, subjects with social phobia showed significantly slower response latencies on invalid trials than psychologically normal participants, but only when threatening cues were used. Others have shown similar biases for threatening words in state-anxious individuals (e.g. Fox, Russo, Bowles & Dutton, 2001) and for threatening pictures in trait-anxious subjects (Yiend & Mathews, 2001), which implies that these groups have difficulty in disengaging attention from threatening material.

Attentional capture and disengagement effects by meaningful stimuli are evidently dependent on stimulus identification and must therefore build on earlier selection processes such as feature integration and perceptual load. Moreover, these attention biases are clearly different from distractor interference in response-competition tasks (e.g. Eriksen & Eriksen, 1974; Eriksen & Hoffman, 1972, 1973), as they do not depend on target-distractor congruency. This suggests that they are not located at the level at which task-relevant responses are activated, but reflect a stimulus's ability to control action regardless of task demands. The next question is how mechanisms such as feature integration, capacity limitation and attention biases relate to face processing. Moving from early to late visual processing, I begin with the role of attention in holistic face encoding.

1.3 The role of attention in holistic face encoding

1.3.1 THE IMPORTANCE OF HOLISTIC INFORMATION

Although faces are often described in terms of the features for which we have particular lexical values (e.g. eye, nose, mouth), considerable evidence suggests that faces are processed in a *holistic* fashion that is dependent on the spatial *configuration* of these features, rather than the features themselves. Thus, Harmon (1973) showed that configural information is sufficient for person recognition when high-frequency information from individual features is disrupted with a blurring technique. Others have shown that face parts are recognized more accurately in the context of a studied face than when they are presented in isolation, a whole-to-part advantage that suggests that faces are not encoded as individual features (Tanaka & Farah, 1993). Face recognition is also impaired by manipulations of configural information, such as stimulus inversion (e.g. Carey &

Diamond, 1977; Diamond & Carey, 1986; Valentine & Bruce, 1986, Yin, 1969) and changes in the internal spacing of facial features (e.g. Haig, 1984; Tanaka & Sengco, 1997). In addition, Young, Hellawell & Hay (1987) found that subjects experienced great difficulty in naming either the top half of one face or the bottom half of another, when both halves were closely aligned to form a 'chimeric' face composite. This *composite effect* has been attributed to interference from the novel configuration of the aligned halves, which does not match the configural information of either of the original faces.

Several researchers have also made the stronger claim that face processing is more dependent on holistic processing than other stimulus categories (e.g. Carey & Diamond, 1977; Diamond & Carey, 1986; Scapinello & Yarmey, 1970; Tanaka & Farah, 1993; Valentine & Bruce, 1986; Yarmey, 1971; Yin, 1969, 1970). For example, Yin (1969, 1970) found that the recognition of airplanes, houses, stick figures, bridges and costumes is less affected by inversion than the recognition of faces. Furthermore, Tanaka & Farah (1993) and Tanaka & Sengco (1997) showed that inverted faces and houses are recognized equally well as whole items as from isolated parts, a result that contrasts the whole-to-part recognition advantage for faces (although a whole-to-part advantage has been reported for some artificially generated objects, e.g. Gauthier & Tarr, 1997; Gauthier, Williams, Tarr & Tanaka, 1998).

Some researchers also argue for the existence of face-specific neural mechanisms that are responsible for the holistic processing of upright faces. These are held to operate to some extent independently of a general-purpose object recognition system, which processes inverted faces and nonface objects. Although this issue remains controversial (see e.g. Kanwisher, 2000; Gauthier & Logothetis, 2000), evidence for such a division has accrued from brain-imaging studies of normal subjects (e.g. Haxby, Hoffman & Gobbini, 2000; Haxby et al, 1999; Kanwisher, McDermott & Chun, 1997), neuropsychological patients with selectively impaired face processing abilities (e.g. Farah, Levinson & Klein, 1995; Farah, Wilson, Drain & Tanaka, 1995; McNeill & Warrington, 1993; Sergent & Signoret, 1992a) or selectively impaired object processing abilities (e.g. Humphreys & Rumiati, 1998; Moscovitch, Winocur & Behrmann, 1997; Rumiati & Humphreys, 1997), and reports that newborn infants prefer intact face stimuli to scrambled faces and nonface objects (e.g. Johnson, Dziurawiec, Ellis & Morton, 1991; Morton & Johnson, 1991).

1.3.2 How does attention relate to holistic face processing?

The importance of holistic information for face processing has led researchers to investigate whether facial features are coded into holistic representations by focused visual attention, an approach that originates from Treisman's feature integration theory. In the first of these studies, Reinitz, Morrissey & Demb (1994) asked subjects to study line-drawn faces under full attention or in a divided-attention condition, in which they were required to count a rapid sequence of dots that were alternating between the top and bottom half of each face. Counting the dots affected the ability to remember the study faces at a subsequent test phase so that participants classified the original faces and conjunction faces, which were constructed by combining an eye-nose set from one studied face with a mouth-hair set from another studied face, as old equally often. Moreover, the faces in both these conditions were classified as old more frequently than completely new faces and conjunction stimuli made from an old and a new face. This led Reinitz et al

(1994) to conclude that faces require attention for the holistic encoding of their constituent features.

However, the results of a subsequent study are inconsistent with these conclusions. Reinitz, Bartlett & Searcy (1997) employed the same full and divided attention conditions as Reinitz et al (1994), but presented participants with a same/different test immediately after viewing each study face. At test participants could be shown a featurally altered face in which one feature (e.g. nose, mouth or eyes) was replaced by a different exemplar of the same type, or a configurally altered face in which eye-mouth distance was manipulated, or an identical face. In contrast to Reinitz et al (1994), this paradigm produced a greater deficit for featurally altered faces than configurally altered faces under divided attention, suggesting that holistic processing was now actually less attentionally demanding than feature perception.

While Reinitz et al (1997) argue that attention may affect face encoding differently in memory and perception, there might be other reasons for these contradictory results. First, Reinitz et al (1997) used images of real colour faces while the earlier study employed line-drawn stimuli. However, although the processing of real face stimuli is particularly dependent on configuration, this type of information appears less important for the processing of artificial face stimuli (Leder, 1996). Second, it has been questioned whether it is reasonably possible to manipulate featural and configural information independently (e.g. Bruce, 1988; Rhodes, Brake & Atkinson, 1993), as replacing one feature with another exemplar will inevitably produce a concurrent change to a face's configuration. This problem may be compounded by the use of different feature-configuration manipulations in these studies. Lastly, the dot-counting task of the divided attention conditions may have affected face encoding by forcing subjects to alternate between the top and bottom halves of the study faces. This could potentially disrupt holistic encoding or feature perception or even both, by directing attention away from relevant face information.

More recently, Boutet, Gentes-Hawn & Chaudhuri (2002) re-examined the role of attention in holistic face encoding with a variation of the composite effect (see Young et al, 1987). In one experiment, stimuli were composed of a face superimposed on a house and subjects were asked to focus on just one of these images during encoding to manipulate attention towards or away from the face. In another experiment with an analogous attention manipulation, a stream of letters was continuously scrolled across a face and subjects were asked to decipher any words within this stream or ignore the letters altogether. Face encoding during these tasks was then assessed with stimuli constructed from the top and bottom halves of either two attended faces, two unattended faces, or two new faces. These halves could be closely aligned to produce a composite face with a novel configuration or misaligned to disrupt configuration. If only the facial features were encoded under divided attention, then recognition of the face halves should have been unaffected by the configuration of the aligned composites. However, although Boutet et al (2002) found that attended face composites were recognized more accurately than unattended composites, misaligned stimuli produced better recognition performance than aligned composites regardless of condition. This indicates that dividing attention interfered with the degree to which the study faces could be remembered but did not disrupt holistic face encoding.

However, this evidence is not entirely persuasive as several studies demonstrate a disproportionate effect of inversion for faces in comparison with photographs of houses (Scapinello & Yarmey, 1970; Valentine & Bruce, 1986; Yarmey, 1971; Yin, 1969). Unlike faces, houses are also recognized equally well from individual parts as from whole items (Tanaka & Farah, 1993; Tanaka & Sengco, 1997). In other words, it appears that houses do not draw on the same (holistic) resources as faces. Task-irrelevant face processing also appears unaffected by task-relevant word processing (e.g. Jenkins, Burton & Ellis, 2002; Lavie, Ro & Russell, 2003; Young, Ellis, Flude, McWeeney & Hay, 1986). This opens the possibility that Boutet et al (2002) failed to manipulate relevant processing capacity away from the study faces as a consequence of the stimulus set used.

This notion receives support from another study that investigated holistic face encoding. Palermo & Rhodes (2002) asked participants to study a central face under full attention, or to match two upright or two inverted flanker faces while studying the central target in a divided attention condition. A speeded two-choice recognition test followed at the end of each trial, consisting of two intact faces, the target and a foil face that differed from the target by one feature (e.g. a pair of eyes), or two isolated exemplars of a particular feature, one of which was extracted from the target. The results showed a whole-to-part recognition advantage for face targets under full attention, indicating that the targets were encoded holistically. Intriguingly, matching upright flanker faces eliminated this advantage but matching inverted flankers did not, even though the latter was more time-consuming. Palermo & Rhodes (2002) concluded that attention is important for encoding facial configuration, but that there might be two distinct processing systems (for similar claims see e.g. Farah et al, 1995; Moscovitch et al, 1997) perhaps even with their own dedicated attentional capacities. The first reflects an object recognition system dedicated to processing individual facial features but also inverted faces and nonface objects. The other system is used for holistic face processing. Matching the inverted faces thus left sufficient capacity for the holistic encoding of the central face target because they did not consume any holistic resources.

1.3.3 OTHER TYPES OF FACIAL INFORMATION

One aspect that is consistently ignored by the studies reviewed in the previous section is that faces convey much more information than just identity, such as sex, emotional expression and facial speech. Models of face recognition specify functionally independent processing routes for these types of information (e.g. Bruce & Young, 1986). Thus it is not necessary, for example, to recognize someone as familiar to determine his or her sex or facial expression. In support of this architecture, there is now evidence for dissociations between sex and expression (Le Gal & Bruce, 2002), sex and identity (Bruce, Ellis, Gibling & Young, 1987; Ellis, Young & Flude, 1990), sex and facial speech (Green, Kuhl, Meltzoff & Stevens, 1991), facial speech and identity (Campbell, De Gelder & De Haan, 1996; Campbell, Landis & Regard, 1986), and identity and expression (e.g. Campbell, Brooks, De Haan & Roberts, 1996; Humphreys, Donnelly & Riddoch, 1993; Sergent, Ohta, MacDonald & Zuck, 1994). Moreover, Calder, Young, Keane & Dean (2000) showed that face stimuli that are composed of two different emotional expressions, such as the top section of an angry face and the bottom section of a happy face, are subject to the composite effect (see Young et al, 1987). Thus, subjects were slower in recognizing the separate expressions when both sections were closely aligned than when they were misaligned. An analogous but independent effect was also found for facial identity when different persons posed for these expressions. This indicates that dissociable types of facial information are coded by different configural processes.

The notion that identity and expression are functionally independent has also been the subject of some controversy. Schweinberger & Soukup (1998) re-examined the extent to which these dimensions could be processed independently, by selectively introducing variations in one dimension during the classification of another. They found that identity decisions were not influenced by variations in expression, but that expression decisions were slowed by variations in identity. Thus, observers could not attend selectively to expression without interference from identity information. Schweinberger & Soukup (1998) were unable to combine these findings with existing research in favour of a bi-directional processing independence between identity and expression (e.g. Campbell et al, 1996; Humphreys et al, 1993; Sergent et al, 1994). Yet, in subsequent research they succeeded in producing further support of an asymmetric relationship (Schweinberger, Burton & Kelly, 1999).

On the whole then, there is evidence that different types of facial information can be extracted independently and that at least some of these types, such as expression and identity, may also be coded separately in faces. Although this issue is convoluted by recent reports of an asymmetric dependency between identity and expression, the notion of a functional independence between these dimensions raises an intriguing question. If attention is important for holistic face encoding, then is it also required to integrate identity and expression information within the same face percept during visual processing?

1.4 Capacity limits in face processing

As described above, there is good evidence that (upright) face processing is particularly dependent on holistic information (e.g. Carey & Diamond, 1977, Tanaka & Farah, 1993; Valentine & Bruce, 1986; Yin, 1969), and that processing upright faces but not inverted faces (Palermo & Rhodes, 2002) or nonface objects (Boutet et al, 2002) disrupts the holistic encoding of another face. One implication of these results is that face processing may have its own processing limits. There are two lines in support of this premise. The first line suggests that people may be unable to ignore a *solitary* face distractor during the classification of a nonface target, even under conditions that usually extinguish distractor processing. The second line hints that face processing may be subject to capacity limits in *multiple* face displays.

Several studies have shown that task-irrelevant face distractors are processed reliably with a concurrently presented nonface target. Young, Ellis, Flude, McWeeney & Hay (1986) examined interference effects with displays composed of a printed famous name and a famous face in a response-competition task. Participants were required to classify the names as pop-stars or politicians while ignoring the face distractor, which could be congruent (e.g. Mick Jagger's name and Mick Jagger's face) or incongruent with the correct response (e.g. Neil Kinnock's name and Mick Jagger's face). They found reliable distractor congruency effects, suggesting that subjects could not prevent semantic categorization of the distractor faces.

Lavie, Ro & Russell (2003) employed a variation of this paradigm to examine the effect of perceptual load on distractor processing. The subjects' task was to categorize famous names as pop-stars or politicians while ignoring a flanking face distractor. In addition, task-relevant load was manipulated by embedding the name targets in displays of one (lowest load condition) to eight letter strings (highest load condition). Response times increased with the number of letter strings, indicating that task-relevant load was successfully manipulated. Remarkably though, the magnitude of face interference was unaffected by variations in relevant load. Furthermore, interference from nonface distractors such as photographs of fruits and musical instruments was extinguished with increasing task-relevant load in a similar task. These results led Lavie et al (2003) to suggest that face processing does not depend on any *general* capacity limits.

Comparable conclusions can be drawn from another study, in which subjects were presented with letter strings that were superimposed on photographs of famous face distractors under different load conditions (Jenkins, Burton & Ellis, 2002). Under low-load, subjects responded to the colour of the letter string, a task that poses minimal attentional demands (e.g. Treisman, 1993). In a high load condition, on the other hand, subjects were required to identify a specific letter target in the string, a manipulation that has previously been shown to eliminate distractor processing (e.g. Lavie, 1995). Distractor processing was then assessed with a surprise memory test for the names of the famous faces (e.g. "was Bill Clinton presented?") and with repetition priming, which is a facilitation in identifying an item due to prior exposure to that item. Although explicit face memory strongly deteriorated under high load, repetition priming was equivalent across conditions.

In addition, there have been a number of reports of prosopagnosic patients who, despite being explicitly unable to recognize familiar faces, nevertheless show the normal pattern of face interference when asked to make semantic classifications of names (e.g. de Haan, Young & Newcombe, 1987; Sergent & Signoret, 1992b). Collectively, these findings suggests that face processing is very robust across different manipulations, and even under conditions that should make this difficult, as long as only a single face is presented at a time. However, these results do not imply that face processing is capacity free. In fact, Lavie et al (2003) cautioned that face processing might be subject to its own capacity limits.

One source of evidence for face processing limits comes from visual search tasks. These studies demonstrate that search for a unique face target, such as a particular face or facial expression, among inverted or scrambled face distractors or upright faces with non-target expressions generates steep search slopes with increasing display size (e.g. Brown, Huey & Findlay, 1997; Kuehn & Jolicoeur, 1994; Nothdurft, 1993). This suggests that face processing limits are severe enough to require sequential identification of the items in these displays. However, what is neglected by these studies is that the number of stimuli that can be perceived simultaneously also depends on visual acuity, which is highest in the centre of the retina (the fovea) but falls off rapidly towards the periphery (see e.g. Anstis, 1974; Curcio & Allen, 1990). Subjects may thus have to foveate across different locations in visual search displays, in particular when large set sizes reinforce small individual items (as in Brown et al, 1997; Kuehn & Jolicoeur, 1994; Nothdurft, 1993). In fact, Näsänen & Ojanpää (2004) measured eye movements during visual search with faces, and found that only two to four faces can be processed during a single eye fixation. Serial search functions for multiple face arrays might therefore only reflect the limits of visual acuity.

Boutet & Chaudhuri (2001) examined face processing in a paradigm that is unlikely to suffer from visual acuity. Observers were shown stimuli composed of two overlapping faces, one rotated 45° clockwise and the other 45° counterclockwise, and had to indicate whether they could perceive both faces as whole and visibly independent entities. This was immediately followed by a test display of two rows of four faces, with each row containing one of the targets. When upright overlapping faces were used, only one of the faces was subsequently recognized. Two inverted faces, on the other hand, were perceived as an ambiguous combination of both. These results are logically similar to Palermo & Rhodes's (2002) findings, reviewed earlier, that only matching upright but not inverted flanker faces impairs the (holistic) processing of a central face target. As typical face processes are disrupted by inversion, these studies hint at a processing limit for upright faces that is independent of any general processing limits.

It should be noted that Boutet & Chaudhuri (2001) used a hypothetical situation that our face processing system is not usually confronted with. Palermo & Rhodes (2002), on the other hand, presented the flanker displays for substantial durations of ≥ 1.5 seconds that may have aided serial face processing. These studies may have also used an inappropriate recognition test to examine face processing, as stimuli may still undergo considerable processing when explicit memory is poor (see e.g. Tipper, 1985; Tipper & Driver, 1988). In fact, a recent priming study indicates that distractor faces might still be processed during a face matching task. Khurana (2000) asked participants to match the second and fourth face in a row of five faces while ignoring the three remaining distractor faces. When the distractors were presented as targets on a subsequent trial, negative priming was found. However, the three distractors were always identical and subjects were presumably scanning across the central distractor (the third face in a row of five) to match the two targets, a situation that is not dissimilar to Palermo & Rhodes's (2002) task. Moreover, targets and distractors were always presented until a response was registered. Under those conditions the faces could have been processed sequentially, again making it difficult to specify any exact face processing limits.

Jenkins, Lavie & Driver (2003) also examined the processing of multiple face distractors, but under better-controlled conditions than Khurana (2000). Subjects categorized the printed names of pop-stars or politicians in displays that were only presented for 200 ms (i.e. too briefly to permit stimulus-responsive saccades), while ignoring a critical famous face distractor that could be congruent or incongruent with the target response. An additional response-neutral distractor (neither pop-star nor politician) of an upright face, a phase-shifted version, an inverted face, or a meaningful nonface object could also be present in the display. They found that interference from the critical face distractor could be diluted by a response-neutral face distractor, but not by any other stimuli. In other words, the processing of a distractor face seemed to be reduced by competition from another face, but not by general competition from different classes of stimuli.

Despite the recurrence of this notion, Jenkins et al (2003) also obtained some evidence that face processing does not proceed *entirely* independent of general processing resources. When they repeated this design with object names as targets and critical object distractors, they found that distractor interference could be diluted by the addition of any visual stimulus, including faces. In turn, it is thus possible that face distractors processing is partly determined by task-relevant nonface load, which makes it also difficult to make a direct inference about face capacity limits from this study.

In summary, there is good evidence that faces are processed reliably alongside nonface stimuli, provided that only a single face is presented at a time (e.g. Jenkins et al, 2002; Lavie et al, 2003; Young et al, 1986), suggesting that face processing may proceed largely independent of any general processing limits. There is also a growing body of research hinting that face processing is not capacity-free, but may be limited in multi-face displays (e.g. Boutet & Chaudhuri, 2001; Palermo & Rhodes, 2002; Jenkins et al, 2003). However, none of these studies were originally designed to examine capacity limits in face processing and none have tested such limits directly. Hence it is difficult to specify the exact nature of any capacity limits in face processing.

1.5 Attention biases to faces

Another question of current interest concerns the influence that faces have on an observer's focus of attention. Evidence from visual search shows that particular faces or facial expression do not pop-out of crowded face arrays (Brown et al, 1997; Kuehn & Jolicoeur, 1994; Nothdurft, 1993). This suggests that faces do not capture attention in multiple face displays. However, there is some evidence that faces may capture attention in competition with other classes of stimuli.

Vuilleumier (2000) studied neuropsychological patients with unilateral visual extinction following brain damage to the right parietal lobe. This deficit is characterized by impaired report of stimuli in the contralesional (left) hemifield when competing stimuli are presented on the ipsilesional side, although neglect patients can still detect stimuli on the contralesional side when they are presented alone. Vuilleumier (2000) utilized this deficit to investigate whether face stimuli are less affected by visual neglect in comparison with other stimuli. He found that visual extinction was reduced for left-sided faces in comparison with meaningless shapes, scrambled faces and names when competing stimuli were presented in the other hemifield. Consequently, Vuilleumier (2000) suggested that faces might possess an advantage in capturing attention and overcoming extinction.

Mack, Pappas, Silverman & Gay (2002) also provide some evidence that faces capture attention. They presented subjects with a stream of visual items presented at a rate of 75 ms/item. Subjects had to identify a line drawing of any of five primary targets (heart, bell, fish, apple, teardrop) within this stream and detect the presence of a closely following secondary target of a happy face icon, an inverted happy face icon, or a tree shape. Under these conditions, the face targets were detected approximately 90% of the time compared to the inverted faces and trees which were only detected on between 40-70% of trials.

Although the use of rather artificial face stimuli in these studies offers little insight into the processing of real faces, which are visually more complex and informative, others have reported a similar advantage for photographs of faces. Ro, Russel & Lavie (2001) alternated displays that were composed of meaningful objects (appliances, clothes, food, musical instruments, and plants) and a solitary human face with blank screens, so that the stimulus displays appeared to flicker. During these alternations, one of the items could suddenly change into another exemplar from the same category. Results indicated that changes were detected more rapidly and accurately in faces than in any of the other objects, an advantage that disappeared when all stimuli were inverted. Consequently, Ro et al (2001) concluded that real faces also have a special capacity for drawing attention.

However, Palermo & Rhodes (2003) contested this interpretation, instead reasoning that these results might reflect an "odd-one-out" advantage as Ro et al (2001) only ever presented one face among a range of nonface objects. In support of this idea, they found a similar change detection advantage when a single nonface object was embedded among several face stimuli. Intriguingly though, they failed to replicate Ro et al's (2001) original findings despite using the same method and stimulus set. Nonetheless, these results suggest that faces behave similar to other objects in a change detection task. Therefore, despite several attempts to demonstrate attention biases to faces, there is no compelling evidence that *realistic* face stimuli capture attention.

Nonetheless, there is also some mixed evidence that the ability to capture attention might depend on the type of face stimuli used. Eastwood, Smilek & Merikle (2001) found shallower search slopes for a sad face icon among neutral face distractors than for a happy face icon. They concluded that emotion information may be perceived outside the focus of attention and can be used to guide that focus to a particular face. Angry schematic faces are also detected very efficiently in visual search, although they do not pop-out (Fox, Lester, Russo, Bowles, Pichler & Dutton, 2000). In addition, it has been shown that trait-anxious subjects respond faster to a target probe when its location is correctly cued by an angry face, an effect that has been attributed to an attentional bias towards threatening stimuli in these individuals (e.g. Bradley, Mogg, Falla & Hamilton, 1998; Mogg & Bradley, 1999).

Emotional expression may also affect the attentional dwell-time or the disengagement of attention from faces. Eastwood, Smilek & Merikle (2003) showed that subjects were slower to count upturned and downturned arches in a display when these were embedded in faces with negative as opposed to positive expressions. Other studies show that trait-anxious or trait-angry individuals also have particular difficulty in disengaging attention from threatening faces (e.g. Bradley et al, 1998; Fox, Russo, Bowles & Dutton, 2001; Fox, Russo & Dutton, 2002; van Honk, Tuiten, de Haan, van den Hout & Stam, 2001). For example, Fox et al (2001) found that responses to a dot probe were delayed in trait-anxious individuals compared to normal subjects, but only when a threatening face incorrectly cued the probe location. In contrast, response times were evenly matched for neutral and happy face cues. Thus, there is evidence that faces may have some limited ability to seize attention, depending on their emotional connotation and probably also the emotional state of the observer. However, although several studies report similar biases with threatening words and pictures (e.g. Amir et al, 2003; Fox et al, 2001; Yiend & Mathews, 2001), expressive faces have never been compared directly with other stimulus classes within the same experiment. Hence it is unresolved whether a *general* disengagement bias exists for faces compared to other classes of stimuli, independent of threat-related information and observers' emotional traits.

1.6 Structure of this thesis

The aim of this thesis was to investigate the relation between attention and face processing across different selection mechanisms. The first experimental chapter examined whether observers can selectively respond to facial expression and identity, by measuring the effect of systematic variations in one of these dimensions on the classification times of the other dimension (Experiments 1 & 2). Although there is already substantial evidence for the functional independence of these types of facial information, this idea has recently been challenged in a similar task (Schweinberger et al, 1999; Schweinberger & Soukup, 1998). This is followed by an attempt to determine whether attention is required to encode expression and identity information from the same face into a multi-dimensional visual percept (Experiment 3). To assess this, a response-competition task was used in which subjects classified face targets according to particular identity-expression conjunctions while ignoring two distractor faces. Response-critical identity and expression information was either combined in one of these distractors (the conjunctive condition) or dispersed across both (the disjunctive condition). Distractor congruency effects on target response times were then contrasted for these conditions to determine if correct conjunction information was available under inattention.

A response-competition task was used again in Chapter 3, but now to establish capacity limits in face processing. This was done by comparing interference from face and nonface distractors during the classification of face and nonface targets. The first experiment used speeded sex judgements to unfamiliar faces and short names (Experiment 4). Subsequent experiments employed semantic judgements to famous faces and famous names (Experiment 5), famous faces and pictures of national flags (Experiments 6 & 7), and a combination of both (Experiment 8).

The purpose of Chapter 4 was to provide a stricter test for face processing limits. It has been shown that nonface distractors can still give rise to priming when they do not interfere with target classification (Driver & Tipper, 1989). Chapter 4 therefore investigated capacity limits in face processing by measuring distractor priming from multiple-item displays. Two experiments assessed the effect of face and flag target processing on face distractor priming (Experiments 9 & 10). A third priming study used face-like and face-unlike nonface targets to explore the visual properties that may be responsible for eliciting face processing limits (Experiment 11).

The final empirical chapter investigated attention biases for faces in comparison with a range of nonface objects in a simple detection task. Over three experiments, subjects were required to shift attentional resources from the location of a face or a nonface distractor to the location of a peripheral line target (Experiments 12-14). Contrasting the effects of different distractor types on target RTs was then used to assess attentional disengagement. A final study examined attentional engagement by faces and nonface objects. The target and distractor locations were now switched so that the subjects were attending to the target at the start of each trial, while the distractors were presented in the visual periphery (Experiment 15).

Chapter 2 Dissociating and Integrating Facial Expression and Identity

Introduction

Established models of face recognition postulate separate parallel routes for the processing of different categories of facial information, such as sex, emotional expression and identity (notably, Bruce & Young, 1986). In particular, the idea that identity and expression are dissociable cognitive functions has been supported by a number of observations from a range of methodologies. These include neurological studies demonstrating double dissociations in brain-injured participants (e.g. Humphreys, Donnelly & Riddoch, 1993; Parry, Young, Saul & Moss, 1991; Schweinberger, Klos & Sommer, 1995; Young, Newcombe, de Haan, Small & Hay, 1993), functional imaging studies showing spatially-dissociable areas of brain activation during the processing of expression and identity (George et al, 1993; Sergent, Ohta, MacDonald & Zuck, 1994), and cognitive studies of neurologically normal participants, which have shown that observers can selectively attend to facial expression and identity in time-stressed categorization tasks (Bruce, 1986; Calder, Young, Keane & Dean, 2000; Campbell, Brooks, de Haan & Roberts, 1996; Young, McWeeney, Hay & Ellis, 1986a).

Recently, however, the idea that expression and identity processing are functionally independent has become the subject of some controversy. Using Garner's (1974, 1976) selective attention paradigm, Schweinberger & Soukup (1998) re-examined whether facial expression and identity can be dissociated in a speeded categorization task. Within this paradigm, introduced in detail later, identity information was processed independent of facial expression but contributed to expression analysis, an intriguing asymmetric relationship among the perception of facial identity and expression.

In a subsequent study, Schweinberger, Burton & Kelly (1999) examined whether such a relationship is related to differences in processing speed. If identity is perceived faster than facial expression, then variations in identity may be more likely to affect expression processing than vice versa. To manipulate processing speeds, Schweinberger et al (1999) employed a morphing technique to create a photographic continuum between two faces. Depending on the percentage contribution of each original face within any point along this continuum (e.g. 30% versus 70%, 20% versus 80%, etc.), morphs were consistently categorized as the facial identity contributing the most. In contrast, classification RTs increased along and peaked towards the middle of the continuum (i.e. the point where each of the originals contributes 50% to the new image). Thus, when two different person's faces with the same expression were morphed, recognition of identity and expression remained relatively unaffected but identity classification times reflected the decreased perceptual salience of the stimuli. Likewise, if two images of one person depicting different expressions were morphed, expression RTs correlated with the morphing continuum. In this way, Schweinberger et al (1999) selectively manipulated the processing speed of expression and identity. However, despite this manipulation the asymmetric relationship first observed by Schweinberger & Soukup (1998) persisted, suggesting that differences in processing speed cannot account for a functional dependence between these types of facial information.

Intriguingly, Schweinberger et al's (1998, 1999) claims have received little support from other face perception studies with Garner's technique. Etcoff (1984) measured participants' ability to sort cards depicting different expressions and identities, but found that expression and identity were classified without interference from the other. Deprived of advanced timing systems, however, Etcoff (1984) used a manual stopwatch to record sorting times. According to Schweinberger & Soukup (1998), this might account for the discrepancies between Etcoff's (1984) results and their own study, in which reaction times were measured with millisecond accuracy. However, similarly to Etcoff's (1984) findings, Le Gal & Bruce (2002) assessed the independence of *sex* and expression judgements to faces with the Garner technique and found that both dimensions were processed independently. As there is evidence that sex classification also proceeds independent of identification (e.g. Bruce, Ellis, Gibling & Young, 1987; Ellis, Young & Flude, 1990), these findings do not rule out an interaction between expression and identity processing. Therefore, they do not provide sufficient evidence to dismiss Schweinberger et al's (1998, 1999) reports. The important point, as far as the present study is concerned, is that an asymmetric processing-dependence between facial expression and identity has previously not been found with other methodologies, nor with Garner's technique, nor have analogous effects been observed with other face processing routes within Garner's paradigm. Thus, the purpose of Experiments 1 & 2 was to investigate the validity of Schweinberger et al's (1998, 1999) claims. Closer examination of these studies suggests that their response pattern may have resulted from asymmetric treatment effects within the paradigm. In the remainder of the introduction Garner's paradigm is outlined and these treatment effects are then discussed in detail.

Garner's Selective Attention Paradigm

Garner (1974, 1976) originally devised this paradigm to examine whether basic object properties, such as form and colour, require shared or independent processing resources. More specifically, he asked whether selective attention to a task-relevant stimulus dimension is possible when variation is added to a second task-irrelevant dimension. If both dimensions are separable, then selective attention to a relevant dimension should be possible regardless of variations in irrelevant information. In a typical Garner experiment, participants are required to make speeded two-choice judgements to four types of stimuli, consisting of the crossing of two exemplars each of two distinct stimulus dimensions. For example, if the dimensions are colour and shape, participants are instructed to classify colour while ignoring the shape of a stimulus, or to ignore colour while classifying shape.

During classification, these stimuli are presented in three experimental conditions. In the *control* condition, stimuli vary along the relevant dimension while the irrelevant dimension remains constant. To illustrate, in a colour categorization task participants may be shown a block of only squares and a second block of only circles, and the stimuli in both blocks must be classified as either green or blue. In the *orthogonal* condition, stimuli vary along relevant and irrelevant dimensions. So, for example, blue squares, green squares, blue circles, and green circles are intermixed within one block. In the *correlated* condition, relevant and irrelevant information is covaried; for example, only green squares and blue circles.

The influence of the irrelevant dimension on the relevant one is determined by contrasting performance across these three conditions. Increased RTs in the orthogonal condition in comparison with the control condition show that task-irrelevant information interferes with the classification of the task-relevant dimension. This is an indication that both dimensions are processed in an integral manner. In contrast, comparable RTs for control and orthogonal condition suggest that both dimensions are dissociable and processed separately. Garner also claimed that faster RTs in the correlated compared to the control condition can be interpreted as evidence for integral processing of two stimulus dimensions, a so-called redundancy gain. This advantage apparently arises due to the invariant combination of relevant and irrelevant information, which facilitates the perception of both dimensions as a unitary event. However, redundancy gains may also arise when participants strategically choose to process the easier of two correlated dimensions, relevant or irrelevant. Consequently, redundancy may be used to support claims for the integral processing of two dimensions, but is not sufficient to establish such claims.

ASYMMETRIC TREATMENT EFFECTS WITHIN THE GARNER PARADIGM

A major disadvantage of Garner experiments is that they can be sensitive to asymmetric treatment effects. Schweinberger et al (1999) eliminated one such effect as a possible explanation for an asymmetric dependency between identity and expression, that is, relative differences in processing speed of both face dimensions. Yet, these studies contain several other potential asymmetric treatment effects that merit further examination.

i) Picture-based response cues versus face-related information

One criticism of Schweinberger et al's (1998, 1999) studies is the use of an extremely limited stimulus set. Schweinberger & Soukup (1998) used a total of just four faces (2 identities x 2 expressions), and Schweinberger et al (1999) used the same number of stimuli to create twenty-eight morphed faces. These sets were repeated over 600 and 280 trials in the respective studies. As a consequence of excessive repetition, subjects might have developed alternative task strategies rather than engaging in typical face processes. Pictures of a person may, for example, bear superficial similarities such as image brightness and colour contrasts (see Figure 2.1 overleaf). Such similarities could result from all images of one person being taken under particular lighting conditions on the same day or even the same time of day. If such similarities are salient and correlated with identity, then participants might learn to distinguish both stimulus identities on the basis of these picture-based characteristics. And if such salient cues intrude on the classification of expression, even if participants can usually attend selectively to expression without interference from identity information, then this might produce the orthogonal interference reported by Schweinberger et al (1998, 1999). Crucially, facial expressions are less likely to correlate with picture-based cues because they are displayed by both identities. Thus, by facilitating identity but not expression classification, picture-based cues may have contributed to an asymmetric response pattern.

ii) Internal features versus external features

The use of a limited stimulus set raises another potential problem as faces contain a variety of cues to identity. These include the spatial relation of internal features (e.g. eyes, nose, mouth), which can communicate very subtle but unique differences between people (e.g. Haig, 1984; Tanaka & Sengco, 1997; Valentine & Bruce, 1986), and external features, such as hairstyle and face outline, which can change frequently and may be shared by different people. Whereas external cues are insufficient to distinguish between the many faces encountered in everyday life, in an experiment with a limited set size they might provide a salient and simple strategy to classify identity (see Figure 2.1). Comparable cues for the classification of expression would be unavailable. Akin to picture-based cues, external feature processing could thus contribute to an asymmetric relationship by providing identity-correlated information that may intrude on expression perception.

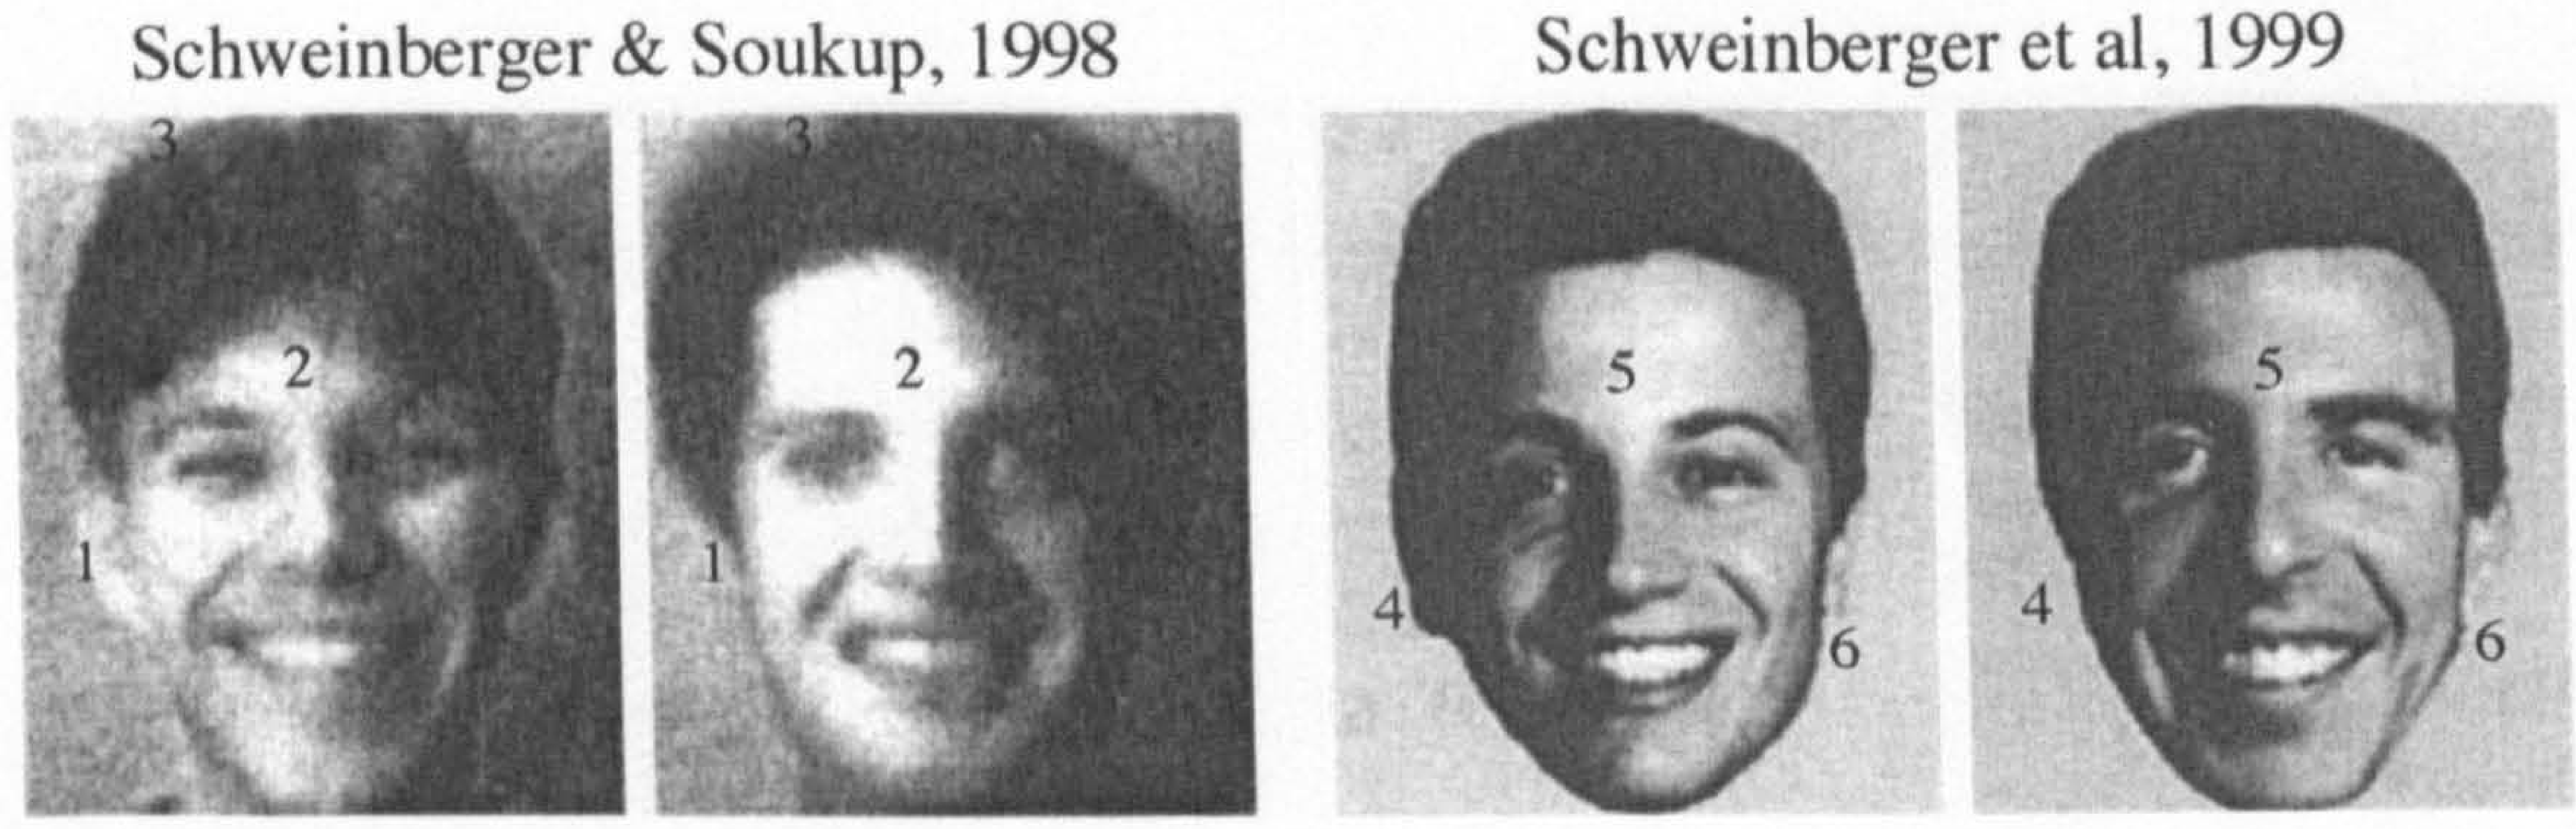

Figure 2.1 Examples of the stimuli used by Scheinberger & Soukup (1998) and Schweinberger et al (1999). The photographs reveal a number of picture-based identity cues and external features that may have been exacerbated by the use of a small stimulus set. These include differences in luminance (2,5), in hairstyle (2,3) and in face outline (for example, note the presence and absence of ears; 1,4,6).

iii) Effects of face familiarity

The potential contribution of external identity cues within the Garner paradigm also draws attention to the role of face familiarity. It is well established that different processes are involved in recognizing familiar and unfamiliar faces (e.g. Burton, Wilson, Cowan & Bruce, 1999; Hancock, Bruce & Burton, 2000). Although we may rely on external features, such as hairstyle and face outline, to identify well-known people, familiar faces are primarily recognized from internal facial features (e.g. Ellis, Shepherd & Davis, 1979; Young, Hay, McWeeney, Flude & Ellis, 1985). Unfamiliar person recognition, on the other hand, tends to rely more on those prominent external features (Bruce, Henderson, Greenwood, Hancock, Burton & Miller, 1999; Ellis et al, 1979; Young et al, 1985). Consequently, one might expect external identity cues to influence expression classification particularly when unfamiliar faces are used. Alternatively, facial identity might only interfere with expression if participants are capable of distinguishing between the stimulus identities. Yet, this may not be the case with unfamiliar faces even when external identity cues are available. Face familiarity might thus play a crucial role in understanding asymmetric interactions between identity and expression. Notably, Schweinberger and associates (1998, 1999) failed to specify face familiarity in their studies. Consequently, it remains unresolved whether their findings reflect general face processes or more specific processes related to either unfamiliar or familiar face processing.

iv) Asymmetric increases in task-relevant information

A fourth problem within the Garner paradigm are asymmetric increases in task-relevant information. As alluded to earlier, integrated processing of two task dimensions is primarily assessed within this paradigm by increasing the ratio of task-irrelevant to relevant information in the orthogonal condition. This is achieved by presenting stimuli in this condition that vary along both, the relevant and irrelevant task dimension. Although such an increase in information may leave the amount of task-relevant information intact when simple shapes and colours are used, expression and identity information may be coded by to some extent overlapping physical features. In other words, increases in irrelevant information in the orthogonal condition may be accompanied by analogous increases in relevant information. This is problematic if participants use substantially different strategies for expression and identity classification. For example, if identity is classified on the basis of external features or picture-based cues, then this may not be overly affected by increasing irrelevant information in the orthogonal condition, as the distinctive physical feature remains relatively intact across different images. To the contrary, the same expression can show considerable variation within as well as between different persons. Unlike identity, increasing irrelevant information in the orthogonal condition may thus increase relevant information when expression is classified. If this results in an increase in task difficulty, one might predict Schweinberger et al's (1998, 1999) asymmetric response pattern. However, this would not reflect orthogonal interference from the task-irrelevant dimension, and should not be interpreted as a functional dependency between independent processing routes.

In summary, Schweinberger et al's (1998, 1999) studies contain several potential confounds that may have contributed to an asymmetric response pattern within Garner's selective attention methodology. These are: i) a severely limited stimulus set, which may have enabled participants to perform identity categorizations on the basis of salient picture characteristics; ii) similarly, the use of external features for identity classifications; iii) effects of face familiarity; and iv) asymmetric increases in task difficulty between identity and expression in the orthogonal condition, especially if identity-correlated task strategies were available. The aim of the present experiments was to investigate whether an asymmetric relationship between expression and identity persist when the potential impact of these confounds is reduced. To address these concerns, the current experiments used a variation of Schweinberger et al's (1998, 1999) task. To discourage the use of picture-based response cues during the classification of identity, a substantially larger stimulus set was now used, consisting of 120 different images. To diminish the influence of external features on identity classification, the stimuli consisted of faces photographed from a variety of viewpoints and, on half of all pictures, with a cap to disguise hairstyle. Because of this large and varied set - each stimulus was only encountered once during the experimental trials of each condition - these changes were also designed to eliminate asymmetric increases in task difficulty from the control to the orthogonal condition. Finally, to examine the influence of face familiarity, identity and expression processing was contrasted across two experiments with participants who were unfamiliar (Experiment 1) and familiar with the stimulus identities (Experiment 2).

Experiment 1

Experiment 1 examined whether the processing of facial expression is contingent on facial identity information within the Garner paradigm, as was suggested by Schweinberger et al (1998, 1999). Rather than revealing a novel functional architecture of identity and expression processing, these findings might reflect asymmetric treatment effects within this paradigm. To address this, the present experiment employed a variation of Schweinberger et al's (1998, 1999) task. Faces were classified according to either expression or identity, while ignoring the other dimension. However, in order to prevent asymmetric treatment effects, a large and varied stimulus set, consisting of 120 digital photographs, was used. In addition, to investigate effects of face familiarity all participants were unfamiliar with the stimulus identities presented in this experiment. If expression processing is affected by task-irrelevant variations of identity in unfamiliar face processing, then an asymmetric interaction analogous to the one reported by Schweinberger et al (1998, 1999) should be found. On the other hand, if previous findings reflect asymmetric treatment effects and if observers can selectively attend to these types of facial information, then variations in identity should not interfere with expression processing.

Method

Subjects Thirty-six undergraduate students from the University of Glasgow, whose ages ranged from 18-25 years, volunteered to participate in the experiment for a small fee. All reported normal or corrected to normal vision and were unfamiliar with the faces they were to encounter in the Experiment.

Stimuli & Apparatus An Apple Macintosh computer presented the stimuli and recorded responses using SuperLab 1.74. Digital photographs of two male employees from the Psychology department at the University of Glasgow served as stimuli. Each model posed for portraits of two emotional expressions (happy & angry) from three different viewpoints (full-face, left, right). To add variation, head-shifts were performed unrestrained during the recording of these images. In addition, both models posed with a cap on half of all photographs to disguise hairstyle (see Figure 2.2 overleaf). In total, 120 photographs were taken with a Fuji FinePix6800 digital camera, consisting of five images under each level of identity (Person A vs. Person B), facial expression (happy vs. angry), viewpoint (full-face, left, right), and hairstyle (cap vs. no cap). These images were displayed in greyscale at a size of 4.5 cm x 6.0 cm.

Figure 2.2 Example displays from Experiment 1. The target set contained two identities, Person A and Person B, depicted from a variety of viewpoints and wearing a cap on 50% of trials to disguise hairstyle. The faces displayed angry facial expressions (see left two columns) and happy facial expressions (right two columns).

<u>Design</u> The experiment had a 2 x 3 mixed design, with the between-subjects variable of group at two levels (identity vs. expression classification) and condition at three levels (correlated, control, orthogonal). Dependent measures were reaction times (RTs) and error percentages.

<u>Procedure</u> The procedure for the identity task was as follows. Prior to the main task, the 18 participants in this group were shown 3 colour photographs (full-face, left, right) of each identity, printed out and mounted on card, for approximately 30 seconds for familiarisation purposes. Participants were then told that the task involved making identity decisions as quickly and as accurately as possible to the faces of these persons presented on a computer screen. In addition, subjects were emphatically instructed to ignore facial expression. Each trial began with a central fixation cross, displayed for 1500 ms. This was replaced by a face stimulus, which remained visible until a response had been made. Following a response, the face was replaced by the fixation cross, marking the start of the next trial. Subjects were requested to respond by pressing the "D" or the "L" key on a standard

computer keyboard. Button-press latencies were measured from stimulus onset and feedback for incorrect responses was given immediately by a short warning tone.

All subjects underwent 2 consecutive blocks for each of the three conditions. For the correlated condition, the 1st block contained happy faces of Person A and angry faces of Person B, and the 2nd block contained angry faces of Person A and happy faces of Person B. In the control condition, the 1st block contained happy faces of both Persons A and B, and the 2nd block contained angry faces of Persons A and B. In the orthogonal condition, all possible combinations of expression and identity were presented within both blocks. Each of the six blocks consisted of 20 practice and 60 experimental trials. Thus, with a stimulus set of 120 images, each image was only encountered once during the experimental trials of each condition. However, a third of these images were also encountered during practice. Trial order was randomized within each block and the order of conditions was counter-balanced across all subjects. Subjects could rest between blocks, initiating the next block by pressing the space bar.

The same procedure was used for the expression classification task, except for the following changes. The 18 participants in this group were also familiarized with the face identities but were instructed to ignore identity while making two-choice expression decisions. In addition, the composition of the correlated and orthogonal conditions remained the same, but the control condition now contained happy and angry faces of Person A in one block, and happy and angry faces of Person B in the other block.

Results

Errors Error rates were generally low. In the identity group, errors were made on 4.1% of all correlated trials, 5.6% of all control trials, and 3.5% of all orthogonal trials. For the expression group, the error rates were 2.4%, 3.2%, and 3.5% respectively. Error rates were not analyzed further.

RTs: Comparisons between classification tasks across conditions. The median correct reaction times (RTs) were computed for each level of group (identity vs. expression) and condition (correlated, control, orthogonal). The averages of these RTs across subjects are shown in Figure 2.3.

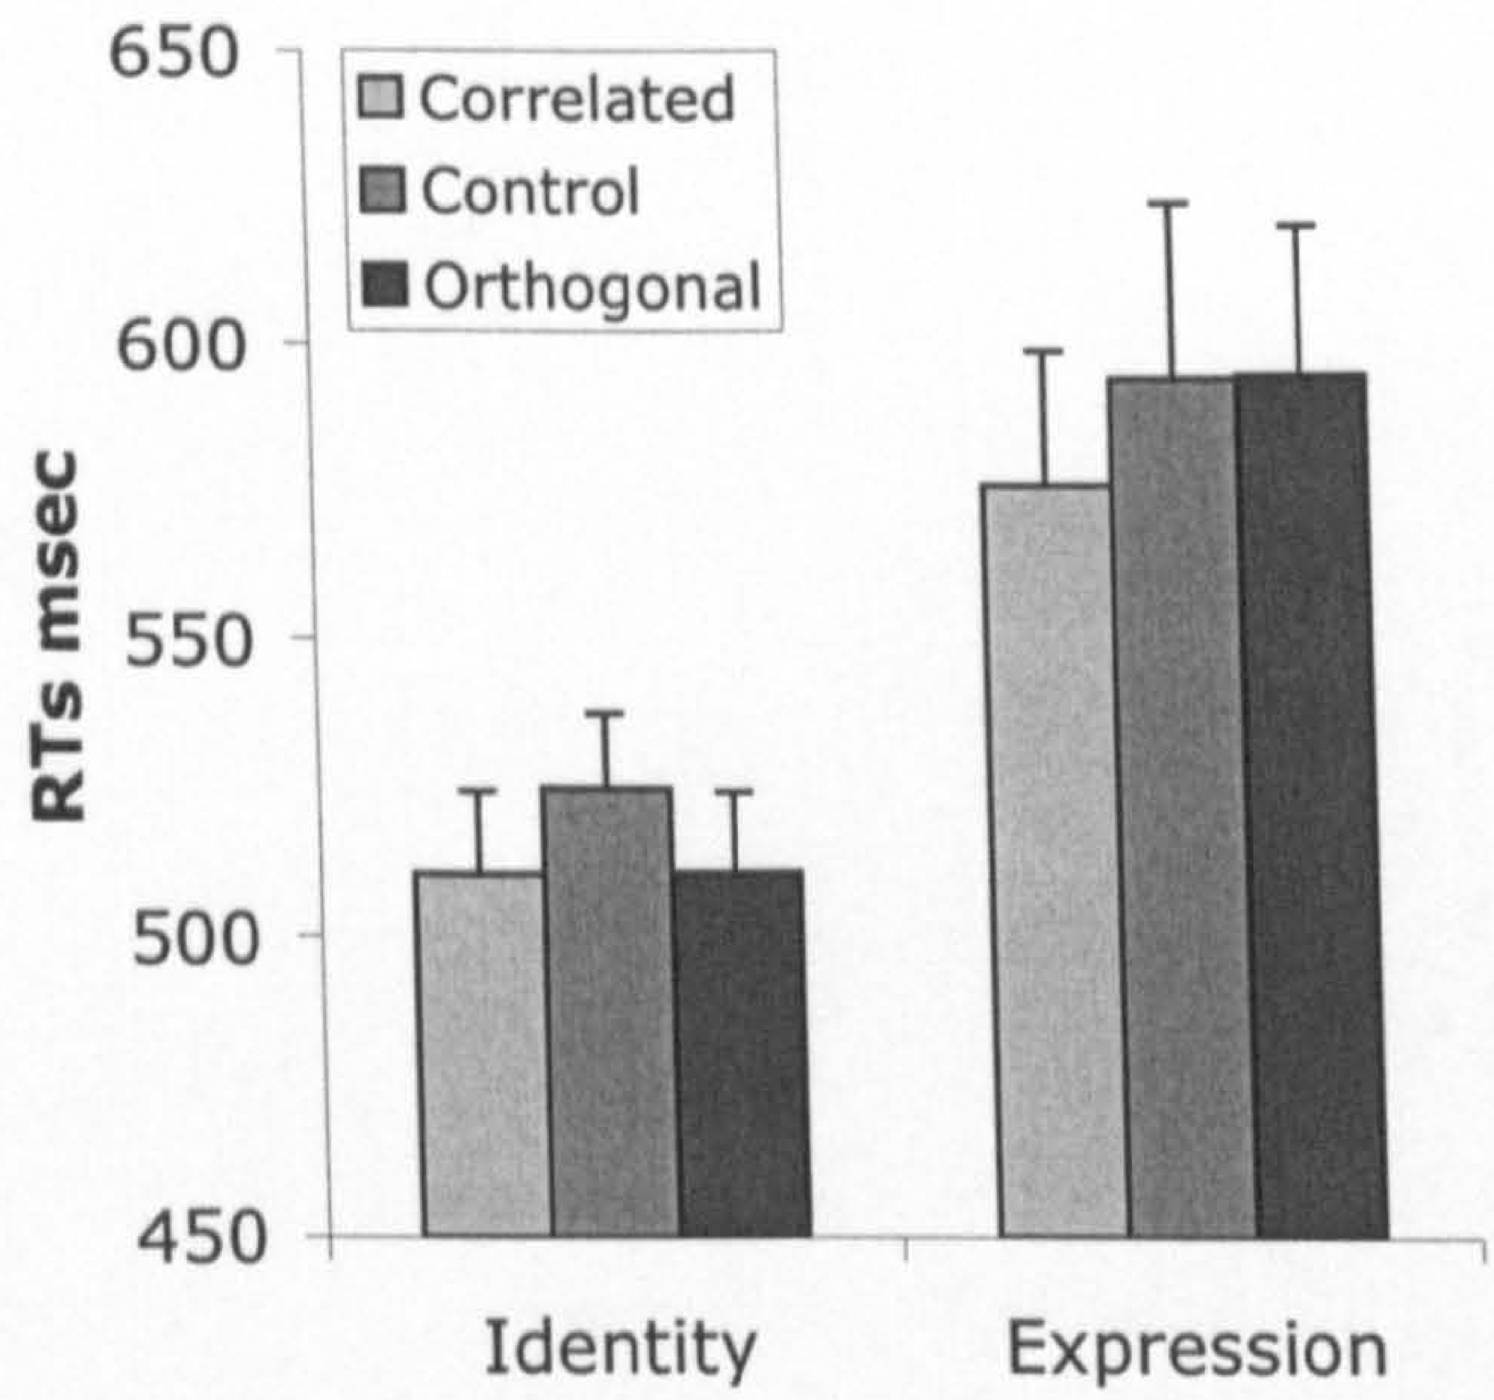

Figure 2.3 Means of the median reaction times (RTs, in msec) as a function of the Garner conditions and group in Experiment 1. Vertical bars represent the positive standard errors of the means.

As can be seen from Figure 2.3, classifications times were faster for identity than expression decisions. However, neither the identity nor the expression group showed an RT advantage in the control condition in comparison with the orthogonal condition, indicating that task-irrelevant information did not influence task-relevant processing in this experiment. These observations were confirmed by a 2 (identity vs. expression) x 3 (correlated, control, orthogonal) mixed analysis of variance (ANOVA), which showed a main effect of group, $F(1,34)=7.31$, $p<.05$, reflecting faster responses to identity than to expression, but no main effect of condition, $F(2,68)=1.49$, and no interaction between group and condition, $F(2,68)<1$.

RTs: Comparisons within each classification task. Separate 3 (correlated, control, orthogonal) x 2 (happy vs. angry) x 2 (Person A vs. Person B) ANOVAs were carried out for each level of group for a more detailed analysis within the Garner conditions. The RTs for every combination of these factor levels may be seen in Figure 2.4.

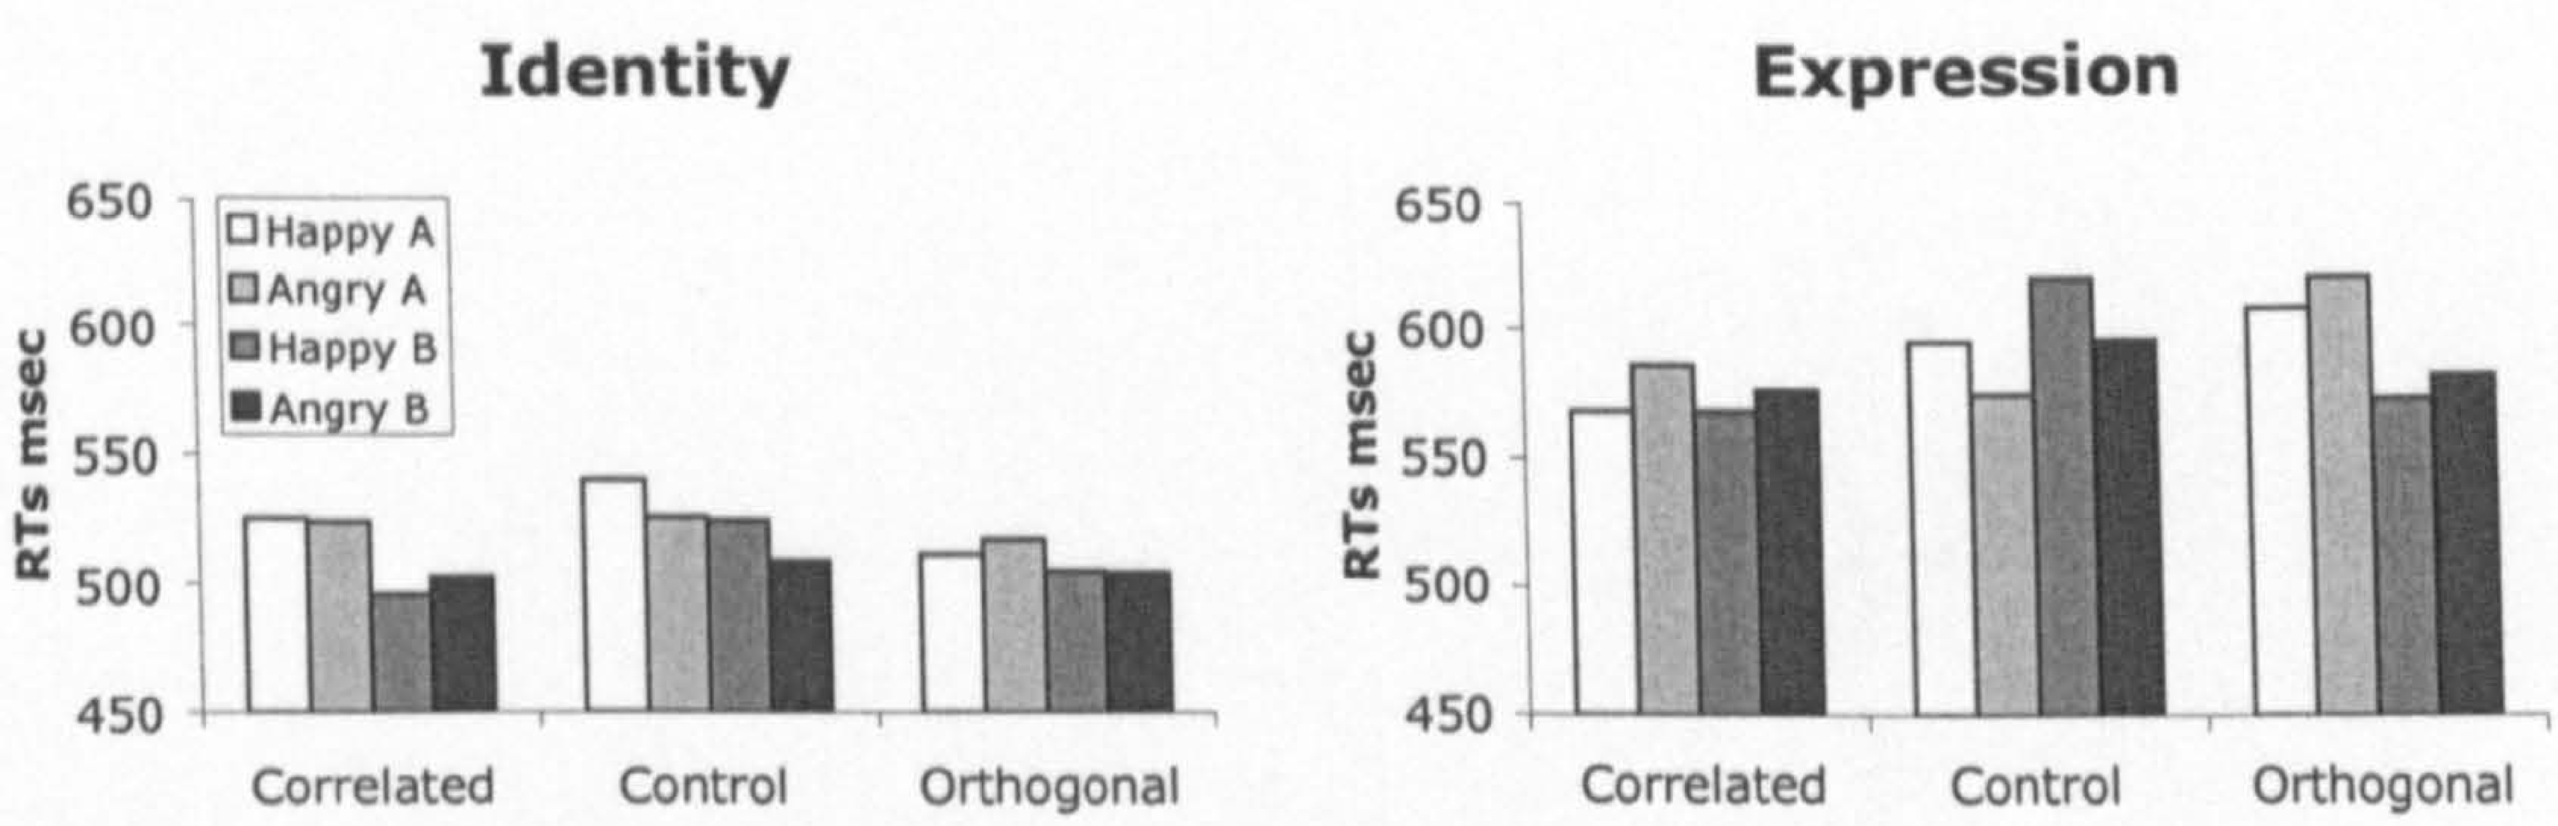

Figure 2.4 Means of the median reaction times (RTs) for every combination of each relevant and irrelevant dimension for Experiment 1. A = face of Person A; B = face of Person B.

For identity classifications, no main effects of condition, $F(2,34)=1.03$, or expression were found, $F(1,17)<1$, again indicating that identity processing was unaffected by expression. A main effect of identity, $F(1,17)=6.73$, $p<.05$, is interpreted as reflecting slightly slower RTs to Person A than to Person B across all identity conditions. No other comparisons were significant.

For expression classifications no main effects of condition, $F(2,34)=1.43$, expression, $F(1,17)<1$, or identity, $F(1,17)=2.03$, were found. However, a significant interaction between identity and condition was observed, $F(2,34)=6.15$, $p<.01$. Simple main effect analysis revealed similar RTs for both face identities in the correlated condition, $F(1,17)<1$, but significantly faster RTs to Person A than Person B in the control condition, $F(1,17)=10.12$, $p<.01$, and the reverse pattern in the orthogonal condition, $F(1,17)=24.85$, $p<.01$. Importantly however, no simple main effects of condition were found for either identity (Person A, $F(2,34)=1.80$; Person B, $F(2,34)=1.78$). Thus, expression classification was not systematically affected by variations in identity across the Garner conditions.

Discussion

Experiment 1 examined recent claims of an asymmetric dependency in face processing, such that expression perception may be contingent upon identity (Schweinberger et al, 1998, 1999). Specifically, the aim was to determine whether the response pattern of previous studies might have arisen from asymmetric treatments effects within the Garner paradigm. Using a modification of Schweinberger et al's (1998, 1999) design, the current experiment sought to diminish the contribution of such effects by using a substantially larger and more varied stimulus set. The RT data show that identity classifications were faster than expression classifications, which indicates that the facial identity may have been more discriminable in the present stimulus set. However, the critical comparisons for establishing an interaction between identity and expression processing, those between the control and orthogonal conditions, revealed no significant differences for the classification of either face dimension. This indicates that participants were able to attend to each dimension selectively, and, at least initially, appears to contradict claims of an asymmetric interaction.

However, Experiment 1 only investigated the processing of unfamiliar faces. There is considerable evidence that unfamiliar face processing, in comparison with familiar face recognition, may rely particularly on external features for person identification, such as hairstyle and face-outline (e.g. Bruce et al, 1999). In the introduction it was suggested that such information might have contributed to Schweinberger et al's (1998, 1999) findings if unfamiliar face stimuli were used. The current results do not support this idea, although the design involved several manipulations to diminish the contribution of external identity cues. Even so, since external features provide only an unreliable source of identity information, the deduction of external identity cues merely serves to underline the validity of the present results.

Experiment 2

The results of Experiment 1 provide support for the hypothesis that identity and expression perception are dissociable cognitive functions, albeit only when unfamiliar faces are processed. However, an objection could be raised as to whether the use of unfamiliar face stimuli is appropriate to investigate the relation of expression and *identity* processing. Although the participants in the identity condition were able to distinguish the different faces quickly and with few errors, the participants in the expression condition may not have learned to do so. Thus it is possible that variations in identity did not interfere with expression classification precisely because the participants in this group had not learned to distinguish the task-irrelevant face identities. To assess whether the absence of any interference may have been due to this, Experiment 2 used participants that were familiar with the stimulus identities. If expression processing is also unaffected by identity information from familiar faces, this would provide further support for a functional independence between expression and identity within this paradigm. On the other hand, if expression processing is contingent on face familiarity, this would provide some support for Schweinberger et al's (1998, 1999) functional interaction.

Method

Subjects, Stimuli & Procedure Thirty-six new subjects, whose ages ranged from 21-33 years, volunteered to participate in the unpaid experiment. All subjects were postgraduate students or research staff from the Department of Psychology at the University of Glasgow and reported normal or corrected to normal vision. The subjects were familiar with the face identities they were to encounter in the experiment and could identify them without delay prior to the task. Apparatus, stimuli and procedure were identical to those of Experiment 1.

Results

Errors As in Experiment 1, error rates were generally low. In the identity group, errors were made on 4.2% of all correlated trials, 4.5% of all control trials, and 4.6% of all orthogonal trials. For the expression group, the error rates were 2.8%, 4.6%, and 4.5% respectively. Error rates were not analyzed further.

RTs: Comparions between classification tasks across conditions. The median correct reaction times (RTs) were computed for each level of group (identity vs.

expression) and condition (correlated, control, orthogonal). The averages of these RTs across subjects are shown in Figure 2.5.

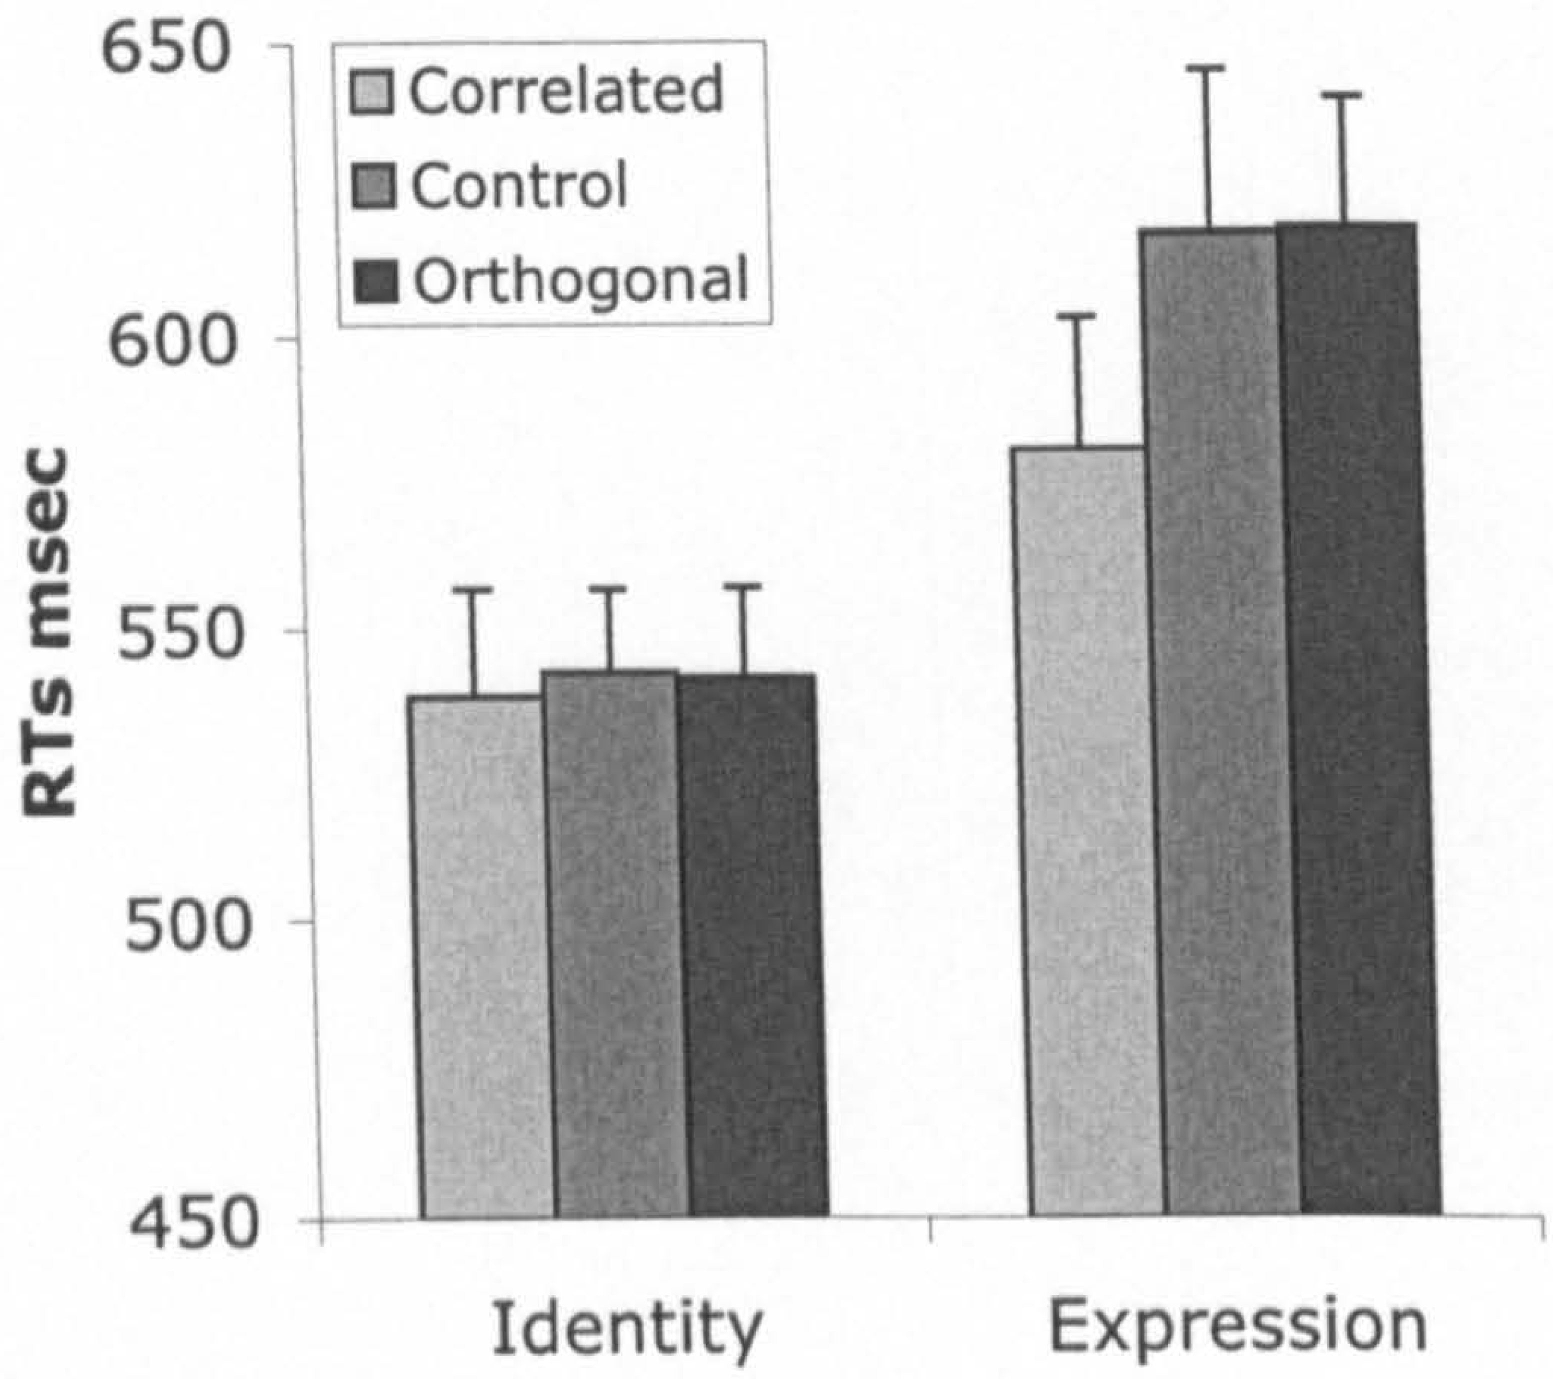

Figure 2.5 Means of the median reaction times (RTs, in msec) as a function of the Garner conditions and group in Experiment 2. Vertical bars represent the positive standard errors of the means.

A 2 (identity vs. expression) x 3 (correlated, control, orthogonal) mixed ANOVA revealed a main effect of group, $F(1,34)=5.52$, $p<.05$, reflecting faster RTs in the identity task than in the expression task, a main effect of condition, $F(2,68)=4.54$, $p<.05$, and an almost reliable interaction between group and condition, $F(2,68)=3.10$, $p=.052$. As can be seen from Figure 2.5, although RTs were noticeably faster in the correlated condition than in the control and orthogonal conditions of the expression group, RTs were evenly matched across all conditions in the identity group. This was confirmed by a simple main effect of condition when making expression decisions, $F(2,68)=7.57$, $p<.01$, but not for identity decisions, $F(2,68)<1$.

RTs: Comparisons within each classification task. Separate 3 (correlated, control, orthogonal) x 2 (happy, angry) x 2 (Person A, Person B) ANOVAs were carried out for each level of group to examine variability across every combination of the relevant and irrelevant dimension. The RTs for these combinations are in Figure 2.6.

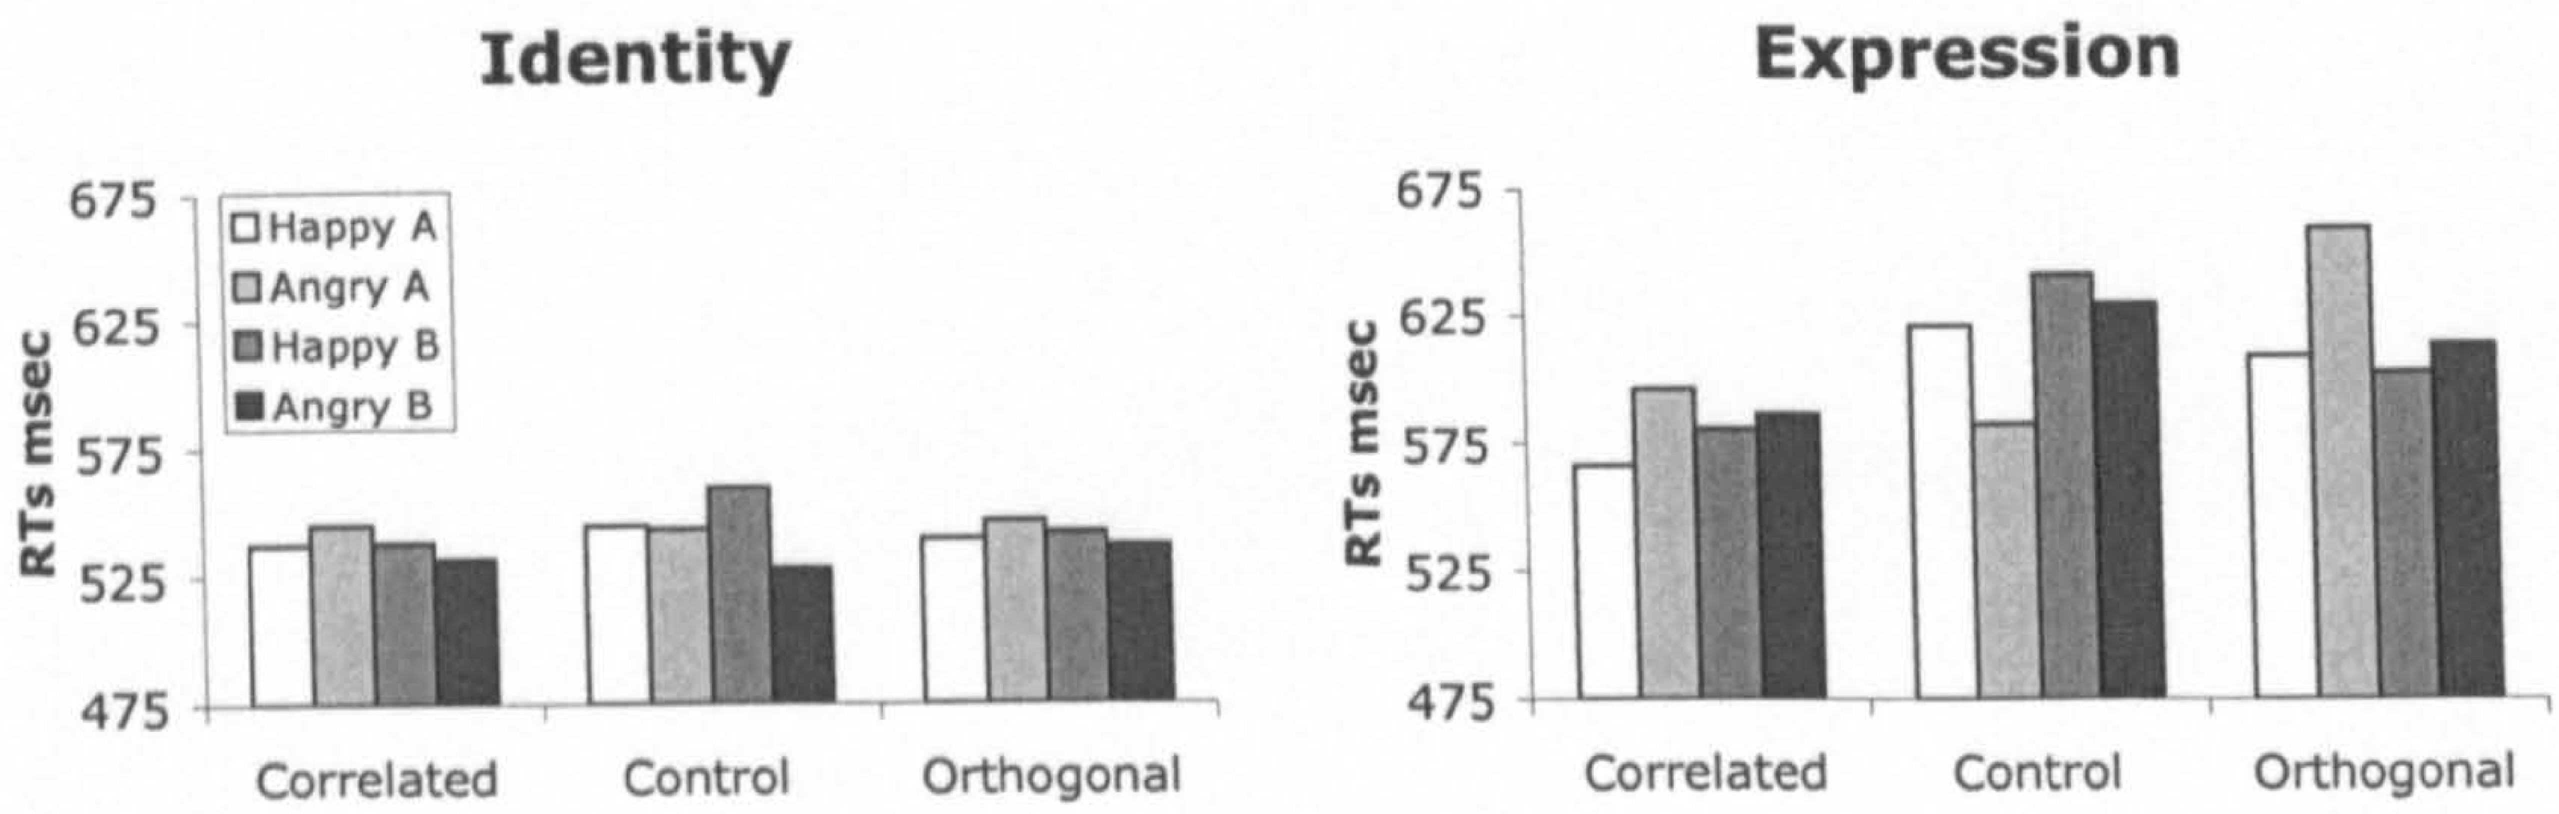

Figure 2.6 Means of the median reaction times (RTs) for every combination of each relevant and irrelevant dimension for Experiment 2. A = face of Person A; B = face of Person B.

For identity classifications, no main effects of condition, $F(2,34)<1$, expression, $F(1,17)=1.00$, or identity were found, $F(1,17)<1$. An interaction between expression and identity, $F(1,17)=6.51$, $p<.05$, reflects slightly slower RTs for angry expressions of Person A than for angry expressions of Person B, $F(1,17)=6.12$, $p<.05$, while RTs to happy expressions were more evenly matched, $F(1,17)=1.41$, and an almost reliable difference between happy and angry expressions for Person B, $F(1,17)=3.94$, $p=.06$.

For expression classifications, no main effects of expression, $F(1,17)=1.35$, or identity were found, $F(1,17)<1$. However, an effect of condition was confirmed, $F(2,32)=5.17$, $p<.05$. Newman-Keuls comparisons (alpha = .05) showed that RTs in the correlated condition were significantly different from the control and the orthogonal condition. More importantly, no differences were found between these latter conditions. The effect of condition was also modified by an interaction of expression and condition, $F(2,34)=6.62$, $p<.01$, reflecting faster RTs to angry than to happy expressions in the control condition, $F(1,17)=4.78$, $p<.05$, and the reverse pattern in the orthogonal condition, $F(1,17)=6.87$, $p<.05$. In addition, happy expressions were classified faster in the correlated condition relative to the control condition, Tukey's HSD test, $p<.05$.

A second complex interaction was observed between identity and condition, $F(2,34)=12.73$, $p<.01$, reflecting faster RTs to Person A than to Person B in the control condition, $F(1,17)=17.45$, $p<.01$, and the reverse pattern in the orthogonal condition, $F(1,17)=10.61$, $p<.01$. In addition, simple main effects of condition for each identity were found (Person A, $F(2,34)=3.92$, $p<.05$; Person B, $F(2,34)= 3.71$, $p<.05$). In depth analysis of these effects revealed faster RTs for Person A in the correlated condition in comparison with the orthogonal condition, Tukey's HSD, $p<.05$, and faster RTs for Person B in the correlated condition in comparison with the control condition, Tukey's HSD, $p<.05$. No other comparisons were significant.

Discussion

As in Experiment 1, the comparisons critical for establishing a functional interaction – those between the control and orthogonal conditions – were not significant for identity or expression decisions. This result therefore further contradicts Schweinberger et al's (1998, 1999) claims of an asymmetric dependency between expression and identity processing, and extends the findings of Experiment 1 to familiar faces. In addition, there is some support for the idea that identity and expression are not just dissociable functions but are processed in parallel (see e.g. Bruce & Young, 1986). This can be seen from a redundancy gain during the classification of expression, reflecting significantly faster RTs in the correlated than in the control and the orthogonal condition. As was reviewed earlier, redundancy gains may be used to support claims of integral processing between two stimulus dimensions, provided that reliable orthogonal interference is found. However, in the absence of orthogonal interference, as was the case in this experiment, redundancy gains most likely arise when participants use correlated irrelevant information to enhance performance. Here, identity decisions were consistently faster than expression decisions, which would suggest that it was indeed possible to use irrelevant identity information to decrease response times to correlated expression decisions. Notably, a similar non-significant redundancy pattern was observed during unfamiliar face processing in Experiment 1, which suggests that the participants in the expression group may have been able to distinguish the task-irrelevant identities to some extent.

The observation of a redundancy gain is interesting in so far as it is an indication of participants' proclivity to employ alternative strategies within the Garner paradigm. Thus, it provides some tentative support for the hypothesis that participants may have developed alternative strategies to produce an asymmetric response pattern in Schweinberger et al's (1998, 1999) studies. In contrast to those studies, the present experiments used a substantially larger and more varied stimulus set to reduce the contribution of such strategies. Nonetheless, Experiment 2 revealed several complex interactions, particularly during the classification of facial expression, which suggest that some stimulus groupings may have been more discriminable than others. However, similar to the overall response pattern, none of the individual stimulus groupings revealed a functional dependency between the processing of expression and identity.

Therefore, the results of Experiment 1 and now also of Experiment 2 suggest that facial expression and identity processing are dissociable cognitive functions. Observers can selectively attend to each type of information without interference from the other, both during the processing of familiar and unfamiliar faces. Although these findings contradict recent claims of an asymmetric interaction between the processing of identity and expression, they support the long-standing view that facial expression and identity are perceived independently and in parallel (e.g. Bruce & Young, 1986).

Experiment 3

The findings from Experiments 1 and 2 converge with numerous claims that observers can selectively attend to facial expression and identity (for a review see Young, 1998). Furthermore, it is thought that these processes occupy spatially dissociable brain areas (e.g. George et al, 1993; Sergent et al, 1994), and rely on distinct types of visual information (Calder et al, 2000). However, considering this wide-ranging dissociation, it is perhaps surprising that, upon encountering a face, expression and identity are perceived as belonging to the same face percept. In fact, the perceptual experience of these facial dimensions appears remarkably integrated: We can accurately extricate a face's identity *and* expression without confusing them with those from another face. This opens the intriguing question as to how dissociable types of facial information are accurately combined within the same face percept during visual processing.

Outside the face domain, visual attention has long been viewed as a crucial resource for full, integrated perception. According to one influential account, attention acts like glue during visual encoding that binds different types of information belonging to the same stimulus (e.g. Treisman, 1988, 1993; Treisman & Gelade, 1980; Treisman & Schmidt, 1982; see also Lavie, 1997). Thereby, object features such as colour and shape are perceived independently under conditions of inattention, but are combined into a multidimensional, conscious percept through focused visual attention.

So far, the role of attention in face encoding has been considered by just a few studies and these have concentrated on only one face dimension - the perception of identity (Boutet, Gentes-Hawn & Chaudhuri, 2002; Palermo & Rhodes, 2002; Reinitz, Bartlett & Searcy, 1997; Reinitz, Morrissey & Demb, 1994). All of these studies examined whether attention integrates featural (or part-based) facial information into holistic percepts, in which these features and their spatial relation are captured as an inseparable source of information. This was done by manipulating attention to or away from faces during learning and by measuring whole/part recognition at a subsequent test phase. As face processing is particularly dependent on holistic information (e.g. Tanaka & Farah, 1993; Tanaka & Sengco, 1997), it seems plausible that distinct facial features, similar to Treisman's (1988, 1993) object features, might be integrated into holistic percepts by attention. However, previous studies failed to produce consistent results, with some suggesting that holistic processing requires attentional encoding (Palermo & Rhodes, 2002; Reinitz et al, 1994), but others reporting that holistic processing proceeds without attention (Boutet et al, 2002; Reinitz et al, 1997).

As even immediate memory for faces appears remarkably poor, these studies may have suffered from possible memory confounds (see e.g. Simons & Levin, 1998). The inconsistencies might also reflect the use of facial features, which were either defined in terms of local face characteristics corresponding to particular lexical values, such as the eyes, nose and mouth (Palermo & Rhodes, 2002; Reinitz et al, 1994, 1997), or particular face regions, for instance the top versus the bottom face half (Boutet et al, 2002). However, this may not be compatible with how facial features are actually represented by the brain, which might not be spatially distinct but perhaps represent different types of configural face information, such as expression and identity (see e.g. Calder et al, 2000). Therefore, if attention is a crucial resource for feature encoding in face perception, then it is possible that it may be involved in the integration of expression and identity information from the same face into a complete percept. To examine this, a variation of Lavie's (1997) response competition paradigm was used.

Similar to Garner's (1974, 1976) selective attention methodology, Lavie's (1997) paradigm was originally applied to basic object attributes, such as colour and shape, to assess the visual integration of such features. During Lavie's (1997) task, participants were required to respond to a central target, while ignoring two flanking distractors. Responses were based on particular conjunctions (combinations) of two features. For example, participants pressed one key for a purple cross or another for a green circle on a critical trial, but withheld responses on non-critical trials when the target consisted of the opposite conjunctions (i.e. a green cross or a purple circle). Thus, colour or shape alone was not sufficient for correct target classification.

Importantly, the unattended distractors were also processed and could influence target responses by way of their response congruency. On a congruent trial, the distractors would contain the same target features, thus facilitating activation of the correct response. On an incongruent trial, on the other hand, response times to a particular target conjunction (e.g. a purple cross) were slowed by the presence of an incongruent set of target features amongst the distractors (e.g. a green circle). The distractors could also influence target classification through the combination of their features. Thus, response-associated distractor information (both congruent and incongruent) was presented either as a conjunction, in which response-critical features were presented within one common distractor, or as a disjunction, in which critical features were separated across both distractors. To illustrate, in a conjunctive congruent trial participants may have been shown a purple cross target flanked by a purple cross distractor on one side and a response-neutral distractor on the other side (e.g. a brown triangle). In a disjunctive congruent trial, on the other hand, a stimulus display may have consisted of a purple cross target flanked by a purple triangle and a brown cross.

Comparing target-distractor congruency effects between conjunctive and disjunctive conditions was critical for establishing the role of attention in colour and shape integration. Thus, if colour and shape information are not separately accessible from unattended objects, then conjunctive stimuli should have interfered more with target classification than disjunctive distractors, which did not match the targets accurately. Alternatively, if conjunction information is unavailable under inattention, then the colour and shape of unattended conjunctive and disjunctive distractors should have interfered equally with target classification, because their individual features were equal in terms of target-distractor congruency. This is precisely what Lavie (1997) found, conjunctive and disjunctive distractors produced equivalent target-distractor congruency effects.

The present experiment examined whether these findings can be extended to expression and identity information from faces. In order to make Lavie's (1997) paradigm relevant for such a task, responses were now based on expression-identity conjunctions. On critical trials, which made up two thirds of all trials, participants responded to the happy face of one person (Person A) or the angry face of another person (Person B). On non-critical trials, on the other hand, participants pressed a single response key for the opposite expression-identity combinations. Additionally, two irrelevant face distractors were presented to the left and right of the target, of which one identity and one expression were either congruent or incongruent with the target. In the *Conjunctive* condition, this identity and expression information was conjoined in one of the distractors (e.g. happy Person A) with the other distractor displaying two response neutral features (e.g. surprised Person C). In the *Disjunctive* condition, the same information was disjoined across both distractors (e.g. happy Person C, surprised Person A).

Displays were also included in which distractors were congruent (or incongruent) in expression or identity alone to assess their specific contribution to distractor congruency effects. It is possible that distractor interference could arise from just one face dimension. For example, in the preceding experiments it was found that identity information was more discriminative than expression. If the same applies here, then distractor interference may predominantly reflect identity information. Equivalent congruency effects between the *Conjunctive*, the *Disjunctive* and the *Identity* condition, but not for the *Expression* condition would reveal this. Finally, in one further condition both distractors in a display consisted of features from outwith the possible target set (the *Neutral* condition), for example, a happy face target of Person A flanked by surprised Person C on one side and sad Person D on the other. This baseline condition was included to examine whether any congruency effects reflect interference or facilitation during target classification.

If identity and expression require attention for visual integration into a single multi-dimensional face percept, the following RT pattern may be predicted. The *Conjunctive* and *Disjunctive* conditions should show reliable target-distractor congruency effects, with slower RTs to incongruent than congruent distractors. These congruency effects should be equivalent, since under conditions of inattention the only way in which these stimuli differ, their conjunctive format, should be inaccessible. Furthermore, if these congruency effects arise from both types of facial information, then the *Identity* and *Expression* conditions should also reveal noticeable distractor interference.

Method

Subjects Eighteen undergraduate students from the University of Glasgow, whose ages ranged from 19-24 years, volunteered to participate in the experiment for a small fee. All reported normal or corrected to normal vision.

Stimuli & Apparatus An Apple Macintosh computer presented the stimuli and recorded responses using Superlab 1.74. The stimuli consisted of greyscale photographs of four males (person A, B, C, and D) from the Ekman and Friesen (1976) pictures of facial affect. Of each person a happy, angry, surprised, and sad picture were used. To remove extraneous background, all images were ellipseshaped and measured 3.5 cm x 4.9 cm (3.4° x 4.7° of visual angle (VA) at a viewing distance of 60 cm, fixed by means of a chinrest). These faces were used to construct stimulus displays containing a central target with one distractor to the left and one to the right. The nearest target-distractor contours were 0.6 cm (0.6° of VA) apart.

There were two main types of stimulus displays, corresponding to the critical and non-critical trials of the experiment. In critical trials, the targets consisted of angry Person A or happy Person B. These targets were combined with distractors under five conditions. In these conditions, with the exception of the *Neutral* condition, some of the distractor features could be either congruent (same response category) or incongruent (different response category) with the target. In the *Conjunctive* condition, response-congruent (or incongruent) expression and identity features were conjoined within one distractor, with the other distracor containing two response-neutral features (i.e. Person C or D, sad or surprised expressions; see Figure 2.7 overleaf). In the *Disjunctive* condition, these features were disjoined between both distractors (see Figure 2.7). In the *Identity* condition, the distractors contained just one congruent (or incongruent) identity feature, with the remaining three distractor features displaying response-neutral features. Similarly, in the *Expression* condition, the distractors contained just one congruent (or incongruent) expression feature while the remaining three distractor features were response neutral.

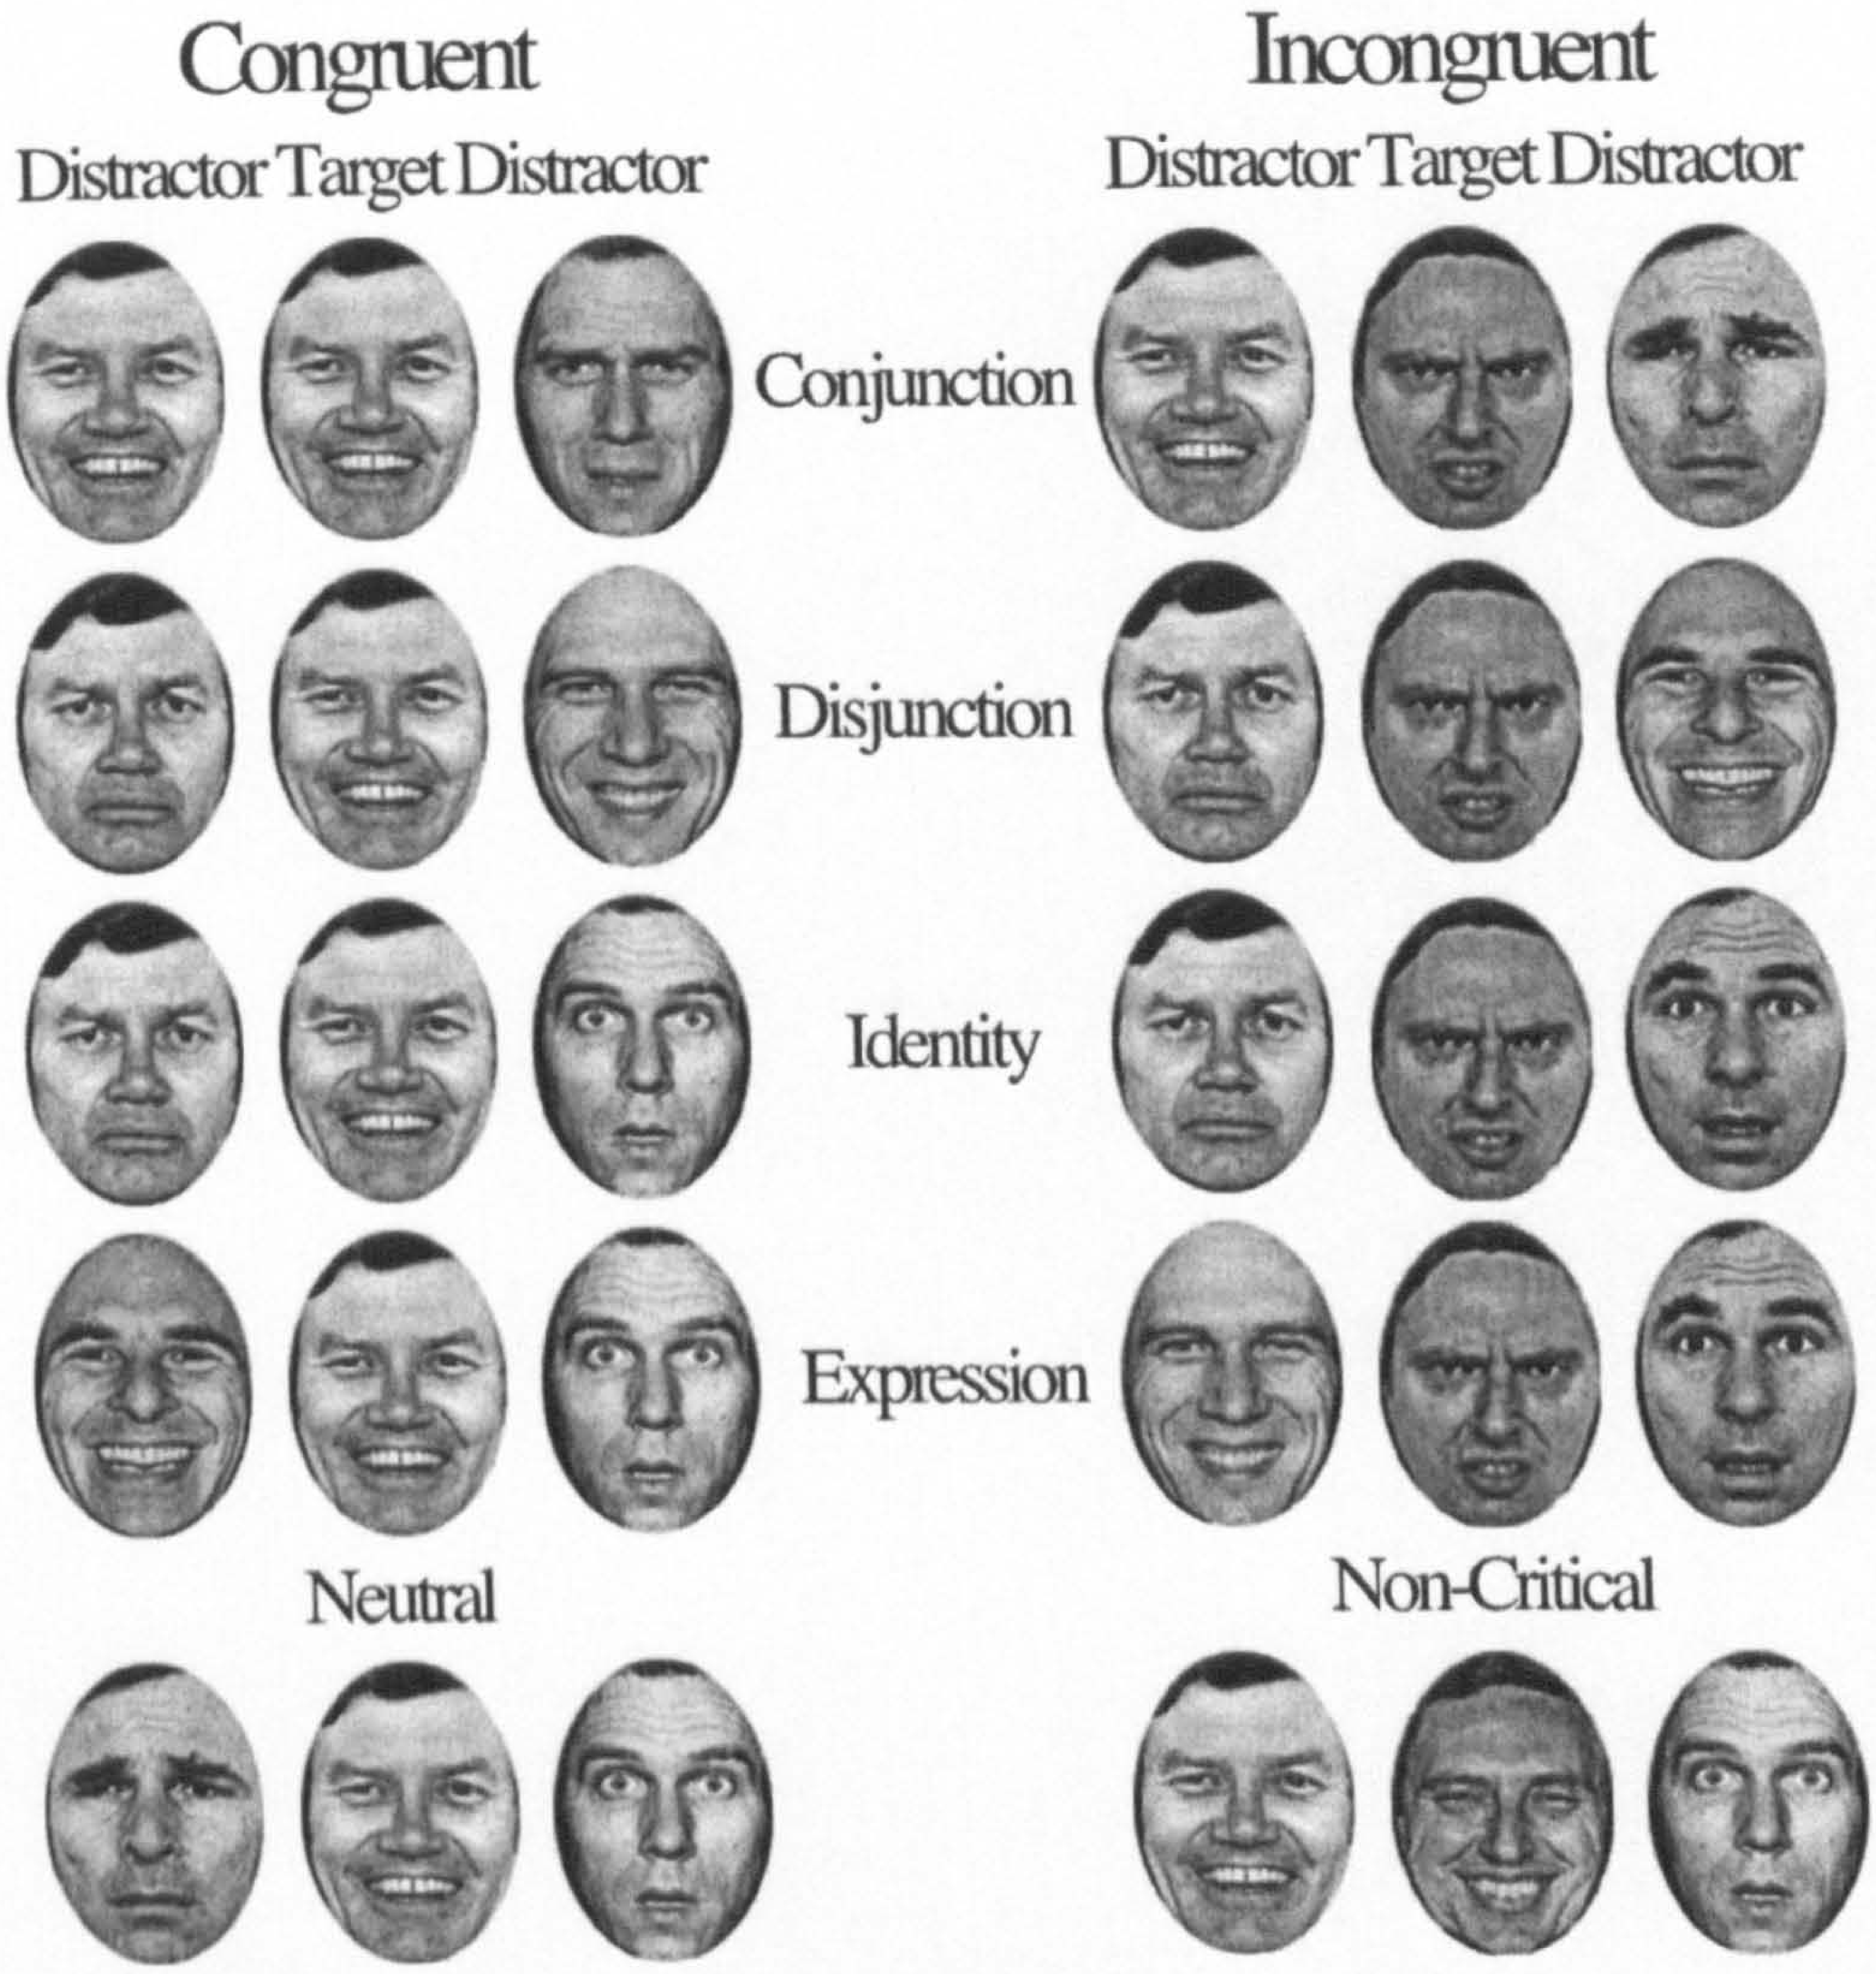

Figure 2.7 Example displays from Experiment 3. The critical targets were angry Person A (right column) and happy Person B (left column). In the *Conjunction* and *Disjunction* conditions, one distractor identity and one distractor expression were either congruent (same response-category) or incongruent (different response category) with the target. In the *Identity* condition, one of the distractors displayed a response-associated identity (Person A or Person B). In the *Expression* condition, one of the distractor faces displayed a response-associated expression (happy or angry). The remaining distractor features were always response-neutral (i.e. Person C, Person D, sad, surprised). In the *Neutral* condition, all distractor features were response-neutral (left column, bottom row). Non-critical targets consisted of the opposite expression and identity combinations (e.g. happy Person A; right column, bottom row).

For non-critical displays, the targets consisted of the opposite expression and identity conjunctions (i.e. happy Person A and angry Person B; see Figure 2.7), but contained exactly the same distractors as the critical trials. As a consequence, target-distractor compatibility was ambiguous on non-critical trials. Note, however, that the purpose of these trials was singularly to ensure that participants were responding to the correct combinations of both target dimensions, expression and identity, and response-congruency effects were not analyzed for these conditions. Thus, the same distractor combinations were used for non-critical as for critical trials to avoid cuing their presence via the distractors.

Pairing the critical targets (angry Person A, happy Person B) in the *Conjunctive*, *Disjunctive*, *Identity*, and *Expression* conditions, with each of the neutral distractors (Person C and D, sad and surprised), and under each level of congruency (congruent vs. incongruent) resulted in a total of 64 stimulus displays. For non-critical targets, 64 analogous displays were made. In addition, 16 *Neutral* displays (8 critical, 8 non-critical) were created by pairing critical and non-critical targets with only response-neutral distractor features.

Procedure

Prior to the main task the participants underwent a training phase to learn the four face identities: Participants were shown four arrays of the face identities, each array depicting the four faces (Person A, B, C & D) with a different expression (angry, happy, sad, surprised). This was followed by 3 blocks of 64 trials (4 identities x 4 expressions x 4) in which the faces had to be classified according to identity (Blocks 1 & 3) and expression (Block 2). Response accuracy was emphasized in the instructions and feedback for incorrect responses was given by a warning tone. Participants had to achieve less than 10% errors for identity classifications to proceed to the main task. Of 18 participants, 12 completed the training phase twice to meet these criteria.

In the main task, participants were told that the task involved making speeded decisions to expressive face targets, presented at fixation, while ignoring two flanking distractor faces. Subjects were requested to respond to angry Person A by pressing the "D" key and to happy Person B by pressing the "L" key on a standard computer keyboard, and to press the <space> bar to both happy Person A and angry Person B. Each trial began with a fixation cross, displayed for 1000 ms, followed by a target-distractor display for 200 ms (i.e. too briefly to permit stimulus-responsive saccades to distractors), ending with a blank screen until a response had been made. All subjects underwent 12 blocks of 54 randomly-ordered trials (36 critical & 18 non-critical trials). The experimental conditions were randomly intermixed within blocks. Blocks 1 and 2 served as practice and were excluded from analysis.

Results

Errors Incorrect responses were made on 3.1% of critical trials and were evenly matched across all conditions (see Figure 2.8 overleaf). For non-critical displays, errors were made on 7.9% of all trials. As critical trials were twice as likely as non-critical trials, this difference might reflect anticipatory response strategies. Overall, however, participants were accurate. This indicates that responses were based on both target features, expression and identity. Errors were not analyzed further.

RTs The median correct RTs were computed as a function of distractor type (*Conjunctive*, *Disjunctive*, *Identity*, *Expression*, *Neutral*) and target-distractor congruency (congruent vs. incongruent). The means of these RTs are shown in Figure 2.8.

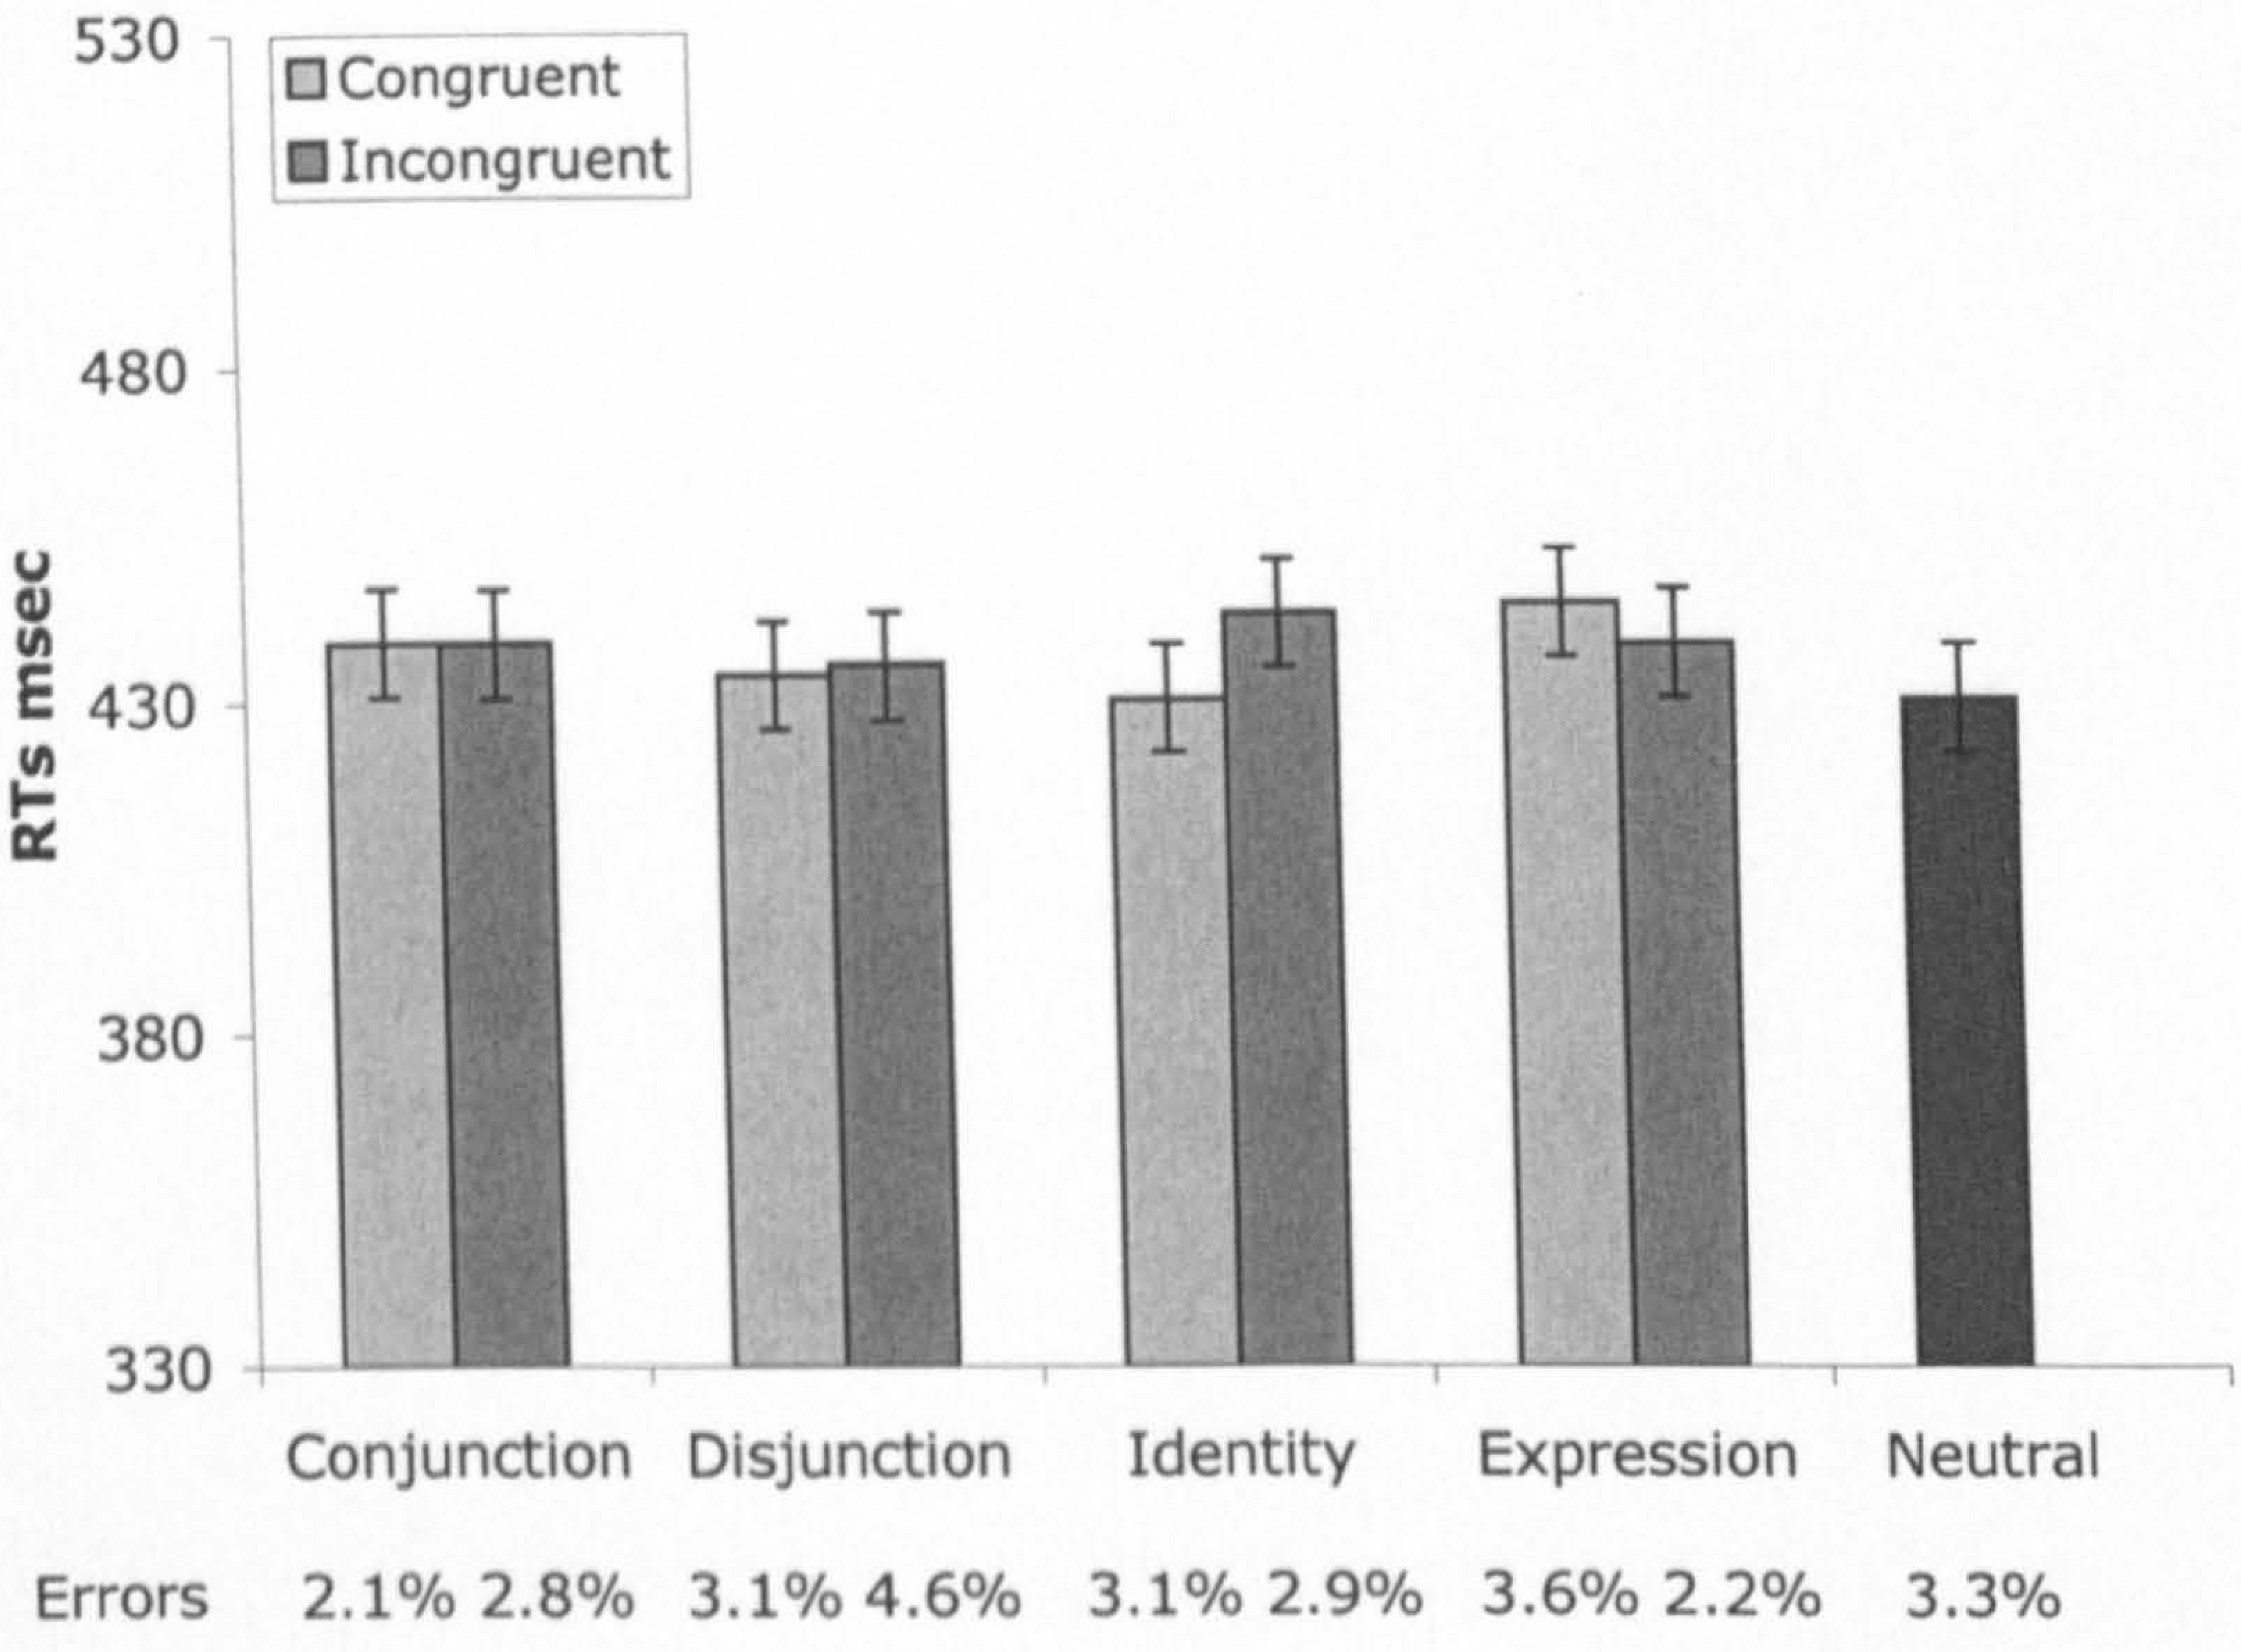

Figure 2.8 Means of the median reaction times (RTs, in msec) and percentage errors as a function of distractor type and congruency for the critical conditions in Experiment 3. Vertical bars represent the standard error (SE) of the means.[1] Non-critical trials: mean RT, 503 ms; error rate, 8%.

Similar to error rates, RTs were slightly slower for non-critical than critical trials, which provides further evidence that participants might have formed expectancies to both types of trials. More important, neither the *Conjunction* nor the *Disjunction*

[1] Throughout this thesis, standard error bars for within-subject designs are based on within-participant variability (see Loftus & Masson, 1994). This provides identical error terms for all conditions by legitimately ignoring between-subject variance.

condition showed a congruency-based response pattern and the *Expression* condition actually revealed a marginal reverse trend (- 6 ms). In fact, only the *Identity* condition revealed markedly faster RTs to congruent than to incongruent displays. However, a 2 (congruent vs. incongruent) x 4 (*Conjunction, Disjunction, Identity, Expression*) within-subjects ANOVA failed to find an effect of congruency, $F(1,17)<1$, condition $F(3,51)<1$, or an interaction between both factors, $F(3,51)<1$, indicating that there were no differences between any of these conditions. A separate one-factor ANOVA of condition (*Conjunction* vs. *Disjunction* vs. *Identity* vs. *Expression* vs. *Neutral*) was conducted to compare performance in the *Neutral* condition with all other critical conditions. This did not reveal differences between any of the conditions, $F(8,136)<1$.

Discussion

The present study employed a variation of Lavie's (1997) response competition paradigm to examine the attentional encoding of facial expression and identity. Face encoding was assessed by measuring interference from distractor features (i.e. expression and identity) during the classification of a face target, whereby the distractors' features could be either congruent (same response category) or incongruent (different response category) with the target. Importantly, congruent (or incongruent) features were either combined in one of two distractors (the *Conjunctive* condition) or separated across both distractors (the *Disjunctive* condition). According to this design two opposing predictions were made. If expression and identity require attention to be integrated into one face percept, then the conjunctive format of the unattended distractor faces should have been inaccessible. Thus, *Conjunctive* and *Disjunctive* conditions should have yielded equivalent congruency effects. Conversely, if attention is not required for perceptual integration, then conjunctive distractors should have produced larger congruency effects than disjunctive distractors by providing exact matches (i.e. the correct combinations of expression and identity) for the target set.

Intriguingly, the RT data do not support either of these predictions, as no congruency effects were found in the *Conjunctive* or the *Disjunctive* condition. Indeed, of all conditions, a slight congruency pattern (of 13 ms) was only observed in the *Identity* condition. However, within this paradigm a congruency effect in the *Identity* condition alone is improbable, as the *Conjunctive* and *Disjunctive* conditions provide the same amount of potentially distracting identity information. This was confirmed by the statistical analysis, which showed that there were no reliable congruency effects in any of the distractor conditions. This is even more striking in comparison with the *Neutral* condition, in which the distractors consisted exclusively of identities and expressions from outwith the target set. Thus, if response-associated distractor features produced *any* target-distractor interference, independent even of congruency, then they should have produced dissimilar RTs with the *Neutral* condition. However, although average RTs were marginally faster in the *Neutral* condition than in the other conditions, none of these conditions differed statistically from each other.

The failure to obtain distractor interference in this experiment means that the role of attention in the integration of facial expression and identity information remains unresolved. Thus, the present study contributes little in terms of data to the existing research concerned with the role of attention in face encoding (Boutet et al, 2002; Palermo & Rhodes, 2002; Reinitz et al, 1994, 1997). Nonetheless, the current design highlights a number of important issues that have been neglected in previous studies, and that may require consideration in future research. First, Reinitz et al (1994, 1997) required participants to study novel faces under divided- or full-attention, followed by a subsequent face recognition test to examine the effects of attention on face encoding. Palermo & Rhodes (2002) and Boutet et al (2002) also examined the effects of unfamiliar face encoding under divided or full attention with a subsequent recognition test. However, even *immediate* recognition memory for unfamiliar faces appears remarkably poor (Simons & Levin, 1998). This implies that previous findings might have been subject to memory confounds. By contrast, the present paradigm measured interference from simultaneously presented faces, thus minimizing any memory demands. Although this particular design was unsuccessful, future studies should also try to separate effects of attention and memory in face processing.

Second, previous studies examined the encoding of facial features into holistic faces, whereby features were defined in terms of lexical values (e.g. eyes, nose, mouth) or face regions (e.g. top half vs. bottom half), even though this may be incompatible with how features are actually presented by the brain (e.g. Ellis et al, 1997). The current design avoided similar, arbitrary definitions by capturing different facial 'features' in terms of the distinctive meaning that they convey, such as expression and identity. This has the added advantage of focusing on more than the recognition of identity, which is just one type of information that can be derived from faces.

In view of these advantages, the complete absence of distractor interference in this experiment is surprising. Particularly, since Lavie (1997) obtained reliable colour-shape interference in a similar design. Of course, this could be explained in terms of the visual attributes of faces, which are visually more complex than simple objects. However, faces were classified quickly and with few errors as targets, and one might expect that the distractor faces could have been processed just as easily. In fact, as only four different face targets were used in Experiment 3, target-distractor interference might have been obtained even just through picture-based cues than identity and expression processing, as was suggested previously as a potential explanation for Schweinberger et al's (1998, 1999) findings (see p. 41). Moreover, target-distractor interference is generally very robust and has been obtained in letter-letter (e.g. Eriksen & Eriksen, 1974), picture-word (e.g. Smith & Magee, 1980), and face-name interference tasks (Young, Ellis, Flude, McWeeney & Hay, 1986). So, how might the absence of distractor interference in the present experiment have been caused? A possible explanation is that the distractor faces did not interfere with target classification because they were not processed at all. According to the perceptual load theory of selective attention (Lavie, 1995, 2000), the processing of task-relevant and task-irrelevant information proceeds automatically until available processing capacity is exhausted. However, irrelevant information is excluded from processing if task-relevant processing consumes all available processing capacity. If the same principles apply to face processing, then it is possible that a relevant face target could monopolize available resources to the detriment of the distractor faces. This would imply a capacity limit for face processing, such that only a single face can be processed at a time. This possibility is examined in the next chapter.

Chapter 3 Capacity Limits for Face Processing: Face Distractor Interference in Sex and Semantic Classification Tasks

Introduction

In the preceding chapter, the role of attention in facial expression and identity processing was assessed in a response-competition experiment (i.e. experiments concerning distractor interference with target classification; see Experiment 3), in which subjects categorized face targets according to specific expression-identity conjunctions while ignoring task-irrelevant distractor faces. Although the target faces were classified fast and accurately, the same faces did not interfere with face target classification when they were presented as distractors. This could be explained by supposing capacity limits in face processing, such that only a single face (i.e. the target) may be processed at a time. On their own, however, these results provide only a hint at such limits as target and distractor stimuli were constrained to faces images. Under these conditions task-irrelevant faces might have been processed even if they did not act as distractors, or alternatively, the face targets might not have been subject to any distractor interference, from face or nonface stimuli. Therefore, the question addressed in the current chapter is whether responses to a face target can be affected by irrelevant distractor faces under conditions that normally allow for distractor interference.

Thus far, several studies have shown that face distractors are processed reliably with a concurrently presented nonface target. Young, Ellis, Flude, McWeeney & Hay (1986) examined interference effects between simultaneously presented photographs of the faces and the printed names of famous people. Using a semantic categorization task (pop-star vs. politician), participants were required to classify either a face or a name while ignoring the distractor, which could be congruent (i.e. same occupation) or incongruent (different occupation) with the target. Names reliably interfered with the classification of face targets. Moreover, faces also interfered with the classification of name targets. Indeed, faces interfered more with names than names interfered with faces.

Recently, Lavie, Ro & Russell (2003) extended this paradigm to investigate the effect of task-relevant load on irrelevant distractor processing. According to Lavie's perceptual load theory of selective attention (Lavie, 1995, 2000; see Chapter 1), the processing of visual information proceeds automatically until available capacity is exhausted. Therefore, irrelevant information is excluded from processing when task-relevant, attended-to stimuli demand all available capacity. To provide a test for this theory with meaningful stimuli, Lavie et al (2003) measured interference from a flanking distractor upon the classification of a central word or a famous name embedded among several letter strings. Perceptual load of the relevant task was manipulated by varying the number of strings in the interference displays. In accord with the load theory, increasing relevant load eliminated congruency effects from meaningful nonface distractors. Intriguingly though, interference from famous face distractors was entirely unaffected by these load manipulations, leading Lavie et al (2003) to suggest that face processing may proceed automatically (for similar claims see Farah, Wilson, Drain & Tanaka, 1995), independent of target processing.

Comparable conclusions can be drawn from a study by Jenkins, Burton & Ellis (2002) in which an irrelevant famous face distractor showed equivalent repetition priming independent of variations in task-relevant load in a letter-string task, even though explicit memory for the faces was markedly affected by this manipulation. There have also been a number of reports of prosopagnosic patients who show the normal pattern of interference from distractor faces when asked to make semantic classifications of names, despite being explicitly unable to recognise familiar faces (e.g. de Haan, Young & Newcombe, 1987; Sergent & Signoret, 1992b), and these findings have been used extensively to inform theories of covert recognition in prosopagnosia (Young & Burton, 1999). The findings of all these studies suggest that face processing is very robust even across manipulations, which should make it difficult. In Experiment 3 one might have therefore expected the normal pattern of interference, with face categorization times varying as a function of target-distractor congruency. However, unlike Experiment 3 none of these studies examined face processing in multi-face displays, and none imply that face processing is entirely capacity-free. Indeed, Lavie et al (2003) suggest that face processing may be subject to its own capacity limits.

So far, evidence for face processing limits has been rather indirect and has accrued mostly from studies that were not originally motivated by this issue. For example, Palermo & Rhodes (2002) asked subjects to remember a centrally presented target face while matching two flanker faces. Memory for the central face was assessed using a two-alternative recognition test, consisting of either two intact faces, the target and a foil image that differed from the target by one feature (e.g. a pair of eyes), or two exemplars of a particular feature, one of which was extracted from the target. Successfully matching the flanker faces resulted in better memory for intact targets than individual features, but only when the flanker faces were presented inverted. Conversely, matching upright flanker faces eliminated this advantage, suggesting a processing limit for upright, intact faces that is independent of any general processing limits. Using a different technique, Boutet & Chaudhuri (2001) observed perceptual rivalry of two upright overlapping faces, one rotated 45° and the other 45° counterclockwise, whereby only one of the faces could be retrieved at a subsequent recognition test. Two inverted faces, on the other hand, were perceived as an ambiguous combination of both, again suggesting upright face processing limits. Finally, Jenkins, Lavie & Driver (2003) examined dilution of congruency effects in a famous name categorization task. They found that interference from a famous face distractor could be diluted by the presence of another (response-neutral) face, but not by phase-shifted faces, inverted faces, or meaningful nonface objects. In other words, processing of the distractor face seemed to be reduced by competition from an additional face, but not by general competition from different classes of stimuli.

If limits on face processing do apply, then it is possible that a distractor face might not influence responses to a target face, as the resources needed to process the distractor would already be engaged in processing the target. This could provide an explanation for the absence of face-face interference in Experiment 3. Note however, that target-distractor interference is a highly robust effect, which has been demonstrated with various classes of stimulus pairs (e.g. letter-letter, Eriksen & Eriksen, 1974; picture-word, Smith & Magee, 1980; face-name, Young et al, 1986). Given this remarkable generality, the absence of any face-face interference, at least under conditions that normally produce target-distractor interference, seems a somewhat counterintuitive prediction. The present chapter provides a direct test for this prediction, by assessing interference from faces and nonface comparisons over a series of five experiments.

The first experiment examined capacity limits in face processing with unfamiliar faces in a sex classification task. Most previous studies that hint at capacity limits for face processing rely on memory for previously unfamiliar faces (e.g. Boutet & Chaudhuri, 2001; Palermo & Rhodes, 2002). However, memory for unfamiliar faces appears remarkably poor, even over a very short time interval (see e.g. Simons & Levin, 1998), and this may have contributed to previous findings. Sex judgements, on the other hand, can be performed very quickly and without difficulty on unfamiliar faces (e.g. Bruce, Ellis, Gibling & Young, 1987). Moreover, faces usually contain some salient external sex-cues such as hairstyle. These face-related cues might produce interference even when subtler types of facial information, such as expression (as for the expression-identity decisions in Experiment 3), do not. On the other hand, it should be noted that previous studies reporting face-nonface interference used semantic decisions, which require access to facial identity but can be made independent of a person's sex (see e.g. Bruce & Young, 1986; Burton, Bruce & Johnston, 1990; Burton, Bruce & Hancock, 1999). Subsequent experiments in this chapter therefore examined face-face interference with semantic decisions to provide a closer analogue to previous face-nonface interference tasks (e.g. Jenkins et al, 2003; Lavie et al, 2003; Young et al, 1986), and to generalize the findings of Experiment 4 from unfamiliar faces in a sex decision to known faces in semantic tasks.

Experiment 4

In this experiment capacity limits in face processing were assessed with unfamiliar faces in a sex classification task. Subjects were asked to classify stimuli presented at fixation as being male or female. These target stimuli were either unfamiliar faces or printed four-letter forenames, and they were flanked by the distractor images of faces and names. Processing of the distractor was assessed via its congruency effects on target RTs (i.e. same sex vs. different sex). However, in contrast to previous studies, which only examined face-name interference (e.g. Jenkins et al, 2003; Lavie et al, 2003; Young et al, 1986), additional conditions were included to assess within category interference of face and nonface stimuli. In total, congruency effects were measured under four conditions. These conditions involved combining a face target and a face distractor (in the *FACE-face* condition), combining a face target and a name distractor (the *FACE-name* condition), combining a name target and a face distractor (the *NAME-face* condition), and combining two names (the *NAME-name* condition). If face processing is capacity limited, one might expect measurable congruency effects with this paradigm, even in the absence of any within-category interference in the *FACE-face* condition. Alternatively, if several faces can be processed simultaneously, then face distractors should also interfere with the classification of face targets.

Method

Subjects Thirty undergraduate students from the University of Glasgow, whose ages ranged from 19-25 years, participated in the experiment in return for a small payment. All had normal or corrected to normal vision.

Design & Stimuli An Apple Macintosh computer was used to present stimuli and record responses, using PsyScope 1.2.5. Photographs of four unfamiliar female and four unfamiliar male models served as face stimuli. These images were cropped to remove extraneous background, but the outlines of all faces including differences in hairstyle were preserved. In addition, 4 four-letter printed female forenames (Anne, Kate, Lisa & Mary) and 4 male forenames (Hugh, John, Paul & Tony), shown in 36-point Times font, served as name stimuli. All faces were greyscale on a black background and measured 3.6 cm x 4.5 cm (subtending 3.4° x 4.3° of VA at a viewing distance of 60 cm). The names were printed white on black and were between 2.4 cm (the shortest name) and 3.1 cm (the longest name) in width (2.3°–3.0° of VA). These sixteen images were used to construct stimulus displays containing a central target image (face or name), flanked by a distractor image (face or name) that could be congruent (same sex) or incongruent (different sex) with the target (see Figure 3.1 overleaf). The nearest target-distractor contours were 1.25 cm apart (1.2° of VA). Distractors were equally likely to appear on the left or right of the target (this manipulation produced no reliable effects or interactions and is therefore not reported further below).

Pairing each of the 16 target stimuli with each class of distractor (face or name) under each level of congruency (same or different sex) resulted in a total of 64 displays. For displays in which target and distractor were of the same sex (e.g. two male faces, or two male names), stimuli of two different persons were used (see Figure 3.1).

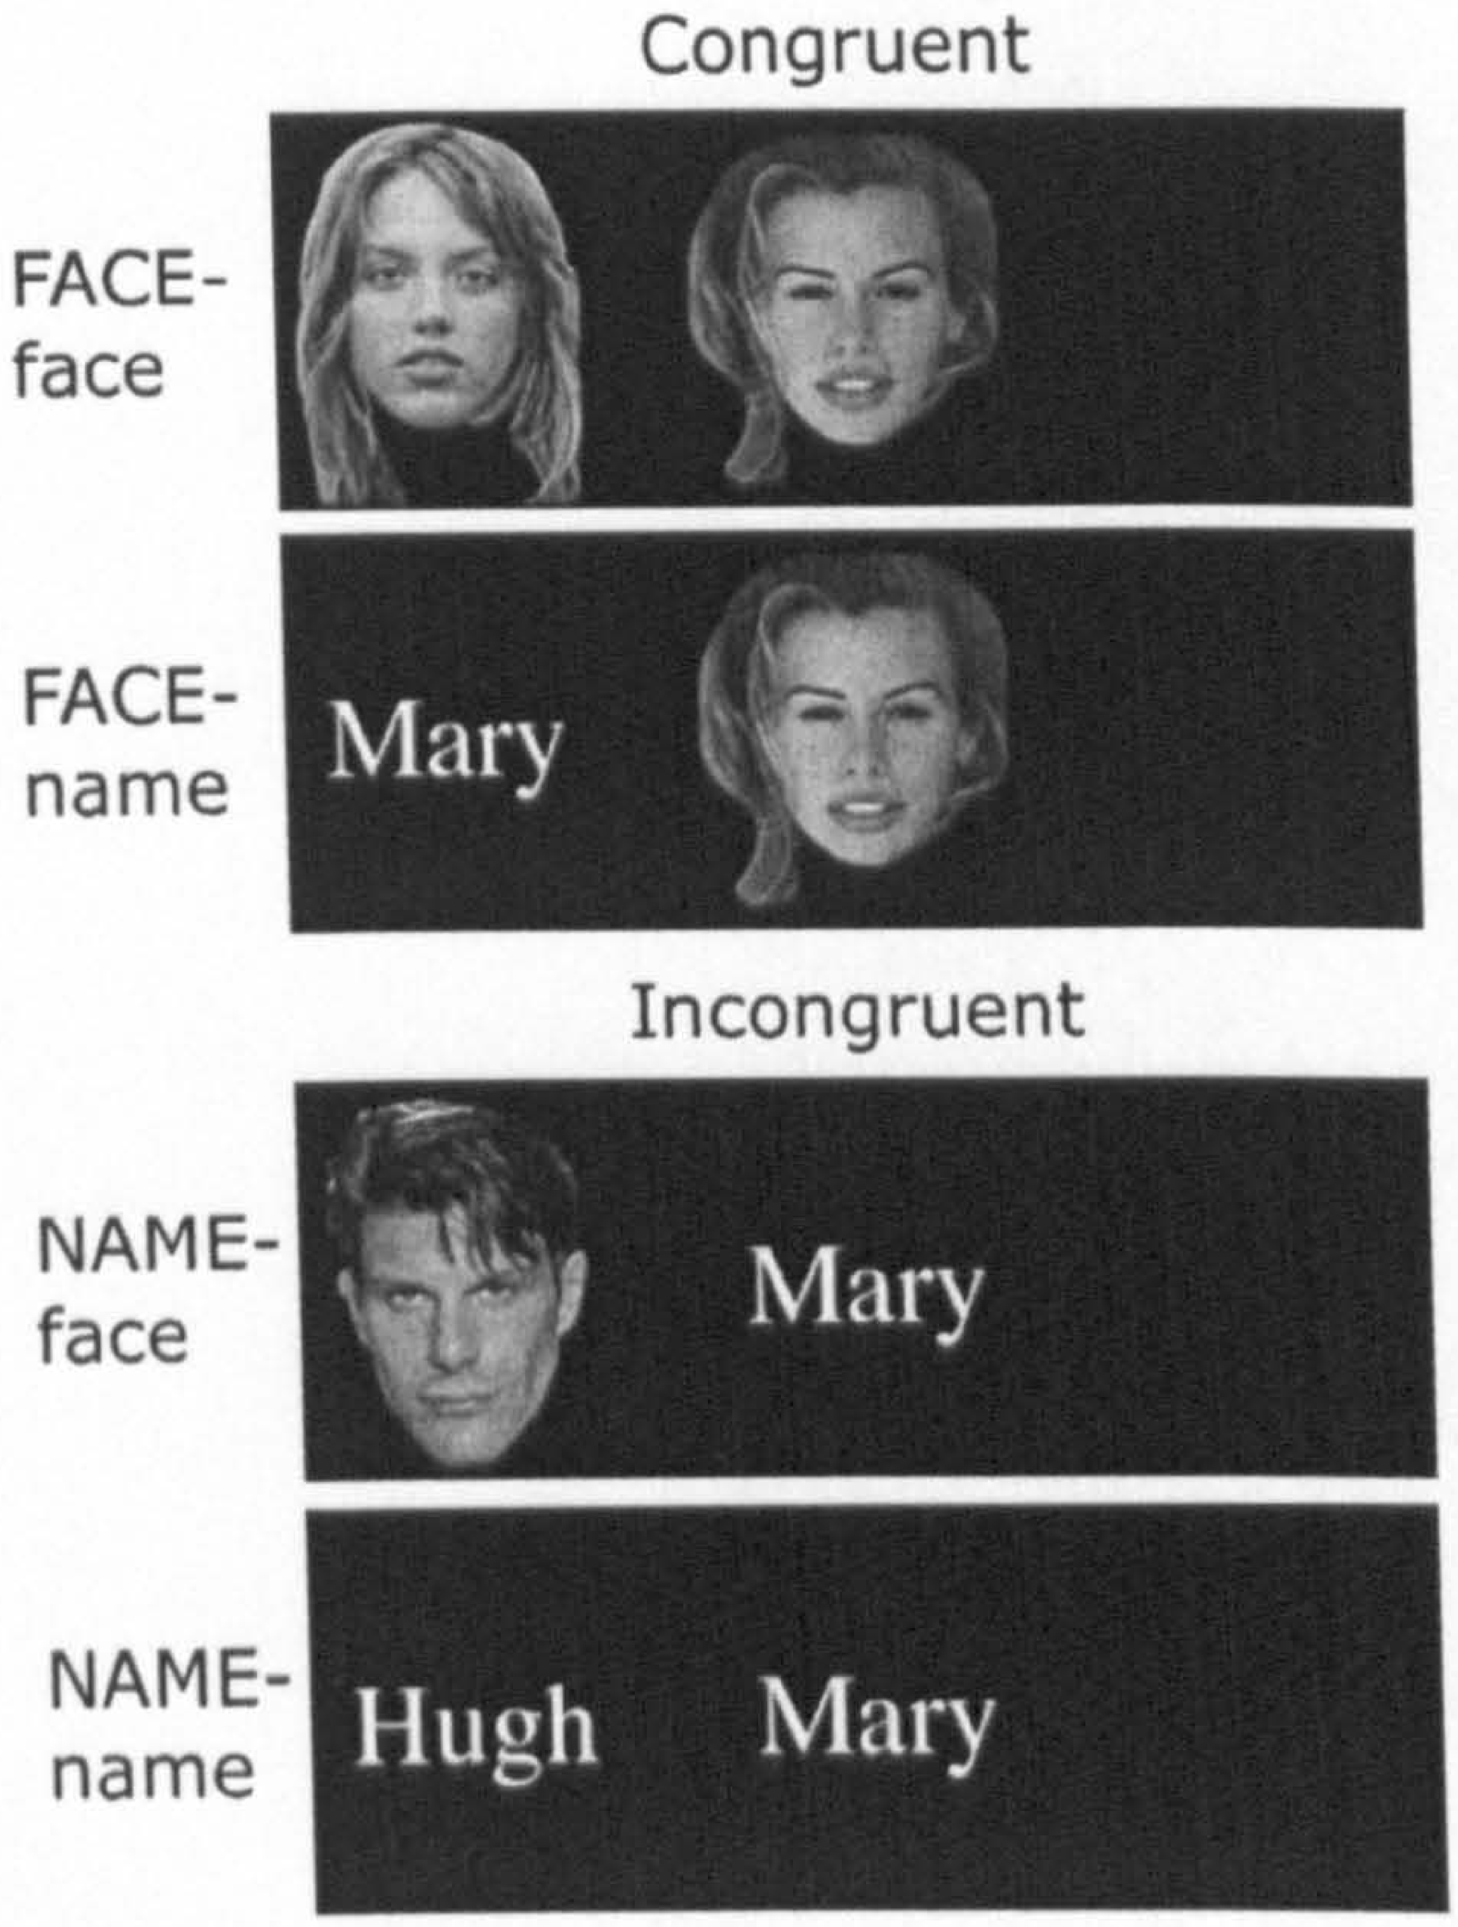

Figure 3.1 Example displays from Experiment 4. The target could be a face or a four-letter forename, and was accompanied by a face or a name distractor, which could be either congruent (same sex) or incongruent (different sex) with the target.

Procedure

Subjects viewed the displays at a distance of 60 cm, which was kept constant by means of a chin-rest. Each trial began with a fixation cross for 750 ms, followed by the target-distractor display for 200 ms (i.e. too briefly to permit a stimulus-responsive saccade to the distractor), and ended with a blank interval until a response was made. Subjects were instructed to classify the target image as a male by pressing the "D" key or as a female by pressing the "L" key on a standard computer keyboard, as quickly and as accurately as possible, while ignoring the distractors. Feedback for errors was given immediately by a short warning tone. Button-press response latencies were measured from stimulus onset. Subjects completed one practice block of 32 trials and 6 experimental blocks of 64 randomly ordered trials, and could take short breaks between blocks.

Results

Figure 3.2 shows the means of the median correct RTs for all conditions. A 2 (face vs. name target) x 2 (face vs. name distractor) x 2 (congruent vs. incongruent) within-subjects ANOVA showed a main effect of congruency, $F(1,29)=23.31$, $p<.01$, with slower responses to incongruent displays, and a main effect of target type, $F(1,29)=32.51$, $p<.01$, with faster responses to face targets than to name targets. No main effect of distractor type was found, $F(1,29)<1$.

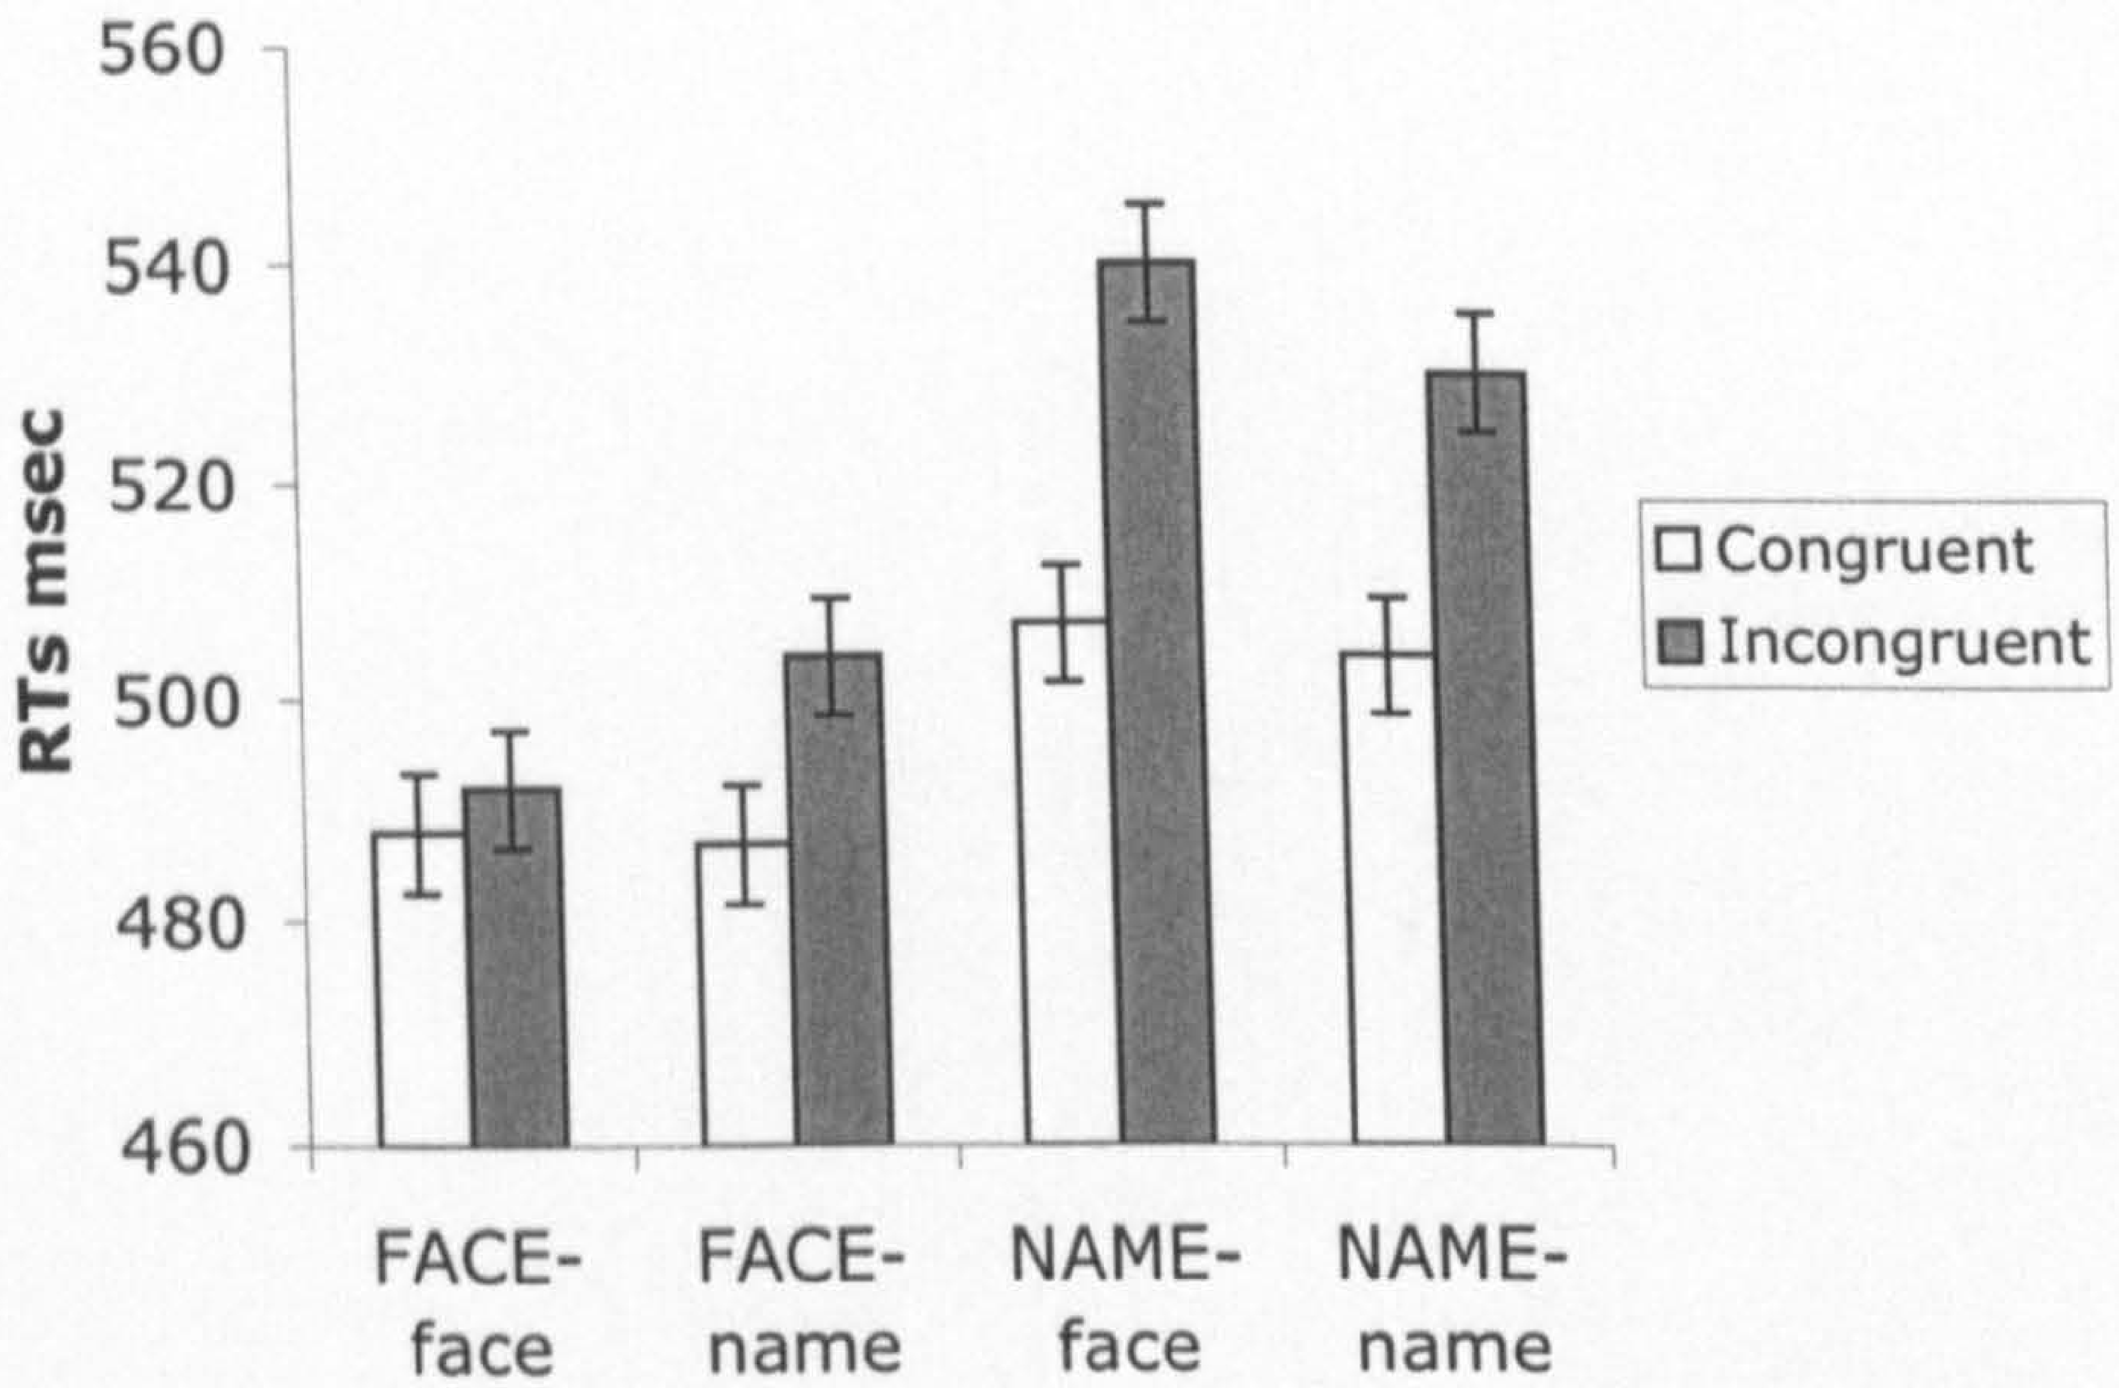

Figure 3.2 Mean reaction times (in msec) across subjects (n=30) as a function of distractor congruency and target-distractor pairings in Experiment 4. Vertical bars represent the standard error of the means.

The effect of target type was modified by an interaction with distractor type, $F(1,29)=6.02$, $p<.05$, an interaction with congruency, $F(1,29)=7.73$, $p<.01$, and a three-way interaction between all three factors, $F(1,29)=5.48$, $p<.05$. As Figure 3.2

suggests, analysis of simple main effects revealed significant congruency effects in the *FACE-name* condition, $F(1,29)=4.52$, $p<.05$, the *NAME-face* condition, $F(1,29)=16.06$, $p<.01$, and the *NAME-name* condition, $F(1,29)=9.62$, $p<.01$. By contrast, there was no effect in the *FACE-face* condition, $F(1,29)<1$.

Error rates were analyzed as the RT data. Incongruent displays resulted in a slight increase in errors in the *NAME-face* condition (incongruent 8.0%, congruent 4.0%) and the *NAME-name* condition (5.7% vs. 5.1%), but no corresponding increase in the *FACE-face* (2.7% vs. 2.8%) or the *FACE-name* conditions (3.8% vs. 3.8%). ANOVA showed a significant main effect of congruency, $F(1,29)=9.91$, $p<.01$, a main effect of target type, $F(1,29)=23.03$, $p<.01$, and an interaction of target type with distractor type, $F(1,29)=6.67$, $p<.05$. In addition, a significant congruency effect was found in the *NAME-face* condition, $F(1,29)=32.56$, $p<.01$. No other comparisons were significant.

Discussion

Similar to Experiment 3, no distractor congruency effects were observed when a face distractor flanked a face target. However, in the current design reliable congruency effects in the NAME-face and the FACE-name condition contrasted this. Moreover, the observed distractor extinction in the *FACE-face* condition does not seem to be a generalized within-category phenomenon, as distractor names also exerted congruency effects onto name targets in the *NAME-name* condition. This pattern of results implies that processing a target face may indeed prevent the processing of a distractor face, and converges with previous findings hinting at capacity limits for face processing (Boutet & Chaudhuri, 2001; Jenkins et al, 2003; Palermo & Rhodes, 2002). However, the absence of face-face interference in a sex classification task is nonetheless surprising, as the face stimuli preserved salient external sex-cues such as hairstyle. On the basis of such cues irrelevant faces could have been classified even without processing actual face information. Thus, these data suggest that the processing of an attended-to face target prevents the processing of all sex-related information from an additional irrelevant face, including even salient external features.

Experiment 5

Experiment 5 was designed to replicate the findings of the previous experiment with familiar faces in a semantic classification task. The retrieval of semantic information, such as occupation and nationality, requires access to a person's identity, which is not necessary for a sex judgement to be made (see e.g. Bruce, 1986; Bruce, Ellis, Gibling & Young, 1987; but see also Rossion, 2002). Consequently, the possibility exists that the results of Experiment 4 will not generalize to tasks that require the identification of familiar faces. Note also that previous studies examining face interference have used semantic judgements, rather than sex or expression decisions, to face-name pairings (Jenkins et al, 2003; Lavie et al, 2003; Young et al, 1986). Experiment 5 thus provides a closer analogue to existing designs that have produced face distractor interference.

Method

Subjects Thirty undergraduate students from the University of Glasgow, whose ages ranged from 19-25 years, participated in the experiment in return for a small payment. All had normal or corrected to normal vision.

Stimuli & Procedure The procedure was identical to Experiment 4, except that the decision to be made was whether the targets were pop-stars or politicians. The surnames and faces of four male pop-stars (Kurt Cobain, Eminem, Michael Jackson & Elvis Presley) and four male politicians (George Bush, Bill Clinton, Colin Powell & Donald Rumsfeld) served as stimuli. The faces were manipulated as in Experiment 4. The surnames were shown in 18-point Arial font, measuring between 1.7 cm and 2.9 cm in width (1.6° - 2.8° of VA). As before, these images were used to construct 64 stimulus displays containing a central face or name target, flanked by a face or a name distractor, which could be either congruent or incongruent (same or different occupation) with the target (see Figure 3.3).

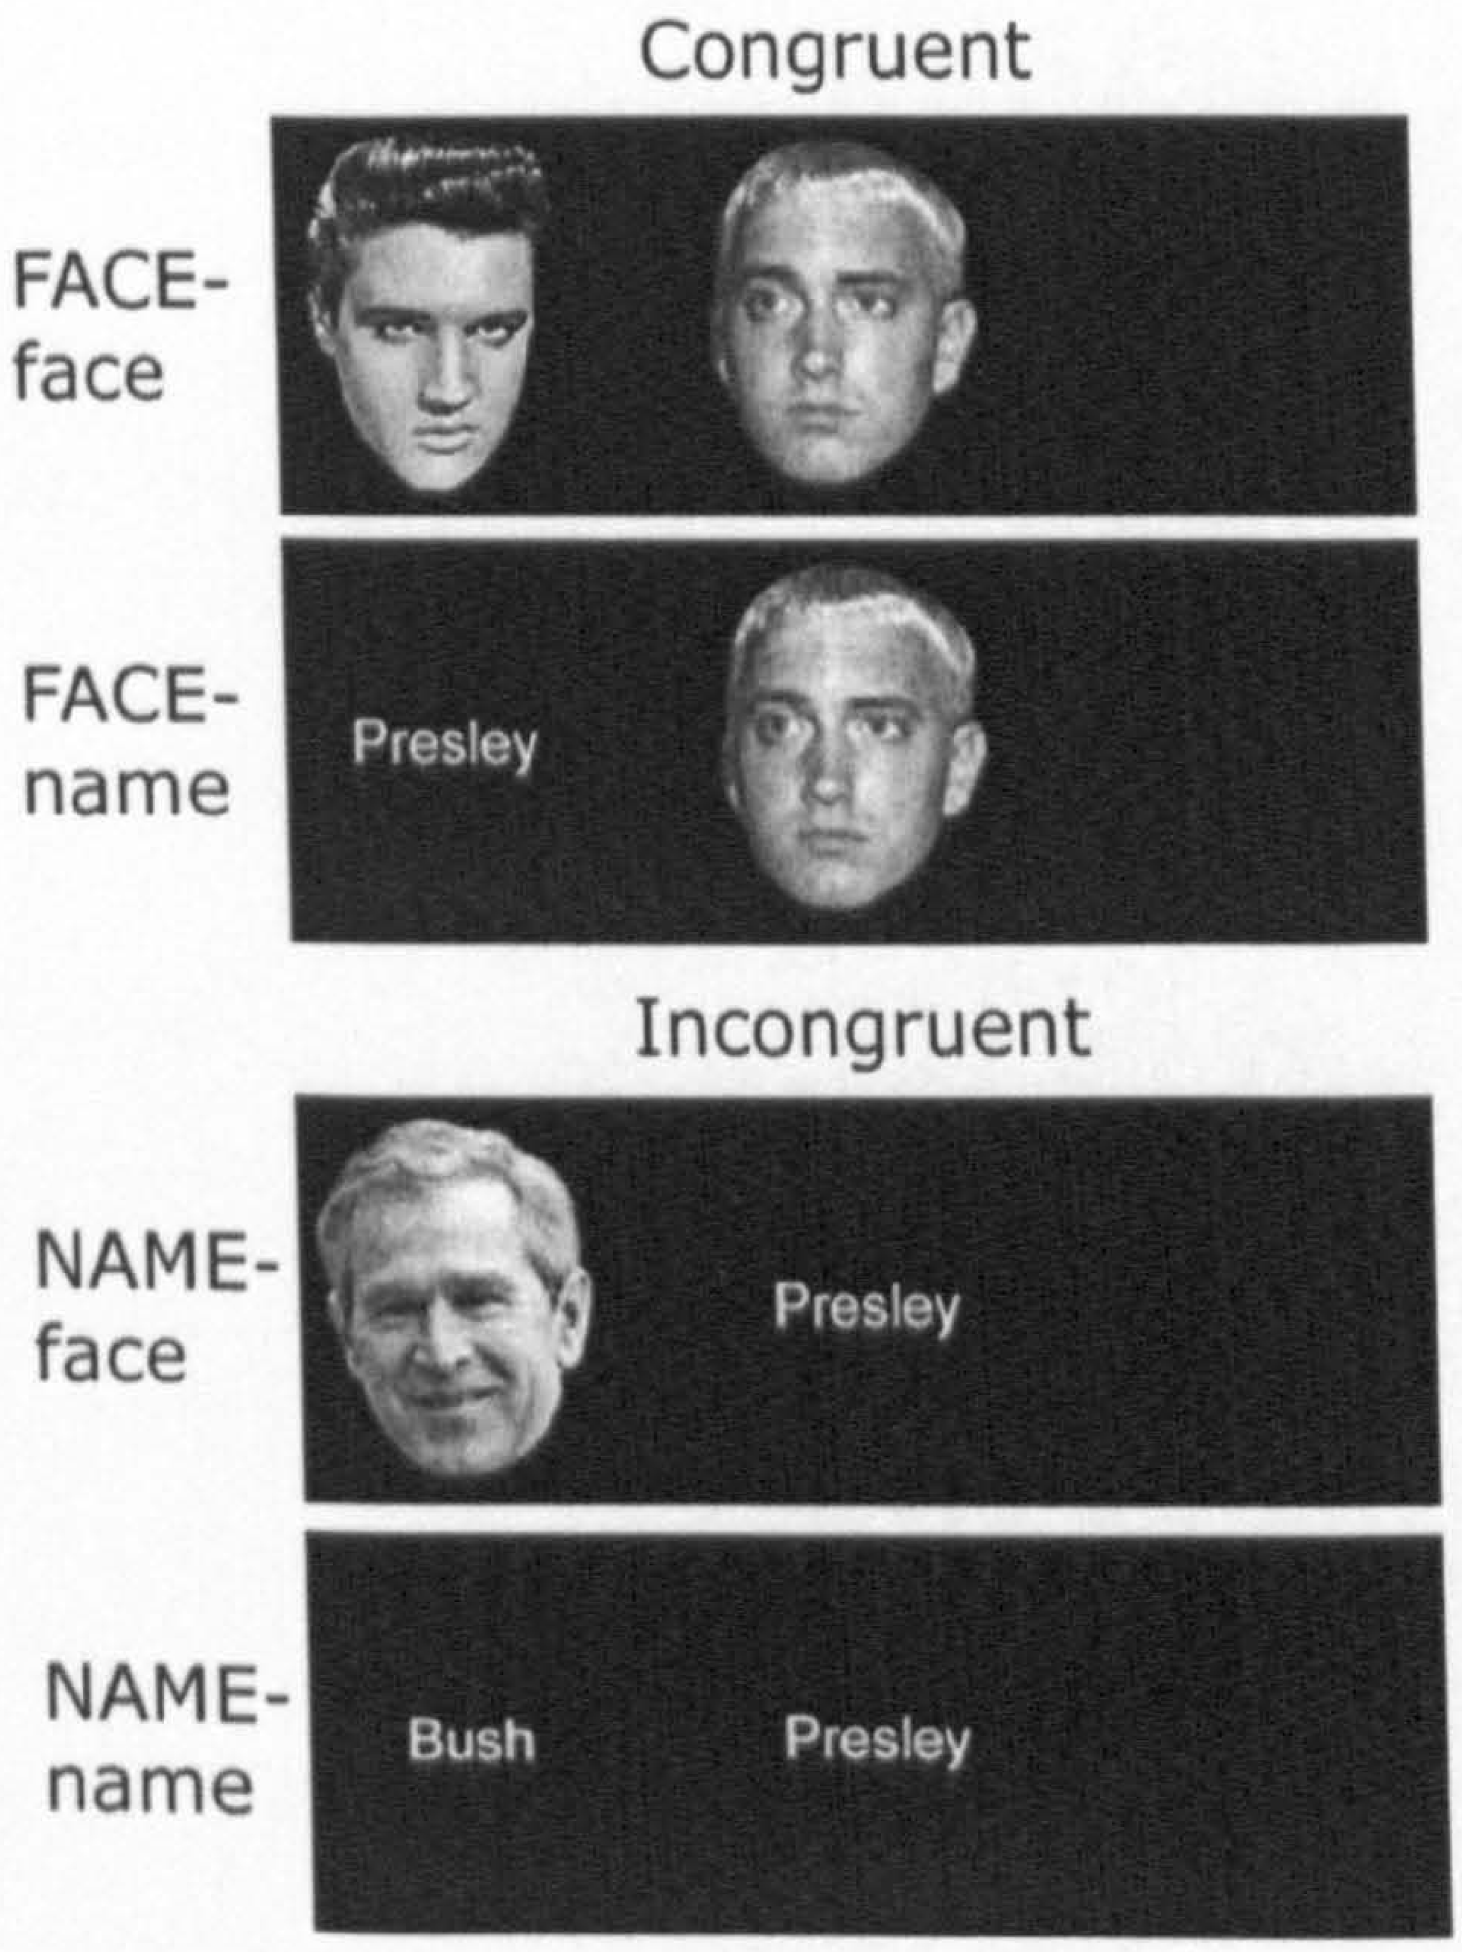

Figure 3.3 Example displays from Experiment 5. The target could be a famous face or a famous name, and was accompanied by a face or name distractor, which could be either congruent (same occupation) or incongruent (different occupation) with the target.

Results

Figure 3.4 shows the means of the median correct RTs for all conditions. As before, a 2 (face vs. name target) x 2 (face vs. name distractor) x 2 (congruent vs. incongruent) within-subjects ANOVA showed a main effect of congruency, $F(1,29)=42.36$, $p<.01$, with slower responses to incongruent versus congruent displays, and a main effect of target type, $F(1,29)=26.02$, $p<.01$, with faster responses to face targets.

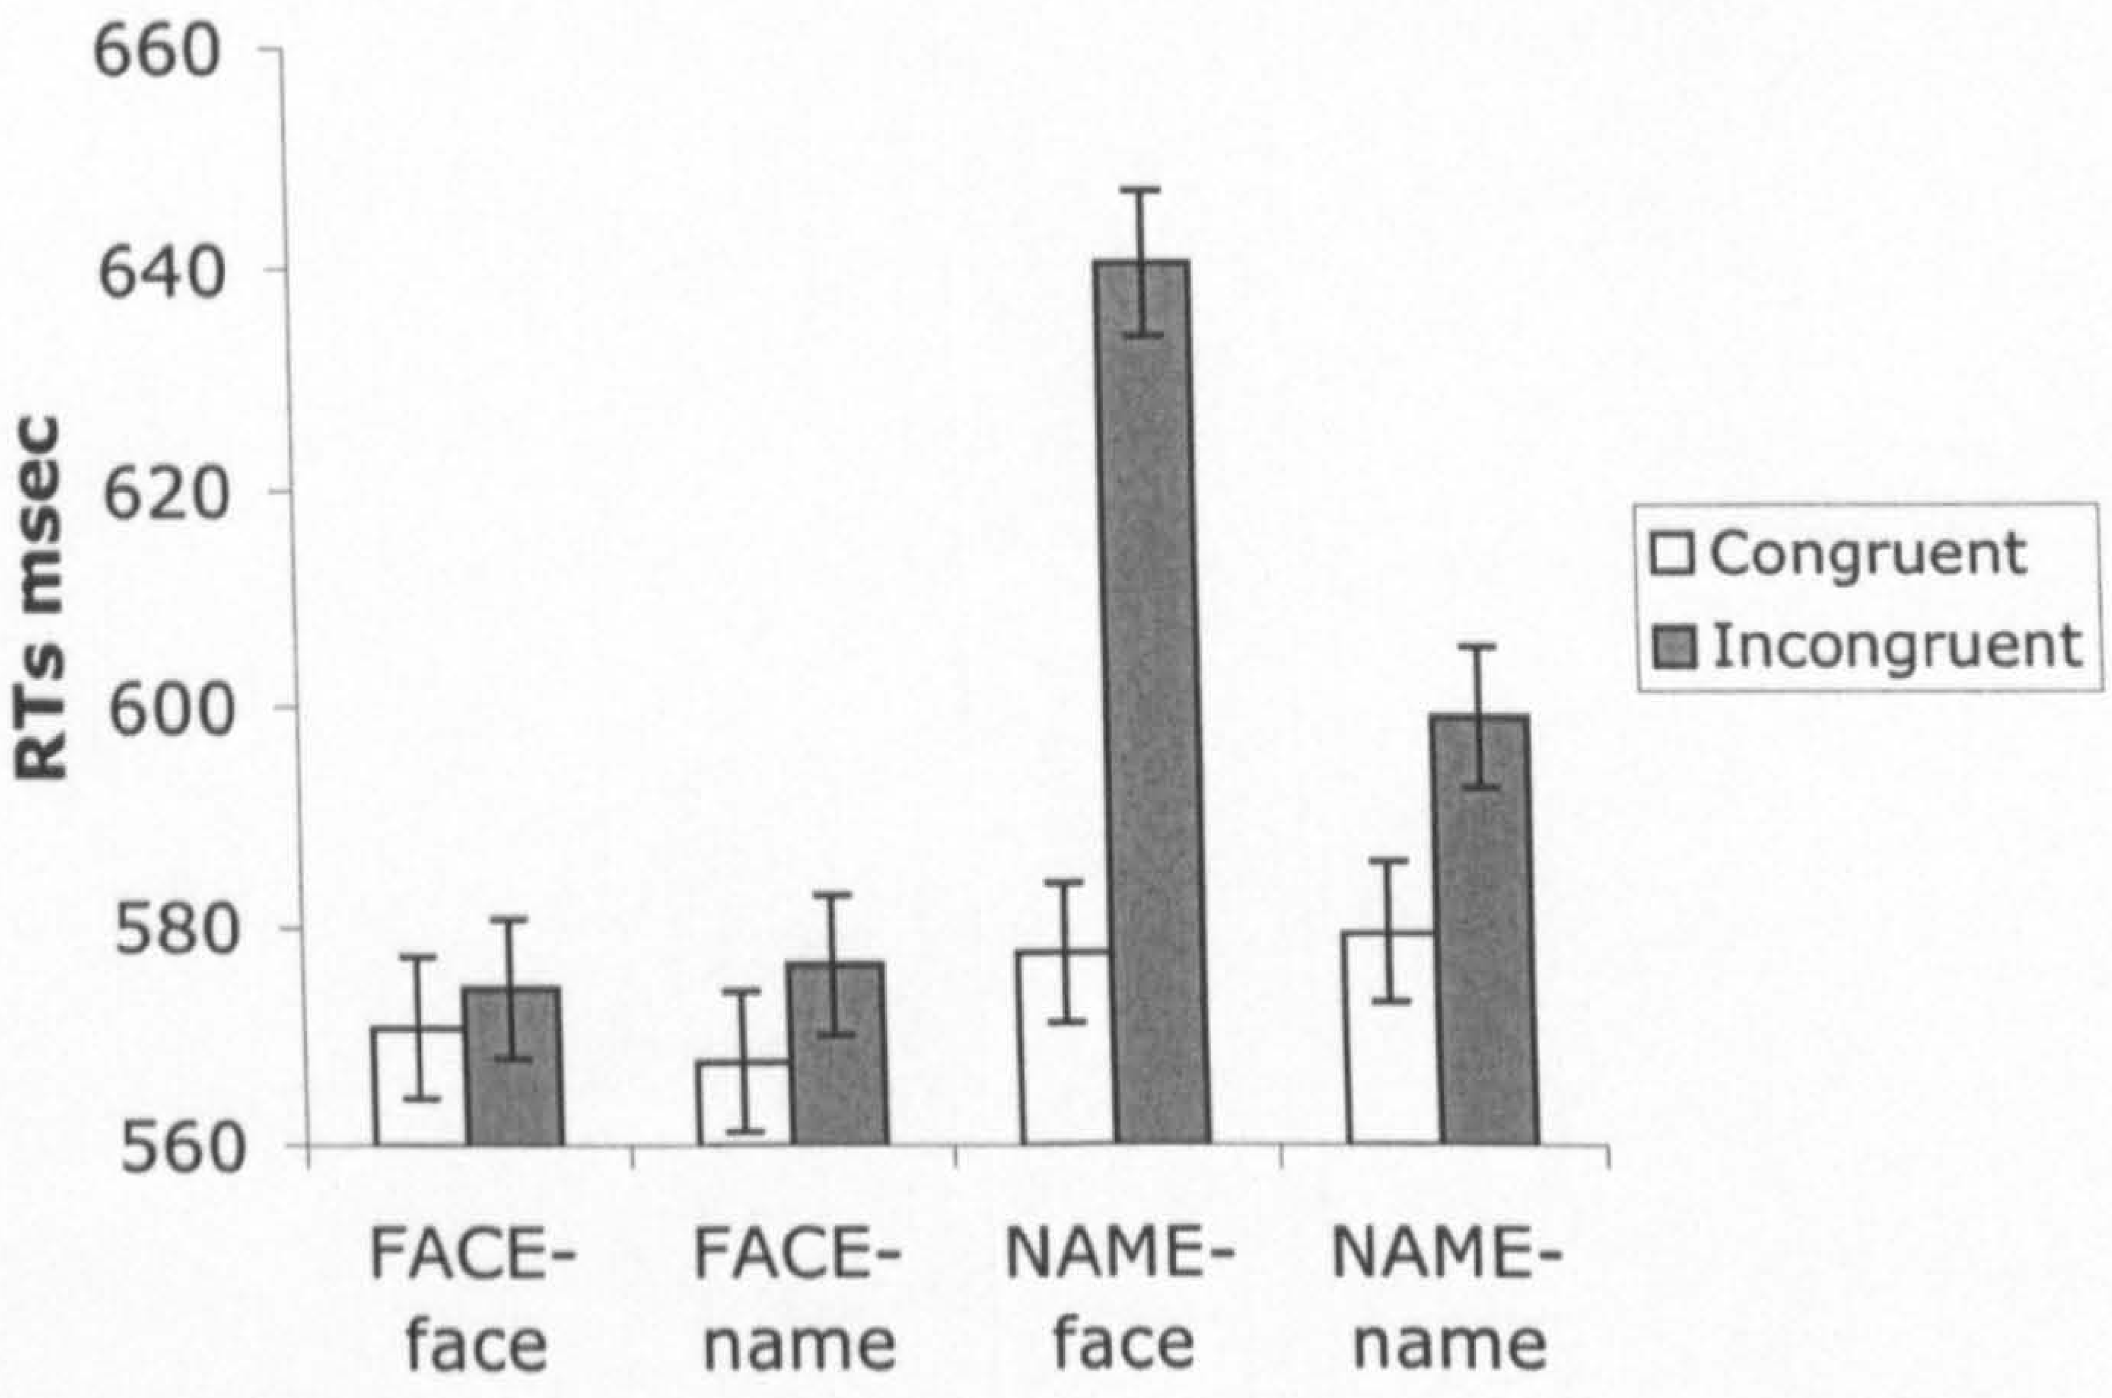

Figure 3.4 Mean reaction times (in msec) across subjects (n=30) as a function of distractor congruency and target-distractor pairings in Experiment 5. Vertical bars represent the standard error of the means.

In addition, a main effect of distractor type was found, $F(1,29)=4.89$, $p<.05$, reflecting slower responses to displays containing face distractors than displays with name distractors. These main effects were modified by two-way interactions between each of the factors [target type x distractor type, $F(1,29)=13.96$, $p<.01$; target type x congruency, $F(1,29)=33.47$, $p<.01$; and distractor type x congruency, F(1,29)=8.09, p<.01], and a three-way interaction between all factors, F(1,29)=17.53, p<.01. Simple main effect analysis revealed significant congruency effects in the *NAME-face* condition, F(1,29)=74.32, p<.01, and the *NAME-name* condition, F(1,29)=7.19, p<.05. However, there were no congruency effects in the *FACE-face*, F(1,29)<1, or the *FACE-name* condition, F(1,29)=1.47.

An analogous analysis of the error rates was carried out. Incongruent displays resulted in an increase in errors in the *FACE-name* condition (incongruent 5.2%, congruent 3.3%), the *NAME-face* (13.3% vs. 6.0%), and the *NAME-name* condition (6.5% vs. 5.6%). However, no corresponding increase was observed in the *FACE-face* condition (3.9% vs. 3.7%). ANOVA revealed main effects of congruency, F(1,29)=33.39, p<.01, target type, F(1,29)=28.01, p<.01, and distractor type, F(1,29)=7.01, p<.05. As for the RTs, there were also interactions between each of the factors [target type x distractor type, F(1,29)=20.41, p<.01; target type x congruency, F(1,29)=15.27, p<.01; and distractor type x congruency, F(1,29)=6.22, p<.05], and a three-way interaction between all factors, F(1,29)=22.28, p<.01. Significant congruency effects were found in the *FACE-name*, F(1,29)=4.45, p<.05, and *NAME-face* conditions, F(1,29)=67.22, p<.01. No other comparisons were significant.

Discussion

This experiment replicates some of the important aspects of Experiment 4 with a semantic decision, which, unlike a sex decision, requires the identification of the face stimuli. As before, the *FACE-face* condition failed to yield a congruency effect. This was contrasted by a reliable congruency effect in the *NAME-face* condition, which indicates that irrelevant faces can nonetheless act as distractors when a semantic task is used. However, unlike the sex classification task of Experiment 4, the RT data failed to yield a reliable interference effect when a face target was flanked by a name distractor (a 9 ms trend in this direction did not approach significance). While the RTs failed to show reliable distractor interference in the *FACE-name* condition, a significant congruency effect in error rates was found. This alone, however, does not support parallel processing of face target and name distractor. Alternatively, it might represent attentional shifts to the distractor locations, which may have enhanced distractor processing to the detriment of accurate target classification.

Note that previous studies also obtained less, albeit significant, interference from name distractors during face classification than from face distractors during name classification (Young et al, 1986). In contrast to Young et al (1986), who presented target *and* distractor centrally, the distractors always appeared in the periphery in this experiment, clearly separated from the target. Although this arrangement was designed to avoid target-distractor confusion, numerous studies have shown that interference can be significantly reduced by increasing spatial separation between a target and a distractor (e.g. Gatti & Egeth, 1978; Merikle & Gorewich, 1979; Hagenaar & Van der Heijden, 1986). Therefore, a possible explanation for the absence of any *FACE-name* interference could to some extent lie in the spatial arrangement of the target-distractor pairings. Nonetheless, the absence of reliable distractor interference in the *FACE-name* condition is potentially problematic, as it raises the possibility that the extinction of *FACE-face* interference does not reflect capacity limits in face processing, but rather that the famous face targets may not have been subject to any distractor interference in the present task. This is explored more thoroughly in the next experiment.

Experiment 6

The purpose of Experiment 6 was two-fold. The first aim was to replicate the interference pattern of the within-category conditions of Experiment 5, and specifically to produce distractor congruency effects in both between-category conditions (i.e. the NONFACE-face and the FACE-nonface conditions). As was discussed in the preceding experiment, the comparison between these latter conditions and the *FACE-face* condition is vital in establishing capacity limits for face processing. The second aim was to examine whether nonface stimuli other than names are subject to interference within this paradigm. To provide an analogue to the semantic task of Experiment 5, images of national flags were used as nonface comparisons and subjects were asked to classify the face and flag targets as being American or British.

Method

Subjects Twenty undergraduate students from the University of Glasgow, whose ages ranged from 18-24 years, participated in the experiment in return for a small payment. All had normal or corrected to normal vision.

Stimuli & Procedure The procedure was the same as for Experiment 5, except as follows. Subjects were now instructed to classify the targets as American or British. Three different images each of Tony Blair (British Prime Minister), George Bush (American President), the Union Jack (British flag), and the Stars and Stripes (American flag) served as stimuli. The faces and flags were cropped to rectangular shapes to produce a closer resemblance but this resulted in the loss of some external features, such as stimulus outline. Faces and flags were then converted to greyscale and sized to 2.2 cm x 3.0 cm (2.1° x 3.9° of VA). These images were used to construct the stimulus displays as in Experiment 5 (see Figure 3.5). Pairing each of the 12 stimuli with each class of distractor (face or flag) under each level of congruency (same or different nationality) resulted in 48 displays. Subjects completed a practice and 8 experimental blocks of 48 randomized trials.

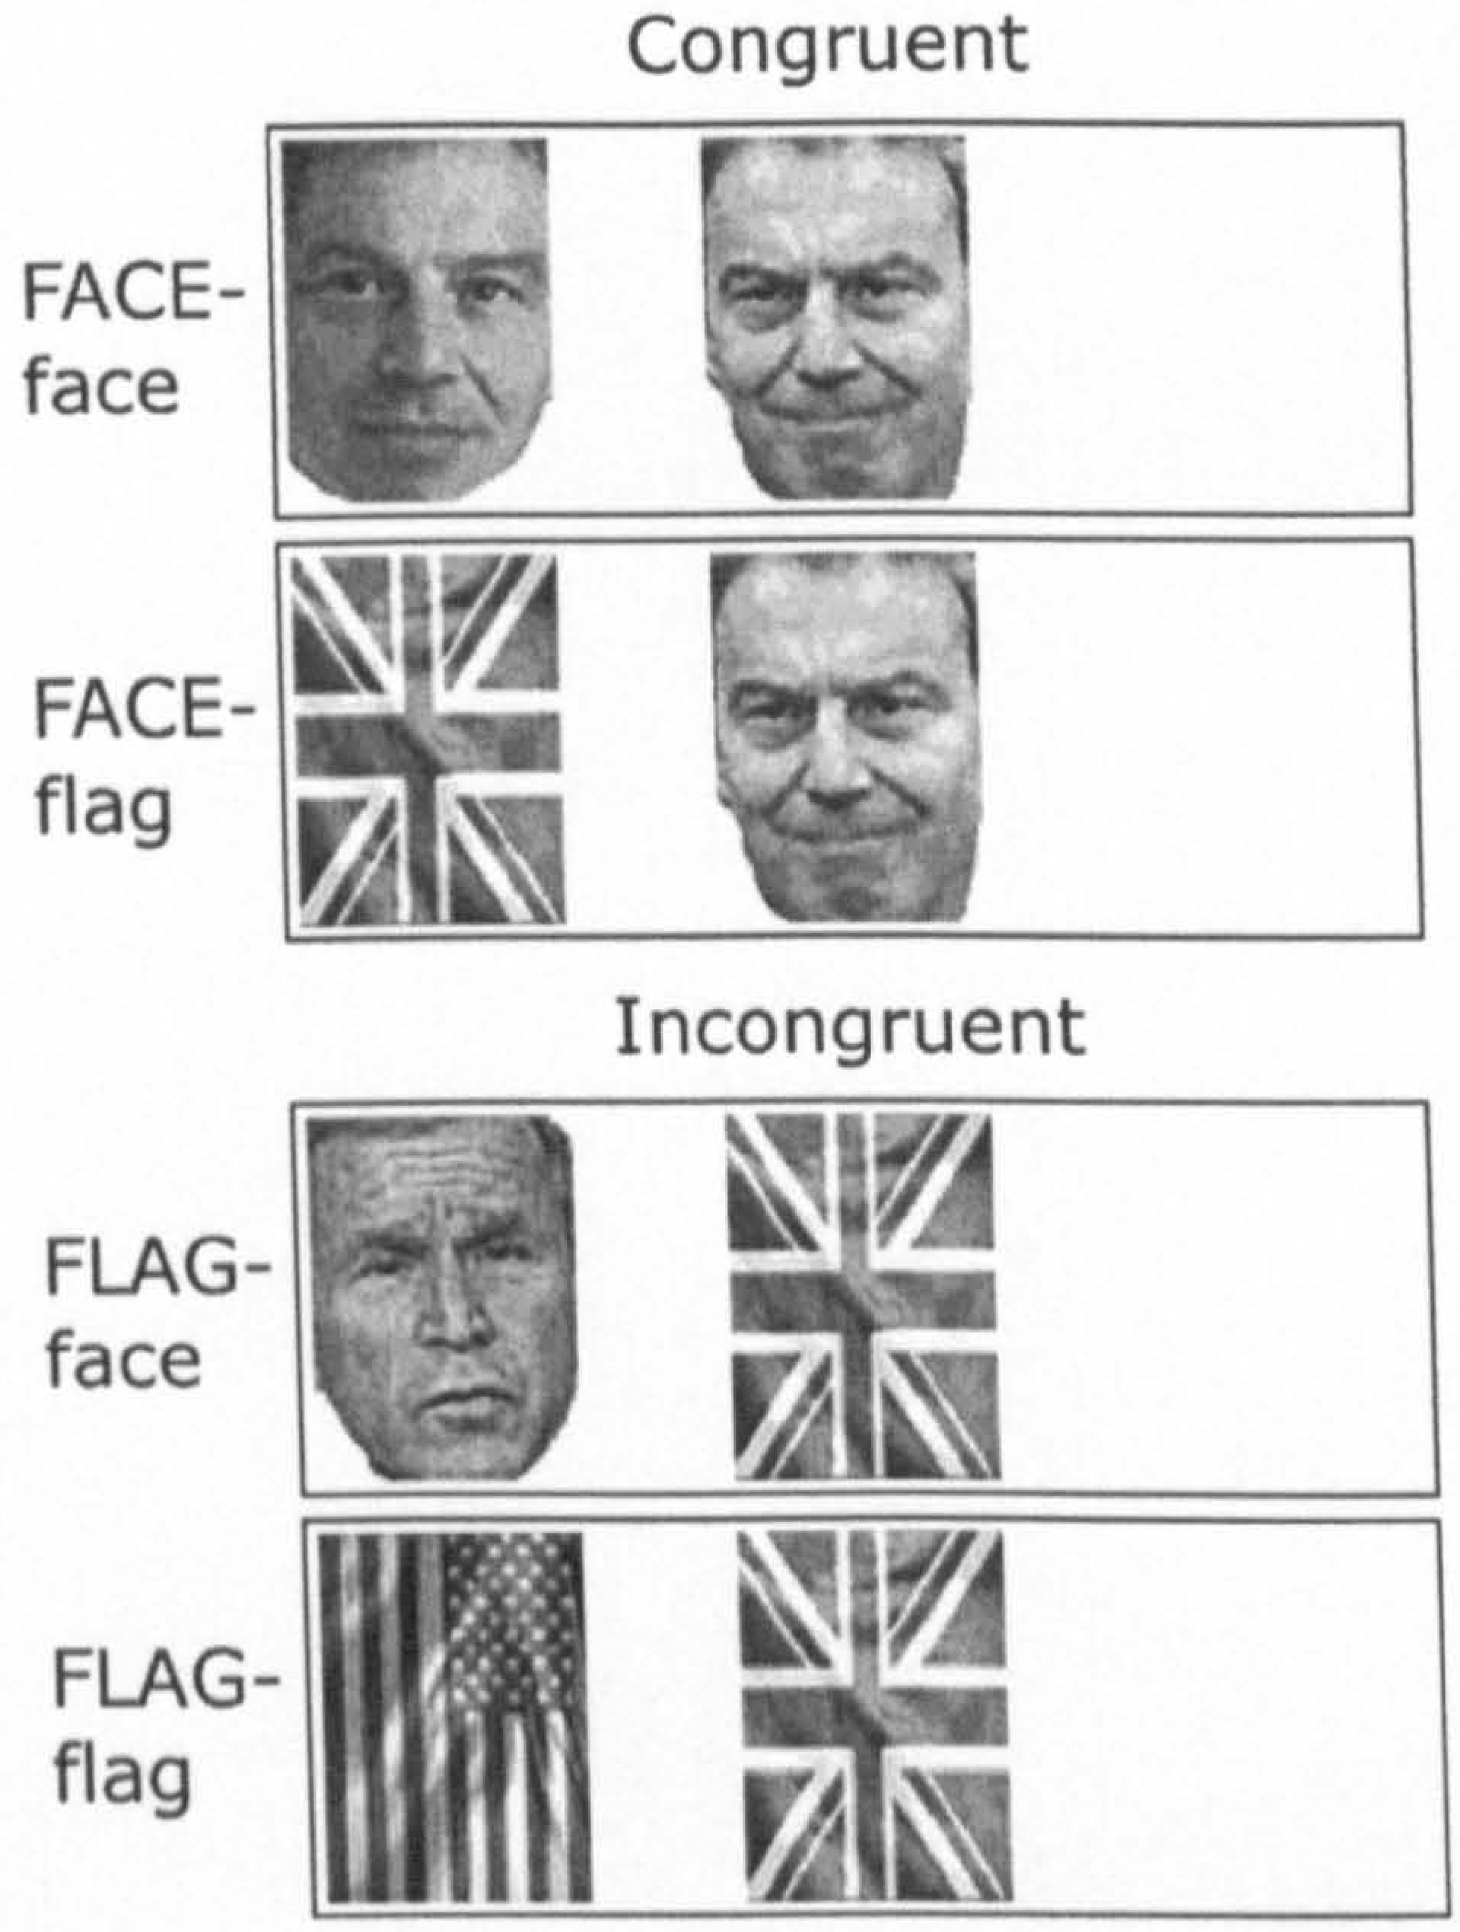

Figure 3.5 Example displays from Experiment 6. The target could be a famous face or a national flag, and was accompanied by a face or a flag distractor, which could be either congruent (same nationality) or incongruent (different nationality) with the target.

Results

As for Experiment 4, the means of the median correct RTs were calculated for all conditions and are shown in Figure 3.6. A 2 (face vs. flag target) x 2 (face vs. Flag distractor) x 2 (congruent vs. incongruent) within-subjects ANOVA showed a main effect of congruency, $F(1,19)=21.29$, $p<.01$, with slower RTs to incongruent displays, but no main effect of target type, $F(1,19)<1$, or distractor type, F(1,19)=1.09. The effect of congruency was modified by an interaction with target type, F(1,19)=6.94, p<.05. As Figure 3.6 suggests, significant congruency effects were found in the *FACE-flag* condition, F(1,19)=5.21, p<.05, the *FLAG-face* condition, F(1,19)=27.71, p<.01, and the *FLAG-flag* condition, F(1,19)=4.82, p<.05. By contrast, there was no effect in the *FACE-face* condition, F(1,19)<1.

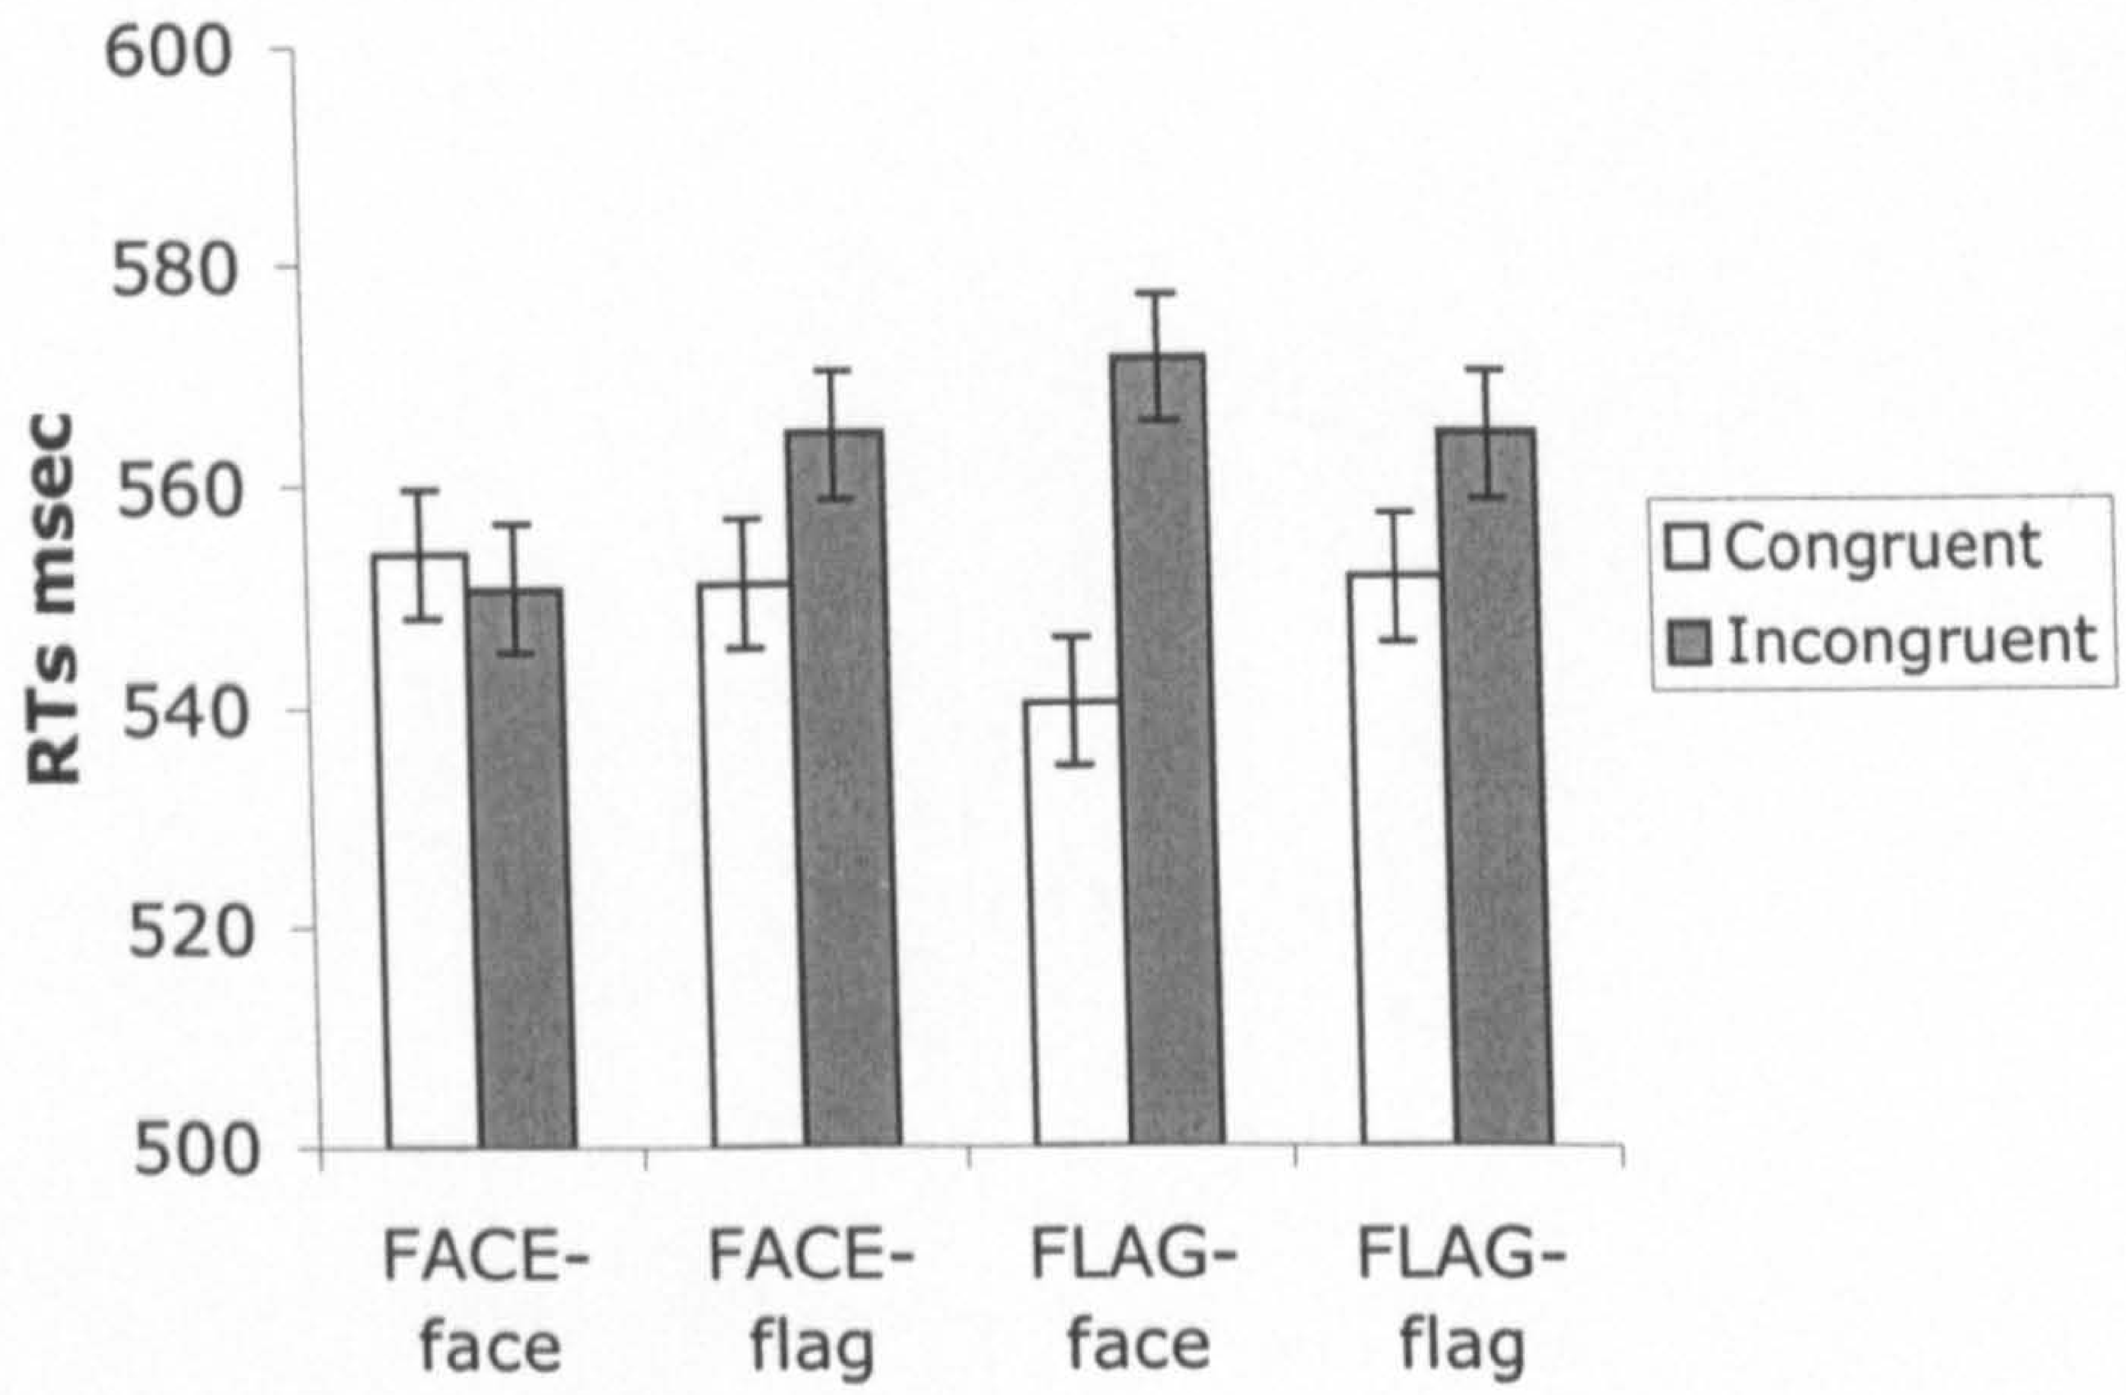

Figure 3.6 Mean reaction times (in msec) across subjects (n=20) as a function of distractor congruency and target-distractor pairings in Experiment 6. Vertical bars represent the standard error of the means.

Error rates mirrored the RT data. Incongruent displays resulted in an increase in errors in the *FACE-flag* condition (incongruent 7.0%, congruent 5.1%), the *FLAG-face* condition (8.5% vs. 5.1%), and the *FLAG-flag* condition (6.9% vs. 4.0%), but no corresponding increase in the *FACE-face* condition (4.4% vs. 4.6%). ANOVA showed a significant main effect of congruency, F(1,19)=4.60, p<.05, and an interaction between target type and distractor type, F(1,19)=6.70, p<.05. No other comparisons were significant.

Discussion

The results show an intriguing pattern. As in Experiment 5, no evidence of distractor processing was found in the *FACE-face* condition. In contrast, however, the present data shows that face targets can be subject to congruency effects in a semantic classification task, as exerted here by the flag distractors. Faces also functioned as distractors. In fact, the largest congruency effect was observed again in the condition in which a nonface target was flanked by a distractor face. In addition, within-category congruency effects were observed for images of flags, suggesting that they are not subject to analogous capacity limits to faces. Thus, these results replicate the pattern that was observed in Experiment 4, and of Experiment 5 with the addition of a reliable congruency effect in the FACE-nonface condition, and extend those findings to images of flags in a nationality task.

Experiment 7

Experiments 4-6 provide converging support for the hypothesis that face processing may be capacity limited, such that a face distractor does not influence target face processing. In nonface paradigms, however, target-distractor interference is seemingly boosted by presenting several congruent (or incongruent) distractors (e.g. letter-letter, Eriksen & Hoffman, 1973). Therefore, to provide a stronger test of the claim that face distractors do not influence target face processing, the number of distractors was increased to four in this experiment, thus increasing fourfold the total amount of congruent and incongruent information in each display. If multiple faces can be processed simultaneously, one might expect this manipulation to boost any influence of the distractors. This might lead to measurable congruency effects even in the *FACE-face* condition, where none were previously found. On the other hand, if target face processing is unaffected by adding further distractors, this would provide additional support for the face processing limits that were observed in the preceding experiments.

Method

Subjects Twenty-two undergraduate students from the University of Glasgow, whose ages ranged from 18-26 years, participated in the experiment in return for a small payment. All had normal or corrected to normal vision.

Stimuli & Procedure These were the same as in Experiment 6, except for the following changes. The former single distractors were replaced by four distractors, positioned around the central target to form a "+" configuration (see Figure 3.7 overleaf). The nearest distractor contours were approximately 1.0 cm (1.0° of VA) horizontally, and 0.9 cm (0.9° of VA) vertically from the target. Twenty celebrities' faces (10 British, 10 American, see Appendix A) and 20 flags (10 British, 10 American) were used as stimuli. In each flanker display, all four distractors were of the same nationality (e.g. four American faces). To avoid confounding semantic information during target classification, the faces were drawn from five occupational categories (pop-star, politician, sports-star, comedian, movie-star), so that no occupation occurred more than once in any face display. Faces were presented with their external features (i.e. hair, face outline) and the flags were cropped to elliptical shapes in order to produce a closer resemblance between the flag and face outlines (see Figure 3.7). Faces and objects measured between 2.1-2.4 cm horizontally and 2.5-3.2 cm vertically (2.0°-2.3° x 2.4°-3.1° of VA).

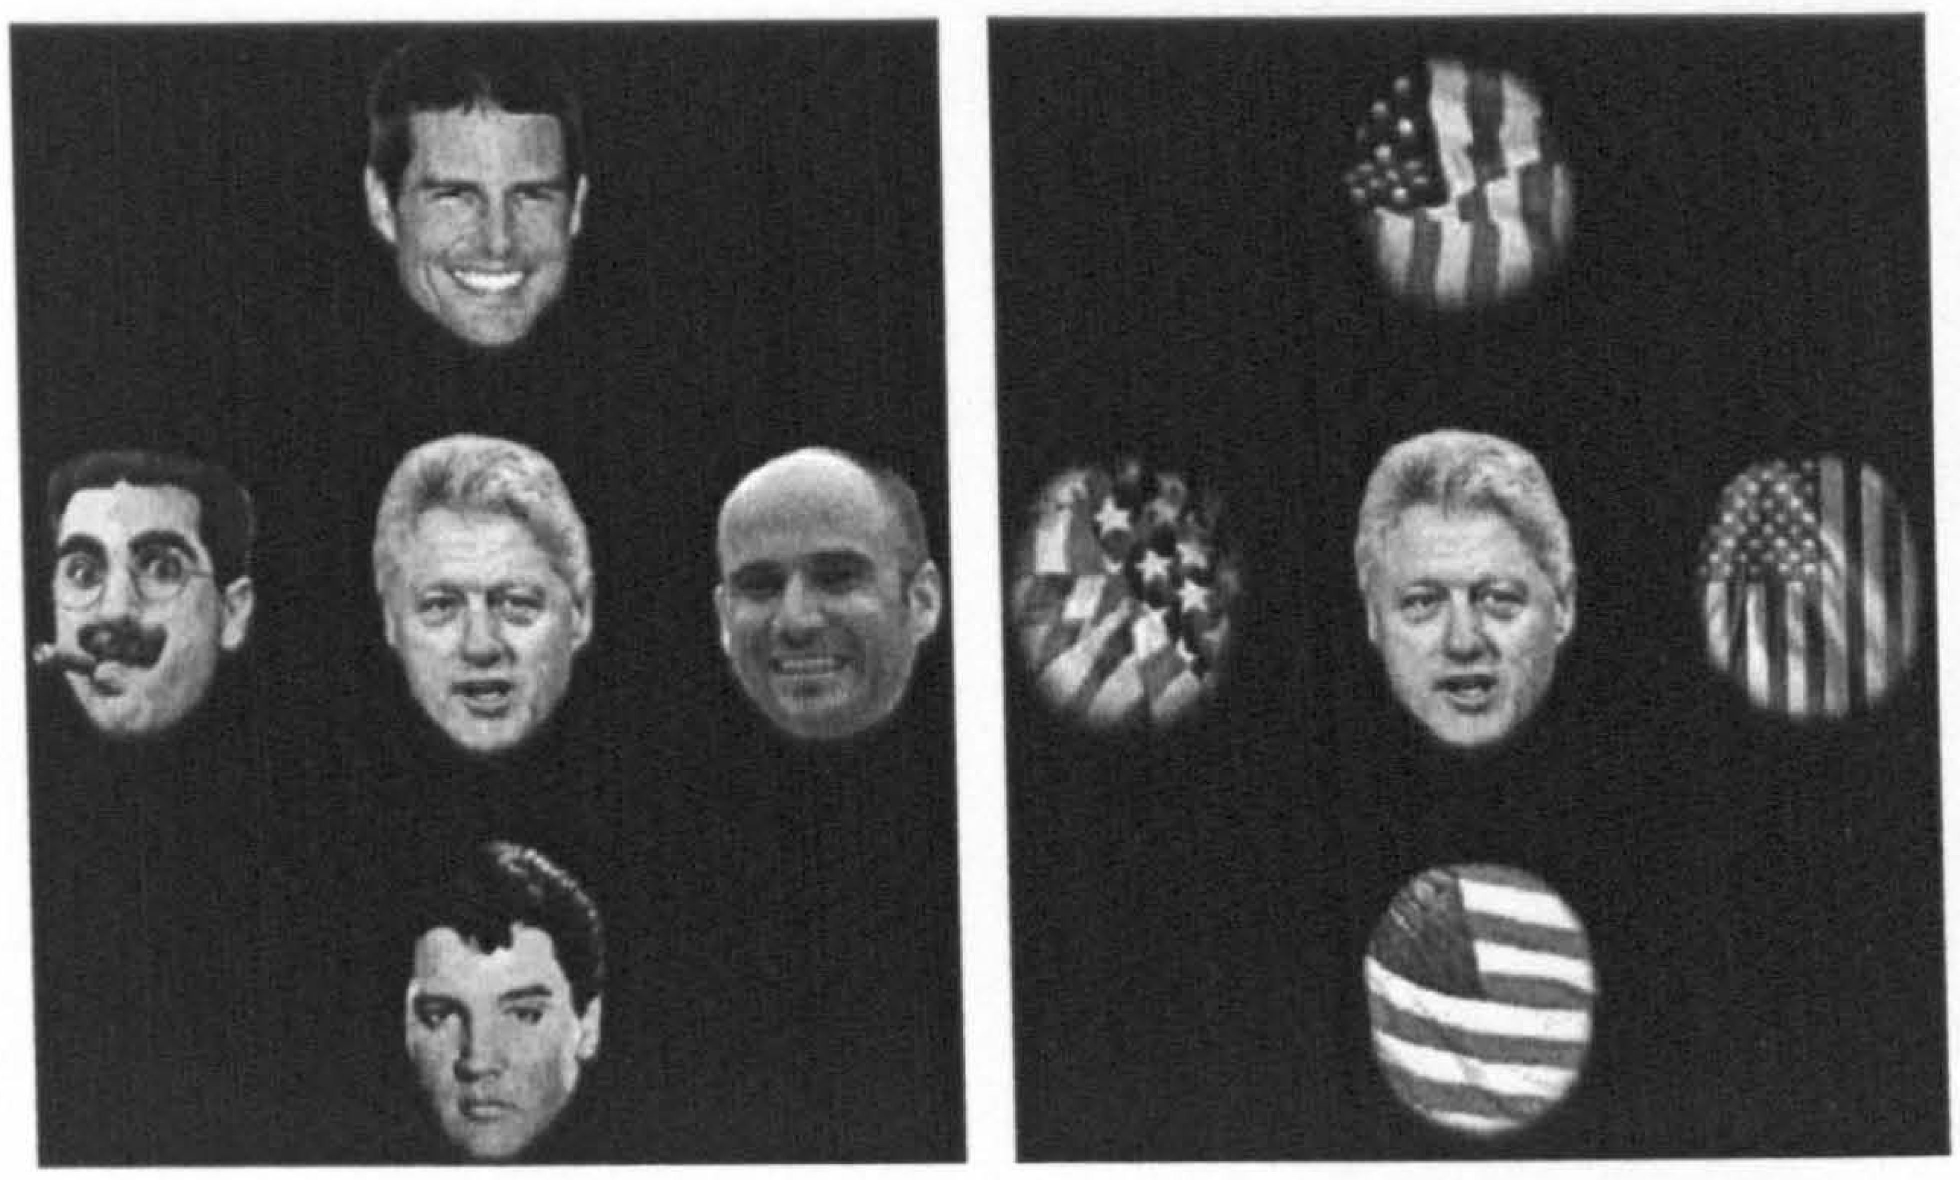

Figure 3.7 Example displays from the *FACE-face* condition (left display) and the *FACE-flag* condition (right display) in Experiment 7. Clockwise from top, the face distractors here are: Tom Cruise, Andre Agassi, Elvis Presley, and Groucho Marx. Target face: Bill Clinton.

Combining each of the 40 targets with congruent and incongruent distractors, under two levels of distractor type, resulted in a total of 160 stimuli. Each subject completed a practice block of 40 trials, followed by eight experimental blocks of 80 trials. Therefore, over the eight experimental blocks each stimulus display was encountered a total of four times. Each condition was equally likely to occur in each block and trial order was randomized in all blocks.

Results

Figure 3.8 shows the means of the median correct RTs for all conditions. A 2 (face vs. flag targets) x 2 (face vs. flag distractors) x 2 (congruent vs. incongruent) within-subjects ANOVA of the RT data revealed a significant main effect of congruency, $F(1,21)=26.00$, $p<.01$, with slower responses to incongruent displays, and a main effect of target type, $F(1,21)=56.39$, $p<.01$, with faster responses to flag targets. These effects were modified by interactions between target type and congruency, F(1,21)=4.87, p<.05, and between distractor type and congruency, F(1,21)=8.77, p<.01. As Figure 3.8 suggests, significant congruency effects were found in the *FACE-flag* condition, F(1,21)=30.54, p<.01, and the *FLAG-flag* conditions, F(1,21)=9.35, p<.01, but not in the *FLAG-face* condition, F(1,21)<1, or the *FACE-face* condition, F(1,21)=1.72.

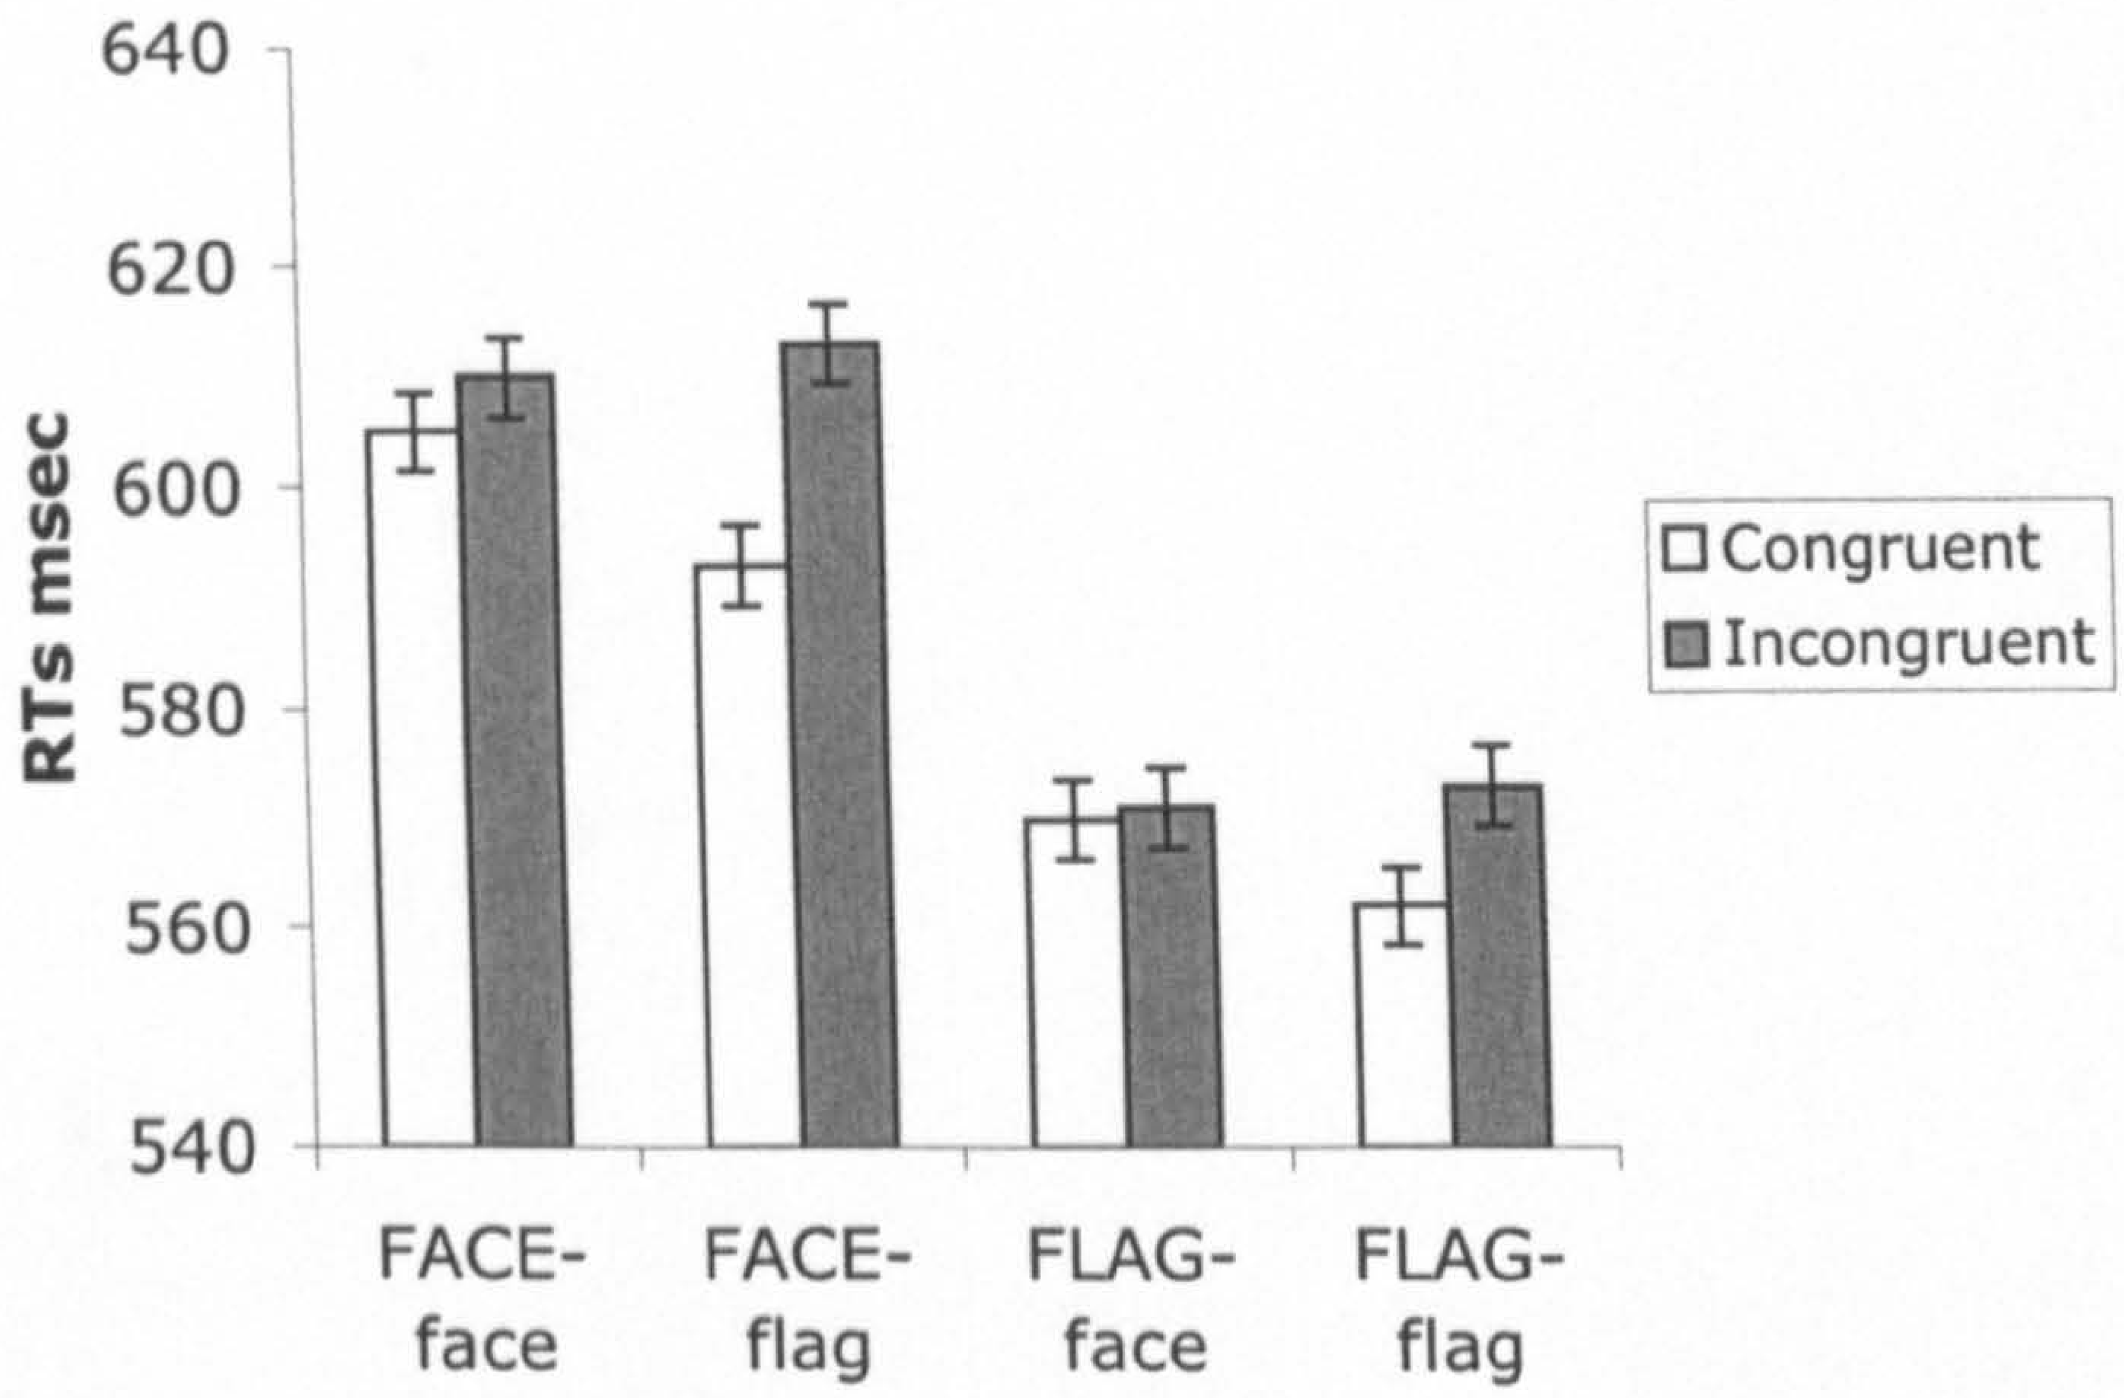

Figure 3.8 Mean reaction times (in msec) across subjects (n=22) as a function of distractor congruency and target-distractor pairings in Experiment 7. Vertical bars represent the standard error of the means.

Error rates followed a similar pattern. Incongruent displays showed increased errors in the *FACE-flag* (incongruent 7.8%, congruent 5.2%), and the *FLAG-flag* conditions (5.5% vs. 3.8%), but not in the *FACE-face* (5.8% vs. 6.4%) or the *FLAG-face* conditions (5.9% vs. 6.6%). Main effects of target type, F(1,21)=3.99, p=.06, and distractor type, F(1,21)=3.78, p=.07, were not statistically reliable but approached significance. In addition, a target type x distractor type interaction, F(1,21)=7.18, p<.05, and a distractor type x congruency interaction, F(1,21)=8.32, p<.01, were found. However, only one congruency effect, in the *FACE-flag* condition, reached significance, $F(1,21)=6.57$, $p<.05$. No other comparisons were significant.

Discussion

To provide a stronger test for face processing limits, the number of task-irrelevant distractor stimuli was increased to four in this study. However, despite a fourfold increase in the amount of potentially distracting information, distractor faces were still unable to influence responses to target faces. In fact, multiple face distractors even failed to produce congruency effects onto the nonface targets, images of flags. This diverges from Experiments 4-6, in which a solitary face distractor interfered strongly with nonface comparison targets. Indeed, the largest congruency effects were observed in the *NAME-face* and the *FLAG-face* conditions in those experiments. Contrary to multiple face distractors, however, and analogous to the solitary distractor flags of Experiment 6, multiple flag distractors exerted congruency effects upon flag targets and face targets alike.

The absence of interference from face distractors onto nonface targets under conditions in which multiple flag distractors produce congruency effects is perhaps surprising, in particular as the processing of just a single of the four distractor faces could have been used to produce the same, strong face-nonface interference of previous experiments. A possible explanation for this finding is that multiple faces may compete for limited processing resources (see e.g. Vuilleumier, 2000; Ro, Russell & Lavie, 2001, for such claims), whereby competition may remain unresolved between several equally task-irrelevant competitors. This suggestion will be discussed in detail at the end of this chapter. Before then, however, it is worth considering other potential explanations for these results.

Experiment 7 used a substantially larger number of stimuli than the preceding experiments. It is therefore possible that the particular stimuli chosen were not sufficiently highly associated with the response category (nationality) to produce interference. However, the face images consisted of well-known celebrities and as targets these faces were classified quickly according to their nationality and with few errors. Moreover, although target RTs were slightly faster to flags than to faces in this experiment, this overall difference cannot explain the complete absence of interference from multiple distractor faces, since multiple flag distractors exerted congruency effects upon both fast flag targets and slow face targets. Indeed, Young et al (1986) also found that slow name targets interfere with fast face distractors during semantic classification. Of course, whenever one is using nonface and face stimuli within the same experiment, it is always possible that effects can be explained in terms of visual complexity or other physical attributes. Although this is always possible, the results of Experiments 4-6, specifically the fact that solitary face distractors interfered more with nonface classification than nonface comparisons interfered with both face and nonface classification, provides at least suggestive evidence that these results might require explanation in terms of competition from multiple face distractors.

Experiment 8

Experiments 4-7 demonstrate that irrelevant face distractors do not interfere with face target classification when two or more faces are presented simultaneously. This indicates that face processing may be capacity limited under these conditions, such that no more than a single face can be processed at a time. This experiment raises a different question than the preceding experiments. A few studies have shown that faces can be recognized after very short exposure durations of less than 70 ms (Ellis, Young & Koenken, 1993; Morrison, Bruce & Burton, 2000), which is considerably less than the 200 ms display time of the present interference paradigms. This opens the possibility that the extinguished (non-interfering) face distractors of Experiments 4-7 may have been processed alongside the face targets, although too late to affect response times, even if they were not processed in parallel. However, despite previous studies reporting face recognition thresholds of less than 70 ms, participants were given ample time to provide a face identification response in the shape of a name or some unambiguous semantic information (Ellis et al, 1993; Morrison et al, 2000). Moreover, Bentin, Deouell & Soroker (1999) showed that ERP markers sensitive to face familiarity occur usually 250 to 500 ms after stimulus onset. Thus, face processing continues considerably beyond the acquisition of a face stimulus from the visual field. This raises the question whether the ongoing processing of a face stimulus is sufficient to extinguish distractor face interference in a subsequent display, or whether face processing limits are only observed in situations in which two faces are displayed simultaneously.

The present study examined this by manipulating the temporal conditions under which face targets and distractor faces were presented. To this end, a cue consisting of either a famous face or a flag was displayed at fixation, followed by a famous name target in the same spatial location and a flanking famous face distractor. As in the preceding experiments, face distractor processing was assessed via its congruency effects on (name) target classification times. However, name classification was contingent upon cue type, such that subjects were instructed to respond to targets that were preceded by a British stimulus (e.g. the face of Prince Charles or a picture of the Union Jack), but to withhold a response following non-British stimuli. Importantly, cues were presented either for a short duration (67 ms) or a long duration (500 ms), followed by a 50 ms inter-stimulus interval (ISI), followed by a face-name interference display of 200 ms. Thus, face cues and face distractors were presented in the same spatial arrangement as targets and distractors in Experiments 4-6 and both relevant face and irrelevant face processing was assessed, but the temporal relationship between these stimuli was changed. If cue and distractor faces can be processed in quick succession, then a face distractor should interfere with a name target following an immediately preceding, briefly presented face cue. Alternatively, distractor processing may be extinguished by the ongoing processing of a briefly presented face cue, relative to when it is displayed for longer.

Method

Subjects Twenty-seven undergraduate students from the University of Glasgow, whose ages ranged from 20-25 years, participated in the experiment in return for a small payment. All had normal or corrected to normal vision.

Design & Stimuli An Apple Macintosh computer equipped with PsyScope 1.2.5 software presented the stimuli and recorded responses. For the cue displays, photographs of four famous British people (Richard Branson, Prince Charles, Lawrence Lwellelyn-Bohen & Jamie Oliver) and of four non-British people (Woody Allen, Jean Reno, Arnold Schwarzenegger & Jerry Springer) served as face stimuli, and images of four Union Jacks (British flag) and the national flags of four other countries (Greece, South Africa, Switzerland & the USA) served as nonface comparisons. Faces and flags were presented in greyscale on a black background, and with their outline intact, and measured maximally 3.6 cm x 4.5 cm (subtending 3.4° x 4.3° of VA). For the interference displays, the full names and faces of six male politicians (Tony Blair, George Bush, Bill Clinton, John Major, Colin Powell & Jack Straw) and of six male pop-stars (Kurt Cobain, Gareth Gates, Mick Jagger, Elton John, Elvis Presley & Justin Timberlake) were used. Names acted as central targets and were shown white on black in 18-point Arial font, measuring between 2.8 cm (the shortest name) and 4.9 cm (the longest name) in width (2.7°-4.7° of VA). The faces served as flanking distractors to the left or right of the name target, and were presented in greyscale at a size of 3.6 cm x 4.5 cm (subtending 3.4° x 4.3° of VA) on a black background. Target-distractor distance varied between 0.6 cm (0.6° of VA, for the longest name target) and 1.6 cm (1.6° of VA, for the shortest name target).

Overall, there were two main types of trials: no-go trials, for which the cue was always a non-British face or flag, and go trials, for which the cues consisted of British faces or flags. This distinction was included to ensure that participants were processing the cue stimuli, for which no direct response requirement was incorporated. In addition to the go/no-go distinction, the cues and the name-face displays were combined in four conditions under each level of target-distractor congruency (i.e. name-face: same vs. different occupation). These conditions involved presenting a face for 67 ms followed by a name-face display (in the *Short Face* condition), presenting a face for 500 ms followed by a name-face display (the *Long Face* condition), or presenting a flag for 67 ms followed by a name-face display (the *Short Flag* condition), and a flag for 500 ms followed by a name-face display (the *Long Flag* condition). Thus, the interference phase was identical across all conditions, which varied only in cue type (face vs. flag) and SOA (500 ms vs. 67 ms).

The specific stimulus pairings for all conditions were created online during the experiment via a PsyScope Factor Table, which was programmed so that each possible combination was equally likely to occur for each participant. The only exceptions were no-go trials, which were only half as likely to occur as go trials. Due to the use of a Factor table for trial generation, for some interference displays the name target and the face distractor were of the same person. The probability of this event was 1/24. Such trials may produce larger congruency effects, than trials in which the name and face of different persons are paired under congruent conditions (see Young et al, 1986). In contrast to Experiments 4-6, however, in which targets and distractors were never taken from the same person to avoid perceptual matching of targets and distractors in the within-category conditions, such strategies cannot account for name-face interference.

Procedure Each trial began with a fixation cross for 1000 ms, followed by a cue item. This was displayed at fixation for 67 ms or 500 ms, and was replaced by a blank ISI for 50 ms, which was in turn replaced by a name-face display for 200 ms, and a final blank interval until a response was registered (see Figure 3.9 overleaf). Subjects were instructed to classify the name target as belonging to a pop-star or a politician as quickly and as accurately as possible while ignoring the face distractors, but only provided that it was preceded by a British face or a British flag. For non-British cues, subjects were instructed to press the space bar following the presentation of the interference displays (to initiate the next trial). A short warning tone again gave feedback for errors. Button-press latencies were measured from the onset of the name-face displays. Subjects completed one practice block of 36 trials and three experimental blocks of 120 randomly ordered trials. Subjects were given breaks between blocks, and initiated each block by pressing the space bar.

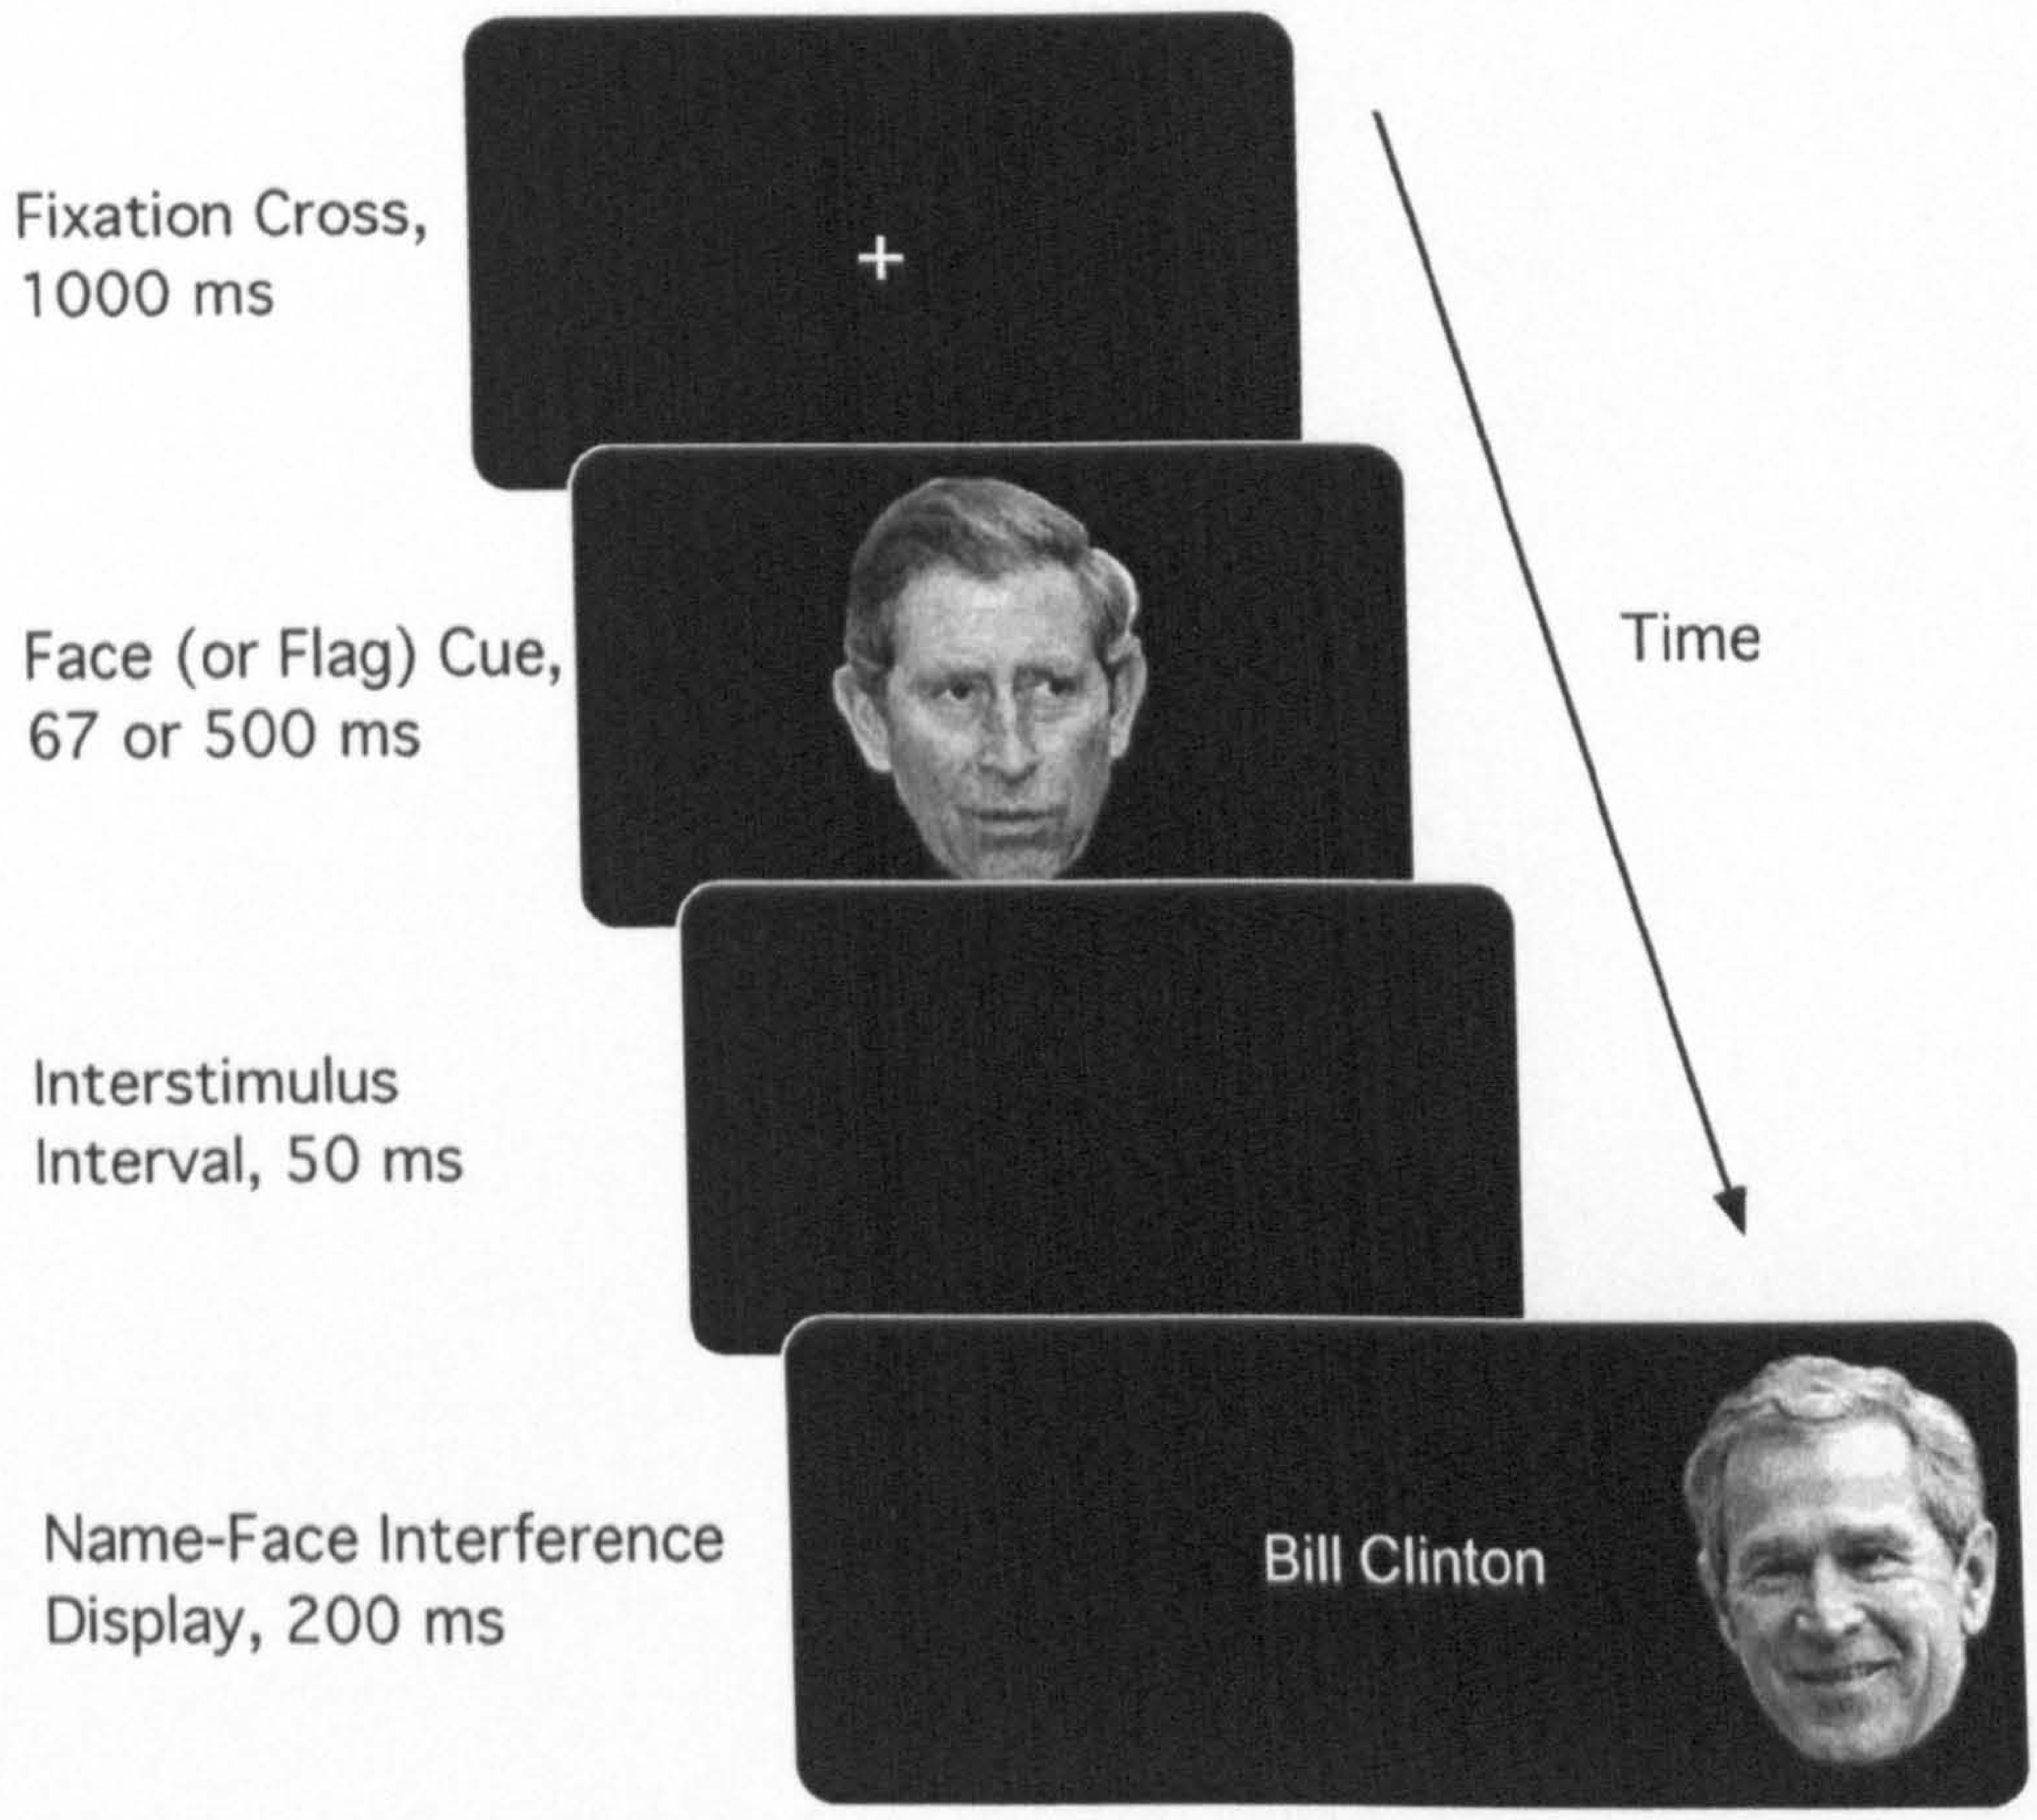

Figure 3.9 Example of a go-trial from Experiment 8. After a 1000 msec fixation, the cue displays were presented for 67 msec (on *Short* trials) or 500 msec (on *Long* trials), followed by a blank screen for 50 msec. The response-competition displays were then presented for 200 msec, followed by a further blank screen until a response was made. The cue could be a face or a flag. Subjects were asked to make an occupational categorization response (pop-star vs. politician) to the name target of the interference displays, but only if it was preceded by a British face or a British flag.

Results

The data from two subjects with overall error rates of 28% and 49% were excluded from the analysis. Accuracy was high for go and no-go face conditions (errors, 4.2% and 2.7%), which confirms that responses were based on information from the cue as well as the interference displays. For go conditions, the means of the median correct RTs and percentage error rates were computed for each level of cue type (face vs. flag), SOA (500 ms vs. 67 ms), and for each level of target-distractor congruency (congruent vs. incongruent occupation). The mean RTs of these conditions are shown in Figure 3.10.

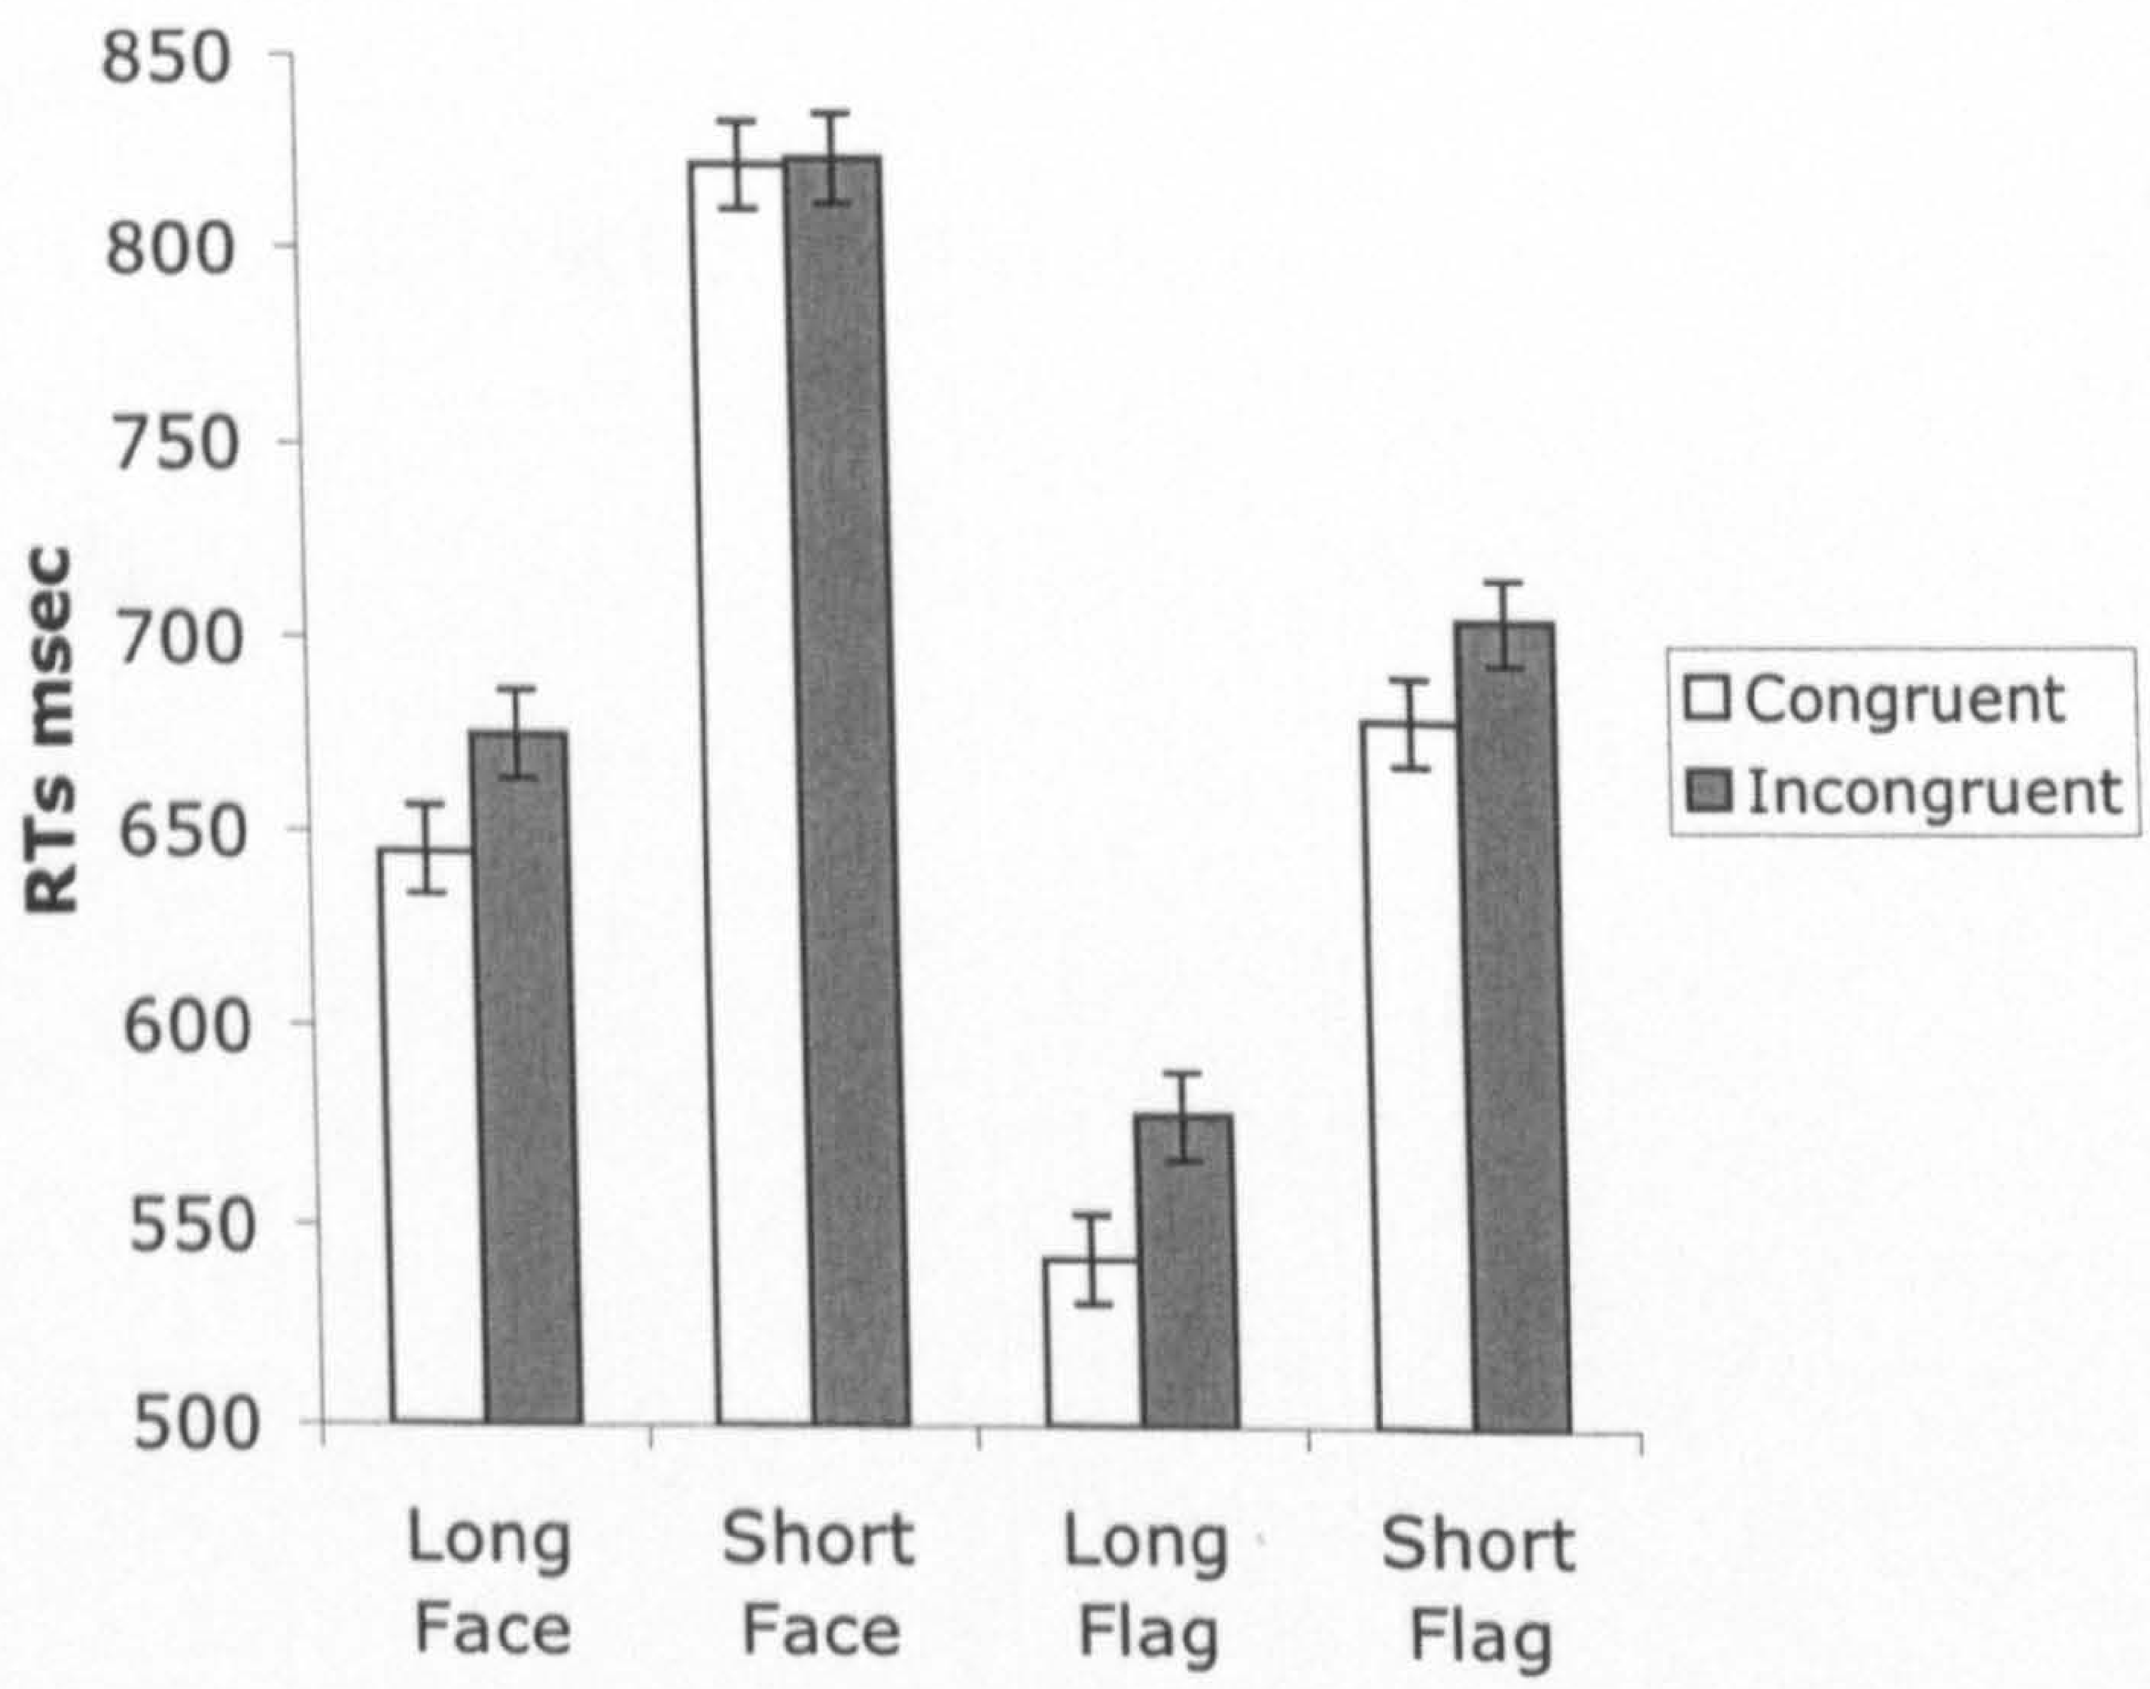

Figure 3.10 Mean reaction times across subjects as a function of distractor congruency, cue type (*Face* vs. *Flag*) and SOA (*Long* vs. *Short*) in Experiment 8. Vertical bars represent the standard error of the means.

A 2 (face vs. flag) x 2 (67 ms vs. 500 ms) x 2 (congruent vs. incongruent) within-subjects ANOVA of the RT data revealed a main effect of cue type,

$F(1,25)=41.74$, $p<.01$, with slower responses following the presentation of face cue, a main effect of SOA, $F(1,24)=211.38$, $p<.01$, with slower responses following short cues, and a main effect of congruency, $F(1,24)=22.79$, $p<.01$, with slower responses to incongruent name-face displays. The effect of cue type was modified by interactions with SOA, $F(1,24)=8.50$, $p<.01$, and congruency, $F(1,24)=7.02$, $p<.01$. Simple main effect analysis revealed congruency effects in the *Short Flag* condition, $F(1,24)=9.96$, $p<.01$, in the *Long Flag* condition, $F(1,24)=12.85$, $p<.01$, and in the *Long Face* condition, $F(1,24)=5.12$, $p<.05$, but no congruency effect was found in the *Short Face* condition, $F(1,24)<1$.

Equivalent analysis was conducted for the error data. Incongruent name-face displays resulted in a small increase in errors in the *Short Flag* condition (incongruent 5.7% vs. congruent 5.3%), in the *Long Flag* condition (3.7% vs. 2.8%), and in the *Long Face* condition (5.1% vs. 4.1%). However, the reverse pattern was found in the *Short Face* condition (incongruent 5.8% vs. congruent 6.7%). ANOVA showed a main effect of SOA, $F(1,24)=9.60$, $p<.01$, with higher errors following short cues. Although marginally higher errors were observed for the face cue conditions, no main effect of cue type was found, $F(1,24)=1.43$. No other comparisons were significant.

Discussion

Experiment 7 examined the temporal conditions under which interference from task-irrelevant face distractors is eliminated. The results show that processing the flags had little effect on subsequent name-face interference: for both short and long flag cues, irrelevant face distractors interfered reliably with name classification. A different response pattern was observed following the presentation of the face cues. Reliable name-face interference was only found following *Long Face* displays. By contrast, face distractor interference was eliminated following immediately preceding, briefly presented faces. This suggests that the ongoing processing of a face is sufficient to prevent the processing of another, subsequently presented face distractor.

However, note that face-cue processing also resulted in longer name categorization RTs in comparison with the flag conditions, which suggests that the face conditions were more difficult in this experiment. Moreover, name RTs increased by almost 200 ms from long face to short face displays. What seems to be happening here is that the ongoing processing of a face's identity was not only extinguishing face distractor interference but also affecting name categorization in a subsequent display. This dramatic decline in performance clearly needs an explanation. The architecture of the IAC model of person recognition (Burton, 1998; Burton et al, 1990, 1999; Burton, Young, Bruce, Johnston & Ellis, 1991; Burton & Bruce, 1993) provides one possible explanation for this effect. According to the IAC model, known faces are recognized by Face Recognition Units (FRUs), one of which exists for each known face. This leads to activation at Person Identity Node (PIN) corresponding to the classification of a person rather than a face. It is at this stage that people can state whether they recognize a person as familiar. Crucially, PINs can also receive activation from persons' names (via Name Recognition Units), and it is held that activation of one PIN leads to the inhibition of others. One possibility is then that activation of a known person's face, the cue item in the case of the present experiment, excites a corresponding person identity node (PIN) in the brain. This not only elicits face processing limits, thus eliminating distractor face processing, but also inhibits person recognition at a general level, resulting in a dramatic increase of name categorization times. However, two findings suggest that this is not a likely explanation. First, although name RTs were slower in the face conditions, erroneous responses to the name targets were reasonably low in all of the conditions. This suggests that face cues did not prevent the processing but only affected the classification of the subsequent name targets. Second, it is difficult to reconcile an explanation in terms of PIN inhibition with the results of Experiment 5, in which face-name and name-name interference for famous persons was found.

An alternative explanation could be that face cues were particularly effective at retaining attentional resources necessary for response to the detriment of the succeeding name targets, even though they did not affect whether the names were actually processed. At present, there is only some indirect evidence that faces have this ability, and this has mostly been obtained with emotionally expressive faces in anxious individuals (e.g. Bradley, Mogg, Falla & Hamilton, 1998; Fox, Russo, Bowles & Dutton, 2001; Fox, Russo & Dutton, 2002). These studies show that classification of a simple perceptual target is slowed when a threatening face in a task-irrelevant location precedes it. I return to this issue in Chapter 5 to investigate whether faces are generally particularly effective at retaining attention in a classification task, in comparison with other stimulus classes.

Overall, the results of Experiment 8 converge with the main claim of Experiments 4-7 that face processing is capacity limited in an interference task such that only a single face can be processed at a time, and extend these findings to the processing of temporally distinct faces. A number of previous studies have already suggested capacity limits for face processing (Boutet & Chaudhuri, 2001; Jenkins et al, 2003;

Palermo & Rhodes, 2002). However, these studies used very different approaches to the experiments reported here. Boutet & Chaudhuri (2001) used overlapping faces, a hypothetical situation that our face processing system is not confronted with outside the laboratory, while the present experiments measured the processing of spatially distinct faces. Palermo & Rhodes (2002) used displays of three faces presented substantially longer ($\geq$ 1.5 s) than the faces in the present experiments (i.e. 200 ms), and required participants to match two peripheral faces before encoding a central target. Under those conditions it is likely, that the three faces were processed sequentially and, consequently, an exact limit in face processing is difficult to specify. Additionally, both these studies tested (immediate) memory for unfamiliar faces, rather than providing a direct test for face processing. Finally, Jenkins et al. (2003) only measured task-irrelevant face processing in a name-face interference paradigm. Without taking resources attributed to task-relevant processing into account, this also makes it difficult to make a direct inference about capacity limits. Thus, the experiments in this chapter add a novel set of data in support of the notion that face processing may be capacity limited.

Previous studies suggesting capacity limits for face processing also observed different patterns for upright and inverted faces (Boutet & Chaudhuri, 2001; Jenkins et al, 2003; Palermo & Rhodes, 2001), suggesting a limit specifically for upright face processing. Similarly, others report that the processing of an irrelevant face seems unaffected by variations in task-relevant processing load of nonface stimuli (Jenkins et al, 2002; Lavie et al, 2003), again suggesting a face processing capacity. Although the present experiments were not intended to examine face-specificity, the results also converge with these suggestions by demonstrating capacity limits in face processing under conditions that allow the processing of multiple stimuli: single face distractors were processed alongside nonface targets in Experiments 4-6 and subsequent to nonface stimuli in Experiment 8, and none of the nonface comparisons (forenames, famous names, images of national flags) in Experiments 4-7 displayed analogous processing limits.

Alternatively, it is conceivable that the processing of faces, which are a visually complex and homogeneous category of stimuli, is simply more demanding of general resources than processing printed names or images of flags. Thus, two faces might exceed general processing capacity, even when two flags or two names, or a face and a flag/name do not. Therefore, some nonface stimuli may be subject to corresponding processing limits. Likewise, some face-specific phenomena have repeatedly been attributed to the fact that we possess a great level of expertise in face processing (see e.g. Gauthier & Logothetis, 2000), which is required to discriminate between highly similar exemplars of a particular class. One might thus expect that other visual stimuli, for which we possess a high level of expertise, may be subject to such processing limits. However, target RTs were as fast or faster for face targets than for nonface targets in Experiments 4-6, which suggests that face processing was no more difficult than the processing of the comparison stimuli, and the error data support this impression. In addition, Jenkins et al (2003) report that an intact, irrelevant face produces no more dilution of object-word interference than a phase-shifted face. If faces are simply a disproportionate drain on general resources, it is hard to see why this should be so; one would expect them to produce disproportionate dilution in that situation. Although many theories of attention already propose separate modality-specific processing limits (see e.g. Pashler, 1998; Schmitt, Postma & de Haan, 2000, 2001), the present results might thus reflect finer, domain-specific subdivisions of processing limits, acting to constrain processing of a particular range of stimuli within the same modality.

Several other aspects of the present findings merit further discussion. First, in Experiments 4-6, face distractors interfered more with the categorization of names and flags than these stimuli interfered with faces. This replicates Young et al (1986) who observed the same pattern in a semantic categorization task. Young et al (1986) suggested that this pattern emerges from the encoding of visual information, whereby faces may be encoded into a form that particularly suits categorization tasks in contrast to names, which may be encoded for naming tasks. In fact, in naming tasks names do seem to interfere more with faces than vice versa (Young et al, 1986). However, according to this explanation it is difficult to see why faces should interfere more with flags in the nationality task of Experiment 5 than flags interfered with faces. If anything, flags should have been coded more readily into a nationality than faces, which are more visually complex than the salient flag patterns and code much more information than a person's nationality (e.g. sex, emotional expression, occupation).

Second, in Experiment 7 multiple face distractors not only failed to interfere with the classification of the face targets but also the nonface targets, images of flags. It is conceivable that multiple face distractors do not interfere with nonface targets for the same reason that a single face distractor produces more interference than a nonface stimulus, with the notable exception of when a face target is paired with a face distractor. There have been numerous recent claims that faces are amongst a class of stimuli capable of capturing attention, even under conditions that deem this unlikely. For example, Vuilleumier (2000) observed that line-drawn faces are detected more frequently in the neglected field of patients suffering from unilateral neglect than written names, meaningless shapes, or scrambled faces. Similarly, faces may have an advantage in capturing attention in neurological normal participants. Thus, normal participants detect an intact schematic face more quickly amongst an array of scrambled faces than a scrambled face amongst intact face distractors (Mack, Pappas, Silvermann & Gay, 2002). Schematic faces are also detected more often in a stream of visual stimuli than inverted faces or nonface comparisons (Mack et al, 2002). In addition, Ro, Russell & Lavie (2001) showed that participants noticed changes concerning pictures of real faces more accurately and more quickly than for other objects in a change detection paradigm (but see Palermo & Rhodes, 2003).

If faces do capture attention, one might expect solitary face distractors to interfere strongly with relevant nonface processing by drawing processing resources to the distractor location (as in Experiments 4-6). And if face processing involved limited resources, one might expect it to fail when confronted with several simultaneously presented competing inputs of equal status, such as the four task-irrelevant face distractors in Experiment 7. In contrast, if competition between simultaneously presented faces can be resolved, for example, by deliberate attention to one of them (e.g. a fixated target face), that privileged face could plausibly monopolize the limited resources, to the detriment of other faces present (i.e. the distractor faces in Experiments 4-8).

The final chapter of this thesis provides a fuller discussion of the question whether the conditions in which faces tend to capture attention are related to those in which capacity limits seem to apply. For the moment, the exact locus of any 'bottleneck' in face processing remains difficult to specify. The present results only imply that face distractor processing, under conditions capable of elucidating capacity limits, stops short of full semantic analysis. Yet, it seems unlikely that these processing limits are located at a semantic level, as faces interfered with names and national flags during semantic classification. The results of Experiment 4 suggest that capacity limits in face processing occur at an earlier, perceptual stage, since even salient face-related sex information was unavailable in face-face displays. However, sex and identity are dissociable face dimensions (Bruce et al, 1987; Ellis et al, 1990), and a processing limit for one of these does not imply the same limit for the other. Moreover, extinguished distractor faces presumably undergo some superficial processing. At the very least they must register as faces at some level, otherwise multiple distractor faces would not compete for processing resources. Therefore, the aim of the next chapter is to examine whether the present face processing limits apply prior to semantics.

Chapter 4 Capacity Limits for Face Processing: Repetition Priming of Distractor Faces from Two-item Displays

Introduction

The experiments in Chapter 3 examined whether responses to face targets can be affected by concurrently presented distractor faces in interference paradigms. In the first study, Experiment 4, participants were required to classify the sex of unfamiliar faces or short forenames, while ignoring a face or name distractor in the display. Subsequent experiments repeated this design with famous faces and famous names (Experiment 5), famous faces and images of national flags (Experiment 6 & 7), or a combination of both (Experiment 8), and with a single distractor stimulus (Experiments 4-6 & 8) or multiple distractors (Experiment 7) in semantic classification tasks. These experiments demonstrated that interference from distractor faces is extinguished by processing a face, but not by processing nonface stimuli (e.g. names, flags). This distractor extinction effect occurred in a context in which faces interfered with nonface targets, and nonface distractors interfered with the classification of both face and nonface targets. Collectively, these findings suggest that face processing may be capacity-limited, such that only a single face can be processed at a time.

The question that is addressed in this chapter is whether these processing limits are still observed in tasks that do not require semantic or sex processing. According to established models of person recognition, the retrieval of personal semantic information is relatively deep and follows face identification (e.g. Bruce & Young, 1986; Burton, Bruce & Hancock, 1999; Burton, Bruce & Johnston, 1990). Several lines of evidence also indicate that sex and identity information from faces are processed independently and in parallel (e.g. Bruce, 1986; Bruce, Ellis, Gibling & Young, 1987). Consequently, the possibility remains that extinguished distractor faces were processed at some level during face target classification. Indeed, in Chapter 3 it was already suggested that these distractors might at least undergo some superficial processing to compete for limited processing resources. Such processing could extend beyond that minimum, perhaps involving access to face identity.

One established method of assessing whether a visual stimulus has been processed is repetition priming. This is a facilitation in processing an item due to prior exposure to that item. In the face domain, such tasks typically consist of a prime phase during which participants are exposed to famous faces, followed by an interval of a few minutes and an unexpected test phase involving familiarity judgements (famous/unfamiliar) to primed and unprimed famous faces and some unfamiliar filler faces. The reliable finding here is that responses to primed faces are faster than to unprimed famous faces and unfamiliar faces (e.g. Bruce & Valentine, 1985; Ellis, Young, Flude & Hay, 1987)

Repetition priming is a robust and long lasting effect, that is even found when different images of the same person's face are used (e.g. Bruce & Valentine, 1985; Ellis, Flude, Young & Burton, 1996), across changes in context (e.g. location, task, Bruce, Carson, Burton & Kelly, 1998), and persists over radically altered but still recognizable representations of a face (e.g. part-face to whole-face, Brunas, Young & Ellis, 1990). However, although repetition priming proceeds independent of the judgement being made at prime phase (Ellis, Young & Flude, 1990), and even when no explicit judgement is required (Jenkins, Burton & Ellis, 2002), it is usually not observed onto sex or expression decisions at test (Ellis et al, 1990; but see Goshen-Gottstein & Ganel, 2000, who obtained priming using part-faces), or onto familiarity decisions when priming of unfamiliar faces is measured (Ellis et al, 1990). In addition, repetition priming is domain-specific when the typical familiarity decision is used. Thus, faces prime faces and names prime names, but one type of stimulus does not prime the other (e.g. Bruce & Valentine, 1985; Burton, Kelly & Bruce, 1998; Ellis et al, 1996). Consequently, it is held that repetition priming operates within the system that responds to facial identity (e.g. Burton et al, 1990, 1999; Ellis et al, 1990), which becomes activated automatically by any recognizable view of a known person's face. Note that this type of priming can be dissociated from cross-domain repetition priming which occurs when semantic judgements are used at prime and test (Burton et al, 1998; McNeill, Burton & Ellis, 2003), reflecting shared semantic access to faces and names following person recognition. Unlike the interference paradigms of the preceding chapter, domain-specific priming can therefore be used to assess face identity processing directly, independent of subsequent semantic processing.

The following experiments utilized this characteristic to provide a further test for face processing limits. Of specific interest was whether repetition priming would reveal any evidence of face distractor processing when it is presented alongside a task-relevant face target, in comparison with the extinction of face-face interference in Chapter 3. Jenkins, Burton & Ellis (2002) report that irrelevant famous face distractors are primed automatically during nonface target processing, even under conditions of high relevant processing load and when participants have no recollection of seeing these faces at prime phase. If several faces can be processed simultaneously, one might thus also expect some repetition priming from task-irrelevant famous distractor faces during task-relevant face processing.

Experiment 9

Experiment 9 examined capacity limits in face processing by measuring repetition priming from two-item displays. To provide an analogue to the interference displays of Chapter 3, the subjects' task was to classify famous face targets or nonface comparisons, in this case images of national flags, as American or British while ignoring a flanking famous face distractor. Within these displays faces could be primed under three conditions: i) as task-relevant face targets (the *Target Face* condition), ii) as irrelevant face distractors that were presented alongside these face targets (the *Face-Face* condition), and iii) as irrelevant face distractors presented alongside flag targets (the *Flag-Face* condition). The extent to which these faces were processed was then assessed in a surprise test phase via speeded familiarity judgements (famous/unfamiliar) to primed famous faces, unprimed famous faces and some unfamiliar filler faces. The unprimed famous faces were included as a baseline (the *Unprimed* condition) to determine whether face distractors were subject to any repetition priming.

As mentioned in the introduction to this chapter, repetition priming can usually be seen even when different images of the same person's face are used at prime and test, reflecting repeated processing within the face recognition system. The present experiment therefore measured cross-image priming to determine whether irrelevant face distractors can be processed to recognition. Face targets should produce the standard pattern of repetition priming, and comparisons with this condition and the *Unprimed* condition should indicate the extent of any distractor priming effects. If identity information from two faces can be processed simultaneously, then face distractors should show repetition priming regardless of target type. Alternatively, if such processing is limited to just a single face, then distractor priming should be eliminated in the *Face-Face* condition. However, since face distractors interfered strongly with nonface target classification in the preceding chapter, the *Flag-Face* condition should allow for reliable distractor priming.

Method

Subjects Thirty-two British undergraduate students from the University of Glasgow, aged 20-28 years, were paid a small fee to participate in the experiment. All reported normal or correct vision.

Design & Stimuli For the prime phase, photographs of forty British and of forty American celebrities (see Appendix B), and of twenty national flags (10 Stars & Stripes, 10 Union Jacks) served as stimuli. All images were greyscale and measured 3.6 cm x 4.5 cm (3.4° x 4.3° of VA at a viewing distance of 60 cm). Faces were presented with their external features intact (i.e. hair, face outline) and the flags were manually cropped to roughly elliptical shapes to produce a closer resemblance with the face stimuli. These images were used to construct displays consisting of a central target, which could be either a face or a flag, and a flanking face distractor (see Figure 4.1 overleaf). Distractors were equally likely to appear left or right of the target, 1.0 cm (1.0° of VA) from the nearest target contours, and were counterbalanced so that target and distractor were of the same nationality and of the same sex (face-face pairings only) in half of the displays. Combining each of the 20 flag targets and the same number of face targets with a face distractor resulted in 40 stimulus displays. Overall, this involved 60 of the 80 face stimuli. The 20 remaining face identities were reserved as unprimed controls for the second phase. The 80 face images were rotated around these conditions so that over the course of the whole experiment, each famous face appeared in each condition an equal number of times.

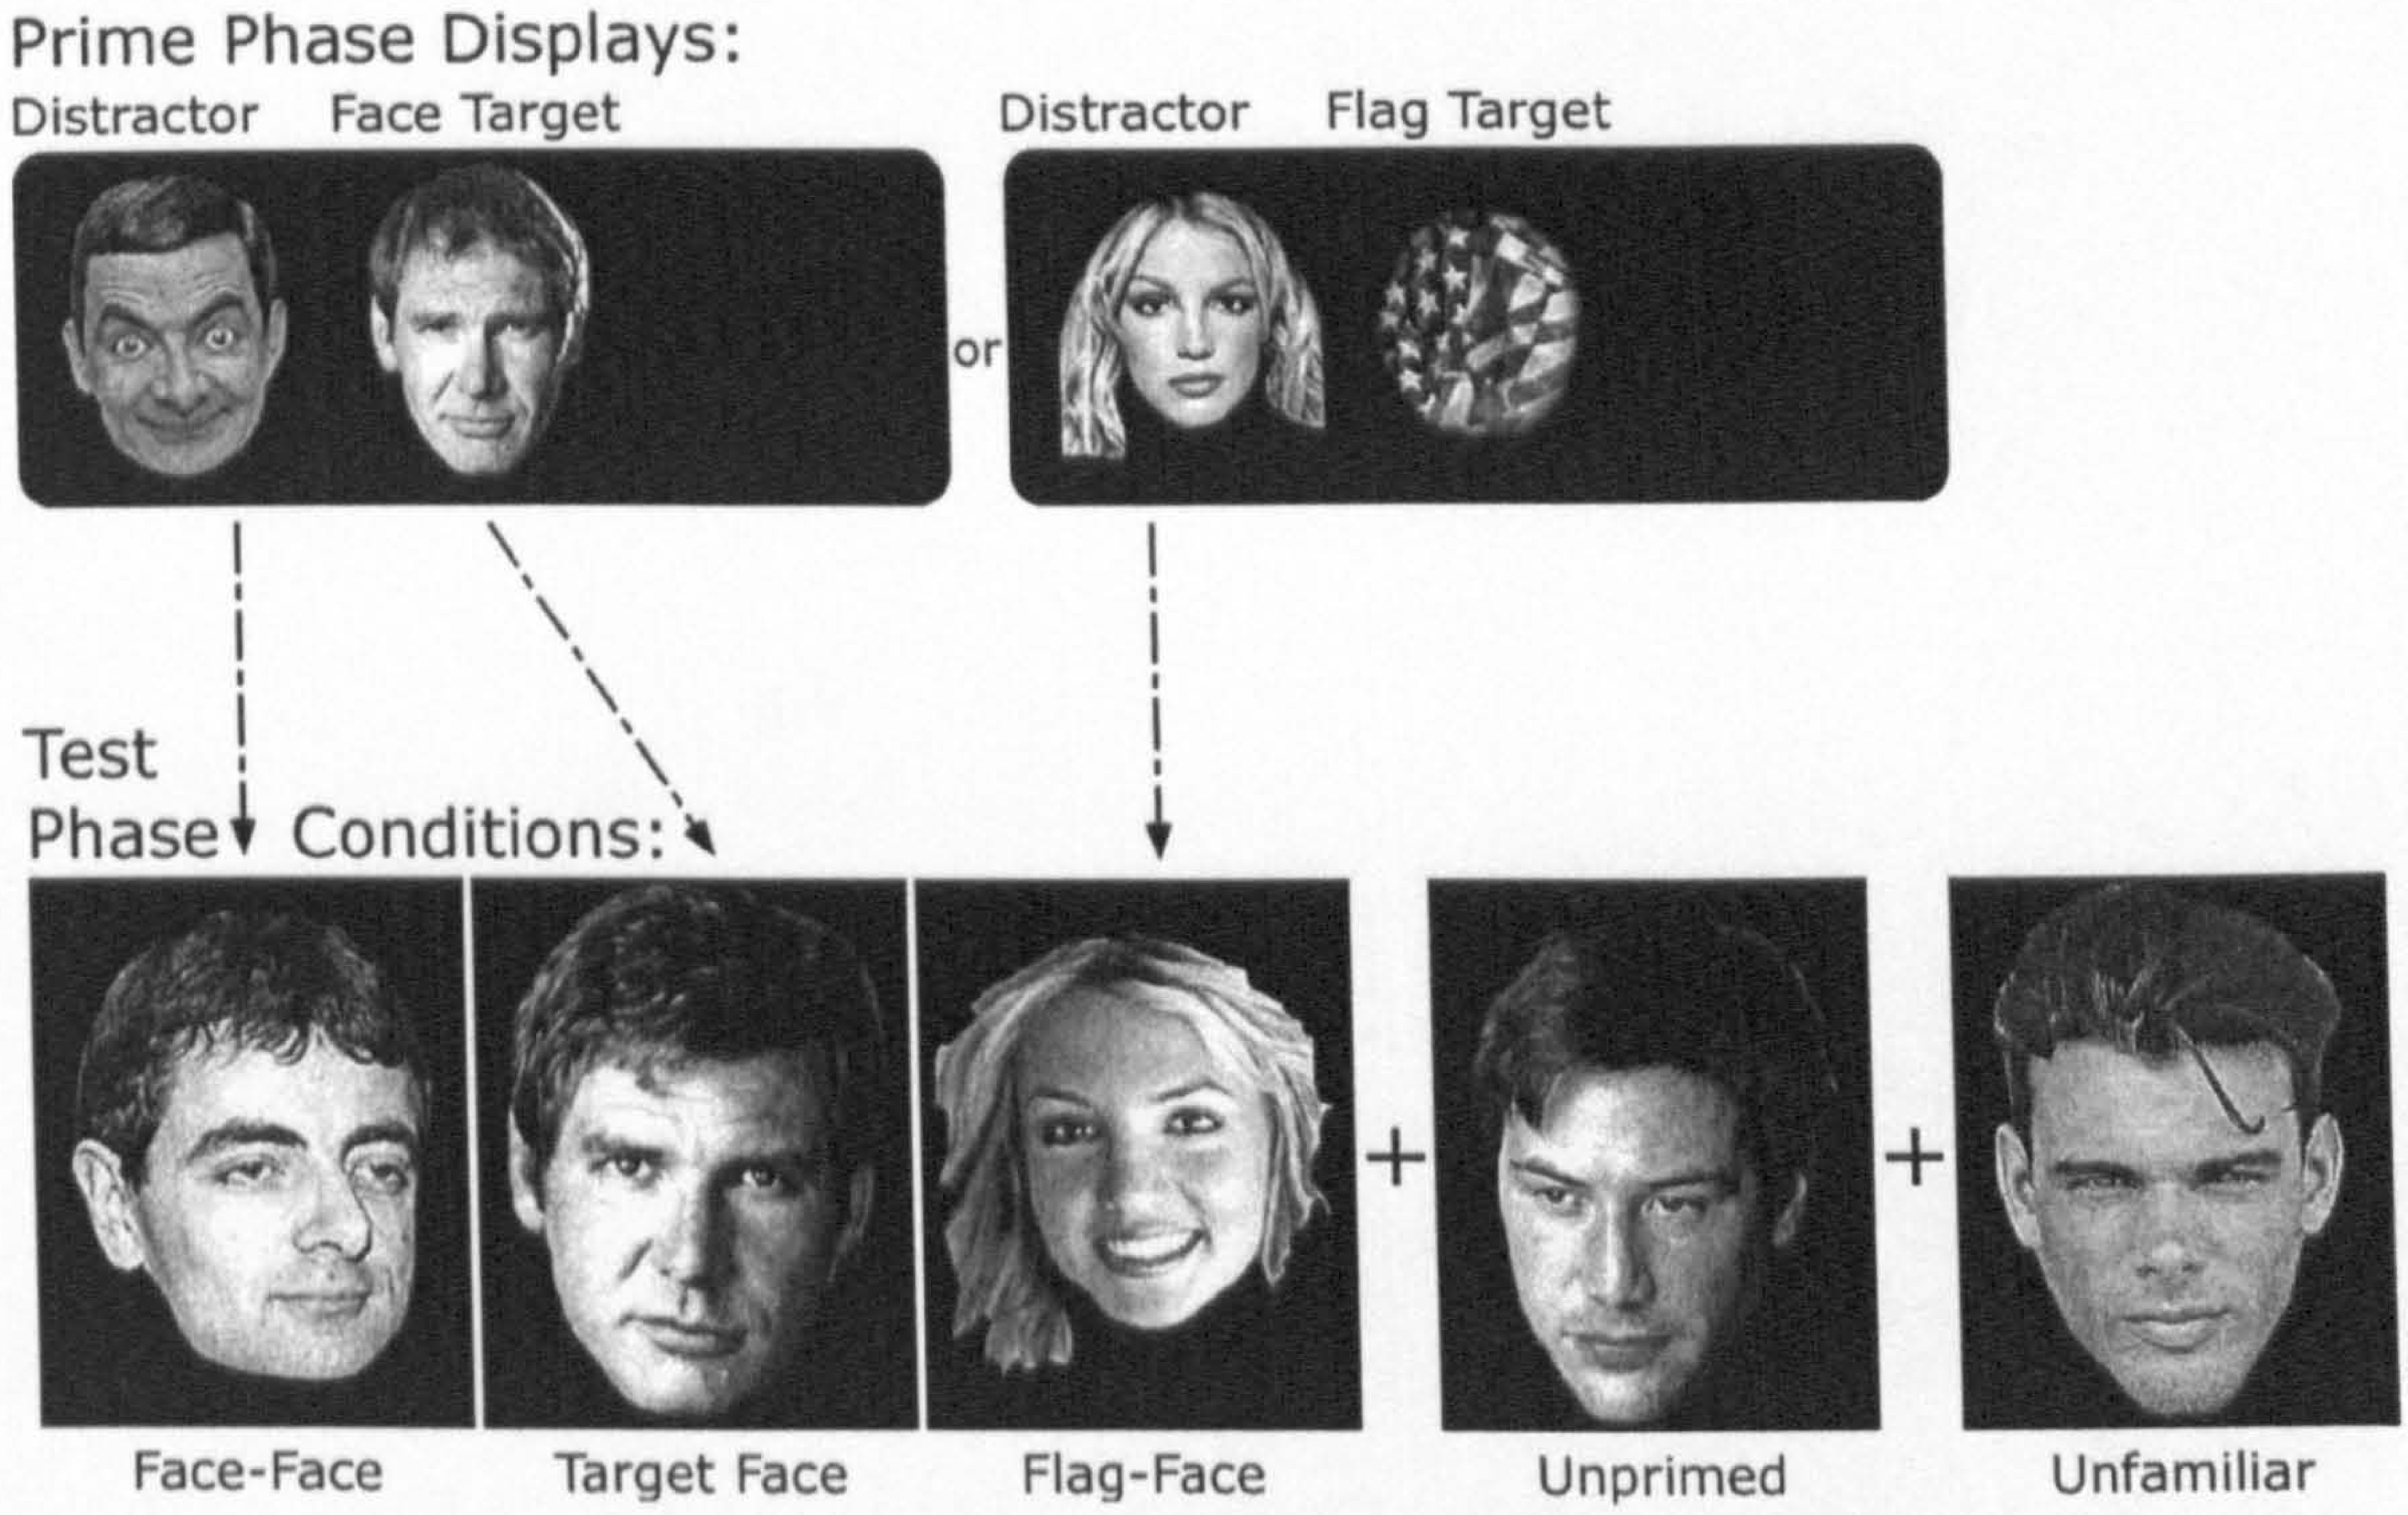

Figure 4.1 Illustration of the experimental conditions in Experiment 9. At prime phase, displays consisting of a face or flag target and a flanking face distractor were shown for 200 msec in an American/British categorization task. At test phase, repetition priming of these faces was assessed via familiarity judgements (famous versus unfamiliar) to different photographs of the primed face targets (the *Target Face* condition) and the distractors (in the *Face-Face* and *Flag-Face* conditions), and to photographs of unseen famous faces (the *Unprimed* baseline condition) and some unfamiliar filler faces. The famous faces here are, from left to right at test phase: Rowan Atkinson, Harrison Ford, Britney Spears, and Keanu Reaves.

Different (unseen) images of the same 80 celebrities' faces, intermixed with 80 unfamiliar faces, were used in the test phase in a speeded familiarity task (see Figure 4.1). The unfamiliar faces were photographs of anonymous male and female models, which provided a close match for the famous faces in terms of image quality, approximate age and good looks. All images were presented in greyscale at fixation, one at a time, at a size of 6.0 cm x 7.5 cm (5.7° x 7.3° of VA) on a dark background. An Apple Macintosh computer was used to present stimuli and record responses using PsyScope 1.2.5. Viewing distance was fixed at 60 cm by means of a chinrest.

Procedure In the prime phase, each trial began with a fixation cross for 750 ms, followed by a target-distractor display for 200 ms (i.e. too briefly to permit stimulus-responsive saccades), and a blank screen which remained on until a response was registered. Subjects made speeded judgements concerning whether the central target was American or British by pressing one of two buttons ("D" and "L") on a standard computer keyboard, but were emphatically instructed to ignore the task-irrelevant distractors. Subjects were encouraged to guess if they were uncertain regarding the correct answer. If no response was made within 2.5 seconds of stimulus onset, the next trial was initiated. All subjects underwent a short practice block of 16 trials, consisting of an additional 4 flags and 12 famous face images. Each of these images was displayed twice during practice and none were encountered subsequently. An experimental block of 40 randomly intermixed trials followed the practice block.

Upon completion of the prime phase subjects were instructed to remain seated at the computer by an onscreen message. The experimenter then entered the laboratory and initiated the (unexpected) test phase. In this phase, each trial consisted of a fixation cross for 750 ms, followed by a single face image, which was displayed at fixation until a response had been made. Subjects were told to make famous/unfamiliar judgements to these faces as quickly and as accurately as possible via a two-choice keypress response ("C" vs. "."). Subjects underwent 4 blocks consisting of 20 famous and 20 unfamiliar trials presented in random order. They were able to rest between blocks, initiating the next block by pressing the space bar.

Results

Priming Phase Accuracy in the prime phase was important for confirming that subjects were focussing on the target stimuli. Incorrect responses were discarded and mean RTs and error rates were calculated for responses to face and flag targets. The cross-subject averages of these means were: face targets 915 msec (error rates 19.4%), flag targets 715 msec (error rates 7.8%). Prime phase data was not analyzed further.

Test Phase The data of principal interest were the responses to primed and unprimed famous faces at test phase. Incorrect responses and RTs exceeding 2 seconds (less than 1% of correct responses) were excluded from analysis. The mean RTs and error rates, averaged across subjects, are shown in Figure 4.2 as a function of experimental condition.

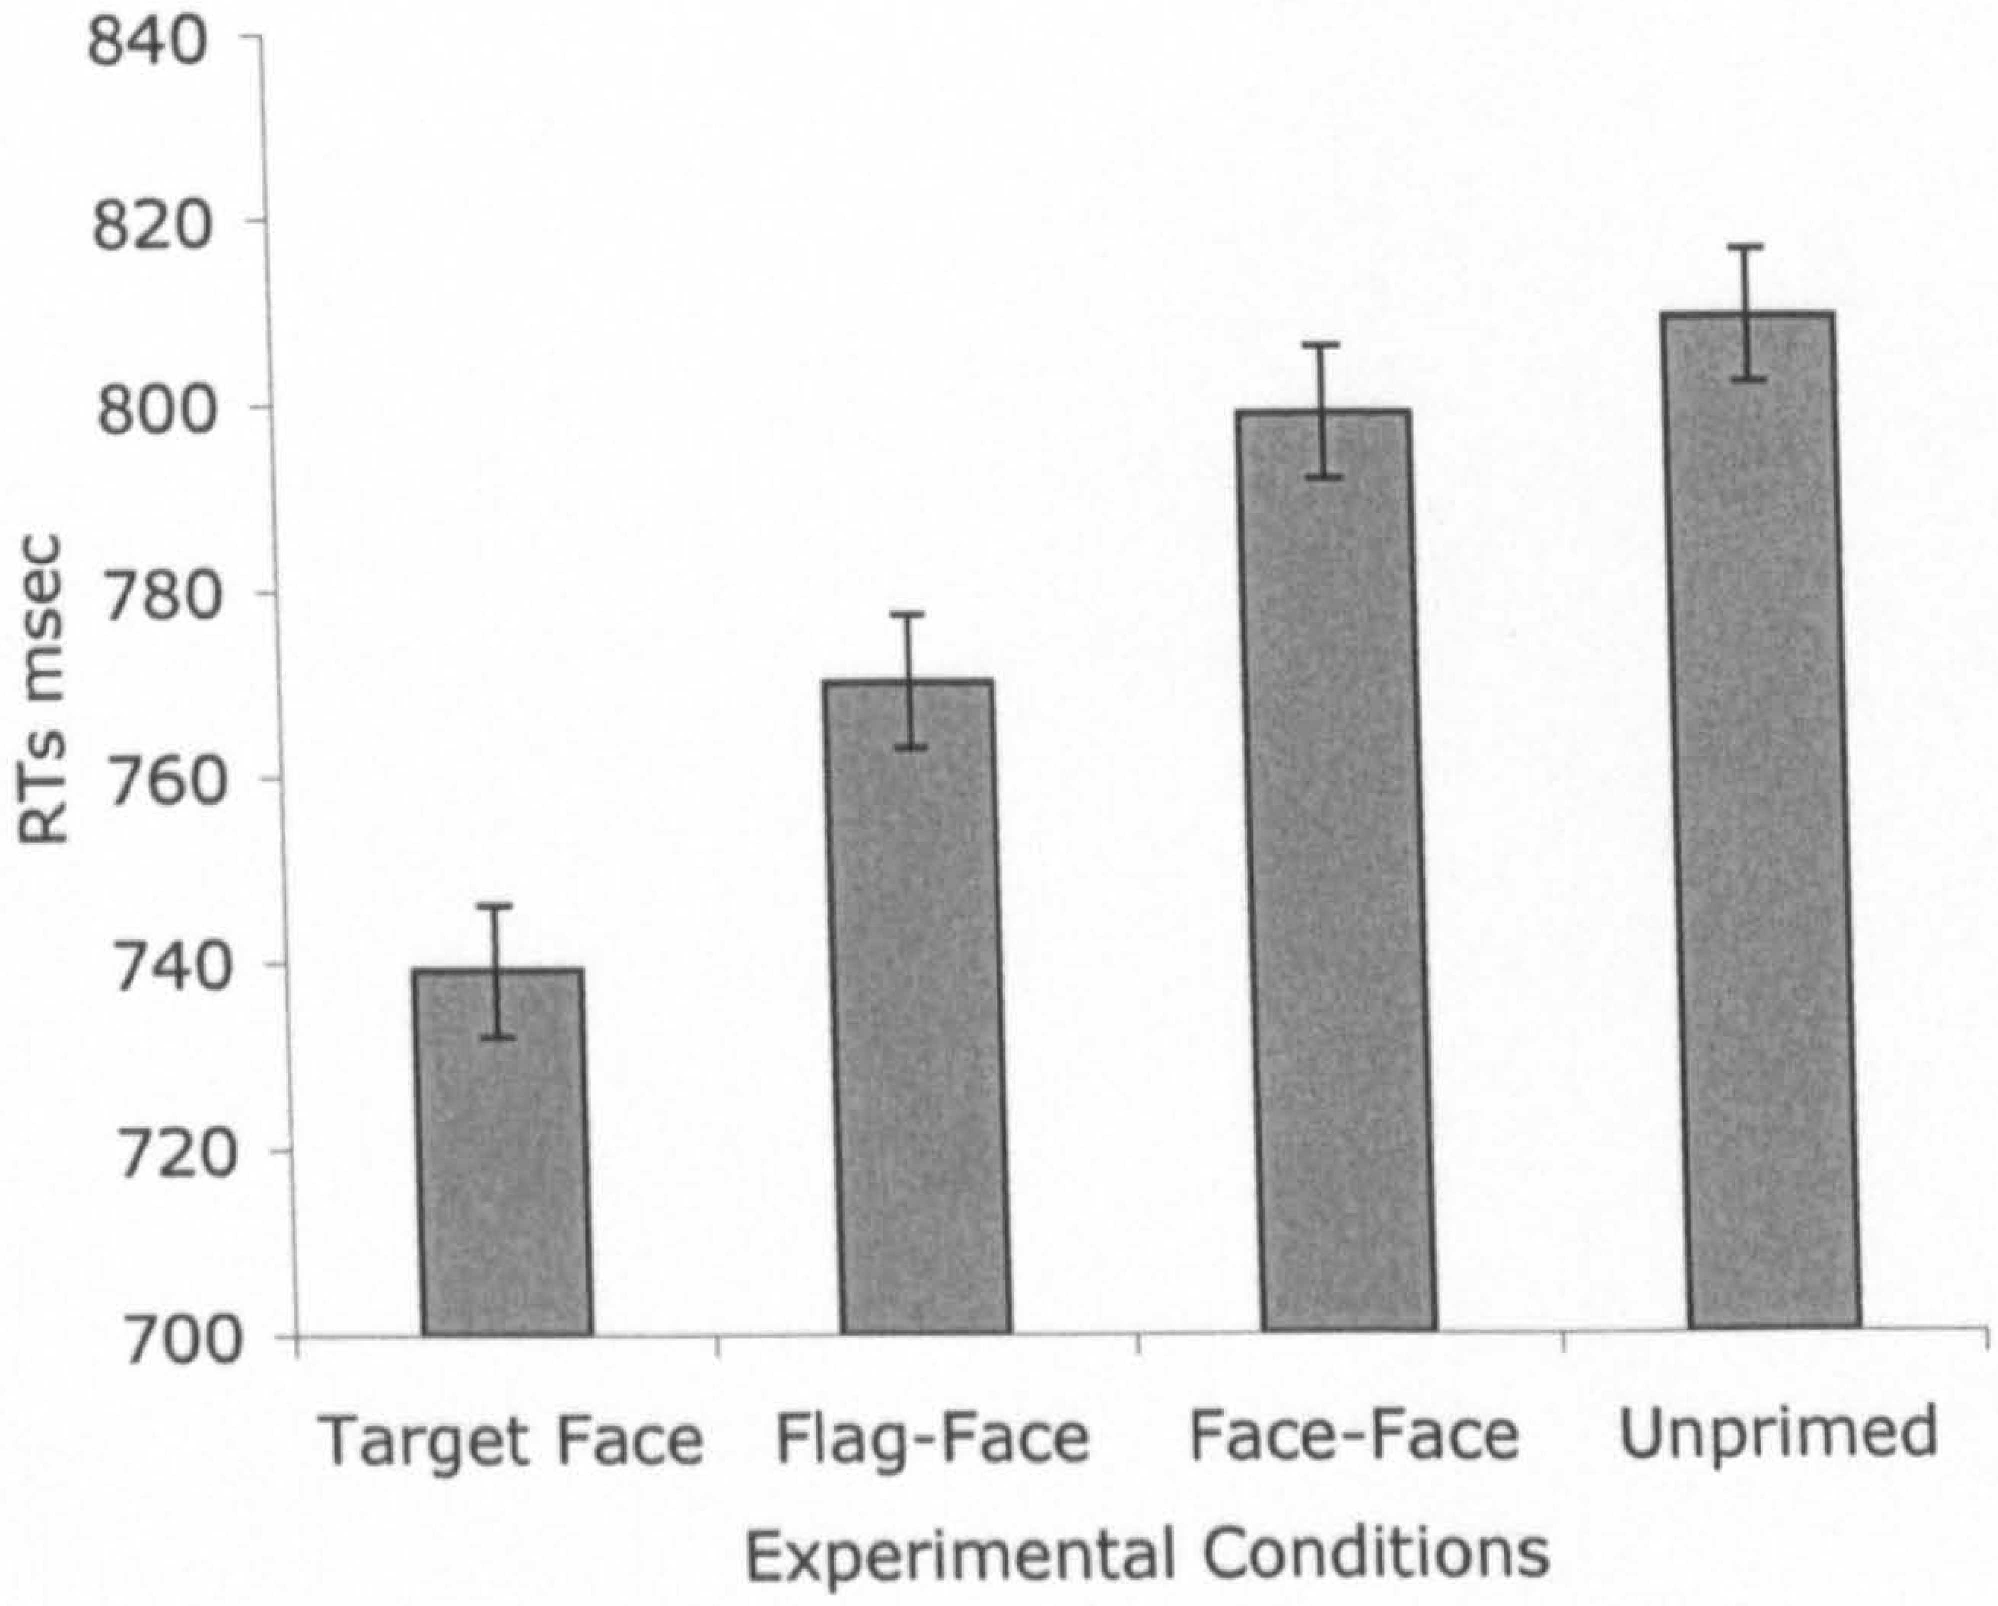

Figure 4.2 Mean correct responses (n=32) to famous faces in the surprise test phase of Experiment 9. Performance is shown as a function of prime type; *Target Face*, *Flag-Face*, *Face-Face*, and *Unprimed*. Standard error bars are shown.

A one-factor within-subjects ANOVA (*Target Face* versus *Face-Face* versus *Flag-Face* versus *Unprimed*) showed an effect of prime type, $F(3,93)=19.64$ $p<.01$. Post hoc comparisons with Tukey's HSD test ($p<.05$) revealed more repetition priming in the *Target Face* condition in comparison with each of the other conditions (*Face-Face*, *Flag-Face*, and *Unprimed*), indicating greatest repetition priming for task-relevant faces. More importantly, the *Flag-Face* condition also showed significant repetition priming in comparison with both the *Face-Face* and the *Unprimed* conditions ($p<.05$). However, there was no difference between the *Face-Face* condition and the *Unprimed* condition.

Errors were made on 11.4% of *Target Face* trials, 13.6% of *Flag-Face* trials, 16.3% of *Face-Face* trials, and on 14.7% of *Unprimed* trials. Error rates were analysed using a one-factor (*Target Face* versus *Face-Face* versus *Flag-Face* versus *Unprimed*) ANOVA, which showed a main effect of prime type, $F(3,93)=3.33$, $p<.05$, reflecting higher errors for *Face-Face* primes than *Target Face* primes (Tukey's HSD, $p<.05$). No other comparisons were significant.

Discussion

The results show maximal repetition priming for face targets, reflecting the task-relevance of these images at prime phase. Irrelevant face distractors also showed repetition priming when presented with flag targets. Although this priming effect was smaller than for the face targets, this processing evidently involved access to the distractors' identities since it survived a change in image between prime and test. By contrast, no priming was found for face distractors during face target processing. Given that face priming has previously been shown to be unaffected by relevant nonface processing, even under conditions of high processing load (Jenkins et al, 2002), the absence of any reliable priming effects in this condition is striking. Overall however, the present data are analogous to the findings of Chapter 3, where face distractor interference was consistently extinguished by relevant face but not by relevant nonface processing. Thus the results further support the claim that face processing is capacity limited. In contrast to Chapter 3, face processing was now measured prior to personal semantic information and independent of facial sex. The present findings therefore exclude the possibility that the previously observed extinction of face-face interference reflects processing limits only for sex information or at a semantic level, and extends these limits to facial identity.

Experiment 10

The preceding experiment showed that face target processing eliminates cross-image repetition priming of a simultaneously presented face distractor. Although these findings converge with the idea that face processing is capacity limited, they are nonetheless surprising as repetition priming onto familiarity decisions usually proceeds automatically at prime phase, thus greatly minimizing any face processing demands (see e.g. Ellis et al, 1990; Jenkins et al, 2002). The basis for the next experiment is the observation that greatest repetition priming is observed when identical images are used at prime and test (e.g. Bruce & Valentine, 1985; Ellis et al, 1987). In this way, it might thus be possible to further optimize the conditions for face distractor priming.

Intriguingly, same-image repetition priming has never been found under conditions that do not allow cross-image priming. However, the possibility exists that a same-image priming advantage might arise to some extent independent of the face recognition system. Some researchers propose that this effect could be attributed to facilitation at a separate "pictorial" memory stage (e.g. Jacoby, 1983; Jacoby & Brooks, 1984) or from a processing overlap in the visual system between prime and test phase (e.g. Blaxton, 1989; Roediger, 1990; for a possible explanation see also Ellis, Burton, Young & Flude, 1997). Such "episodic" accounts emphasize the retrieval of stored event memories, whereby repetition priming is optimized when the processing requirements of prime and test phase are most similar. On their own, these accounts are not sufficient to explain repetition priming of familiar faces. For instance, repetition priming across radically different images of the same person (e.g. Bruce & Valentine, 1985; Brunas et al, 1990), cross-domain priming from faces to names (Burton et al, 1998), and reports that repetition priming is not observed onto sex or expression judgements when the same faces and the same decisions to these faces are used at prime and test (Ellis et al, 1990) all resist effortless accommodation into purely episodic accounts.

Nonetheless, several studies have also reported priming for unfamiliar faces, effects that are unlikely to arise within the face recognition system. In one study, Khurana (2000) asked participants to match the second and the fourth face in a row of five unfamiliar faces while ignoring the three remaining distractor faces. Negative priming was found when the face targets consisted of the to-be-ignored distractors of the preceding trials. In addition, repetition priming was found when distractors were disrupted with high frequency noise or contrast inverted[2]. Others also report some long term negative priming effects with unfamiliar faces in more conventional priming designs (Ellis et al, 1990, Experiment 2; Young, McWeeney, Hay & Ellis, 1986b, Experiment 4), which provides further evidence that information from unfamiliar faces is represented at some level in the visual system. Moreover, negative priming and repetition priming have recently been observed with *novel* object shapes, for which subjects could not possess stored, pre-existing representations (DeSchepper & Treisman, 1996). This provides at least some suggestive evidence that same-image priming might occur at a general processing stage, operating independent of the face recognition system. Even if face distractors are not subject to cross-image priming during face target processing, they might therefore still produce some same-image priming. This was the focus of the present experiment. As before, subjects performed a nationality

[2] Note that these results do not contradict present claims for face processing limits, as displays were presented for a substantial duration, until a response was made. Under those conditions, the faces may have been processed sequentially.

categorization task onto famous faces and images of flags while trying to ignore a flanking famous face distractor, and any face processing at this stage was then assessed via speeded familiarity judgements to primed and unprimed faces. In contrast to Experiment 9, identical face sets were now used at prime and at test phase.

Method

Subjects Thirty-six British undergraduate students from the University of Glasgow, aged 18-30 years, were paid a small fee to participate in the experiment. All reported normal or corrected vision and had not participated in the previous experiment.

Stimuli & Procedure Apparatus, stimuli and procedure were the same as those in Experiment 9, except that the face images used in the prime phase were now the same as those used in the test phase. As in Experiment 1, the participants completed a practice block of 16 trials and an experimental block of 40 trials in the priming phase. This was followed by an unexpected test phase consisting of 4 blocks of 40 familiarity judgements.

Results

Priming Phase Incorrect responses to face targets were discarded and mean correct RTs and error rates were then calculated for responses to face and flag targets. The average of these means across subjects for each target type were: face targets 959 msec (error rates 16.6%), flag targets 811 msec (error rates 10.4%).

Test Phase Incorrect responses and RTs exceeding 2 seconds were excluded from analysis (less than 1% of all correct 'famous' responses). The mean RTs and error rates across subjects are shown in Figure 4.3 as a function of experimental condition.

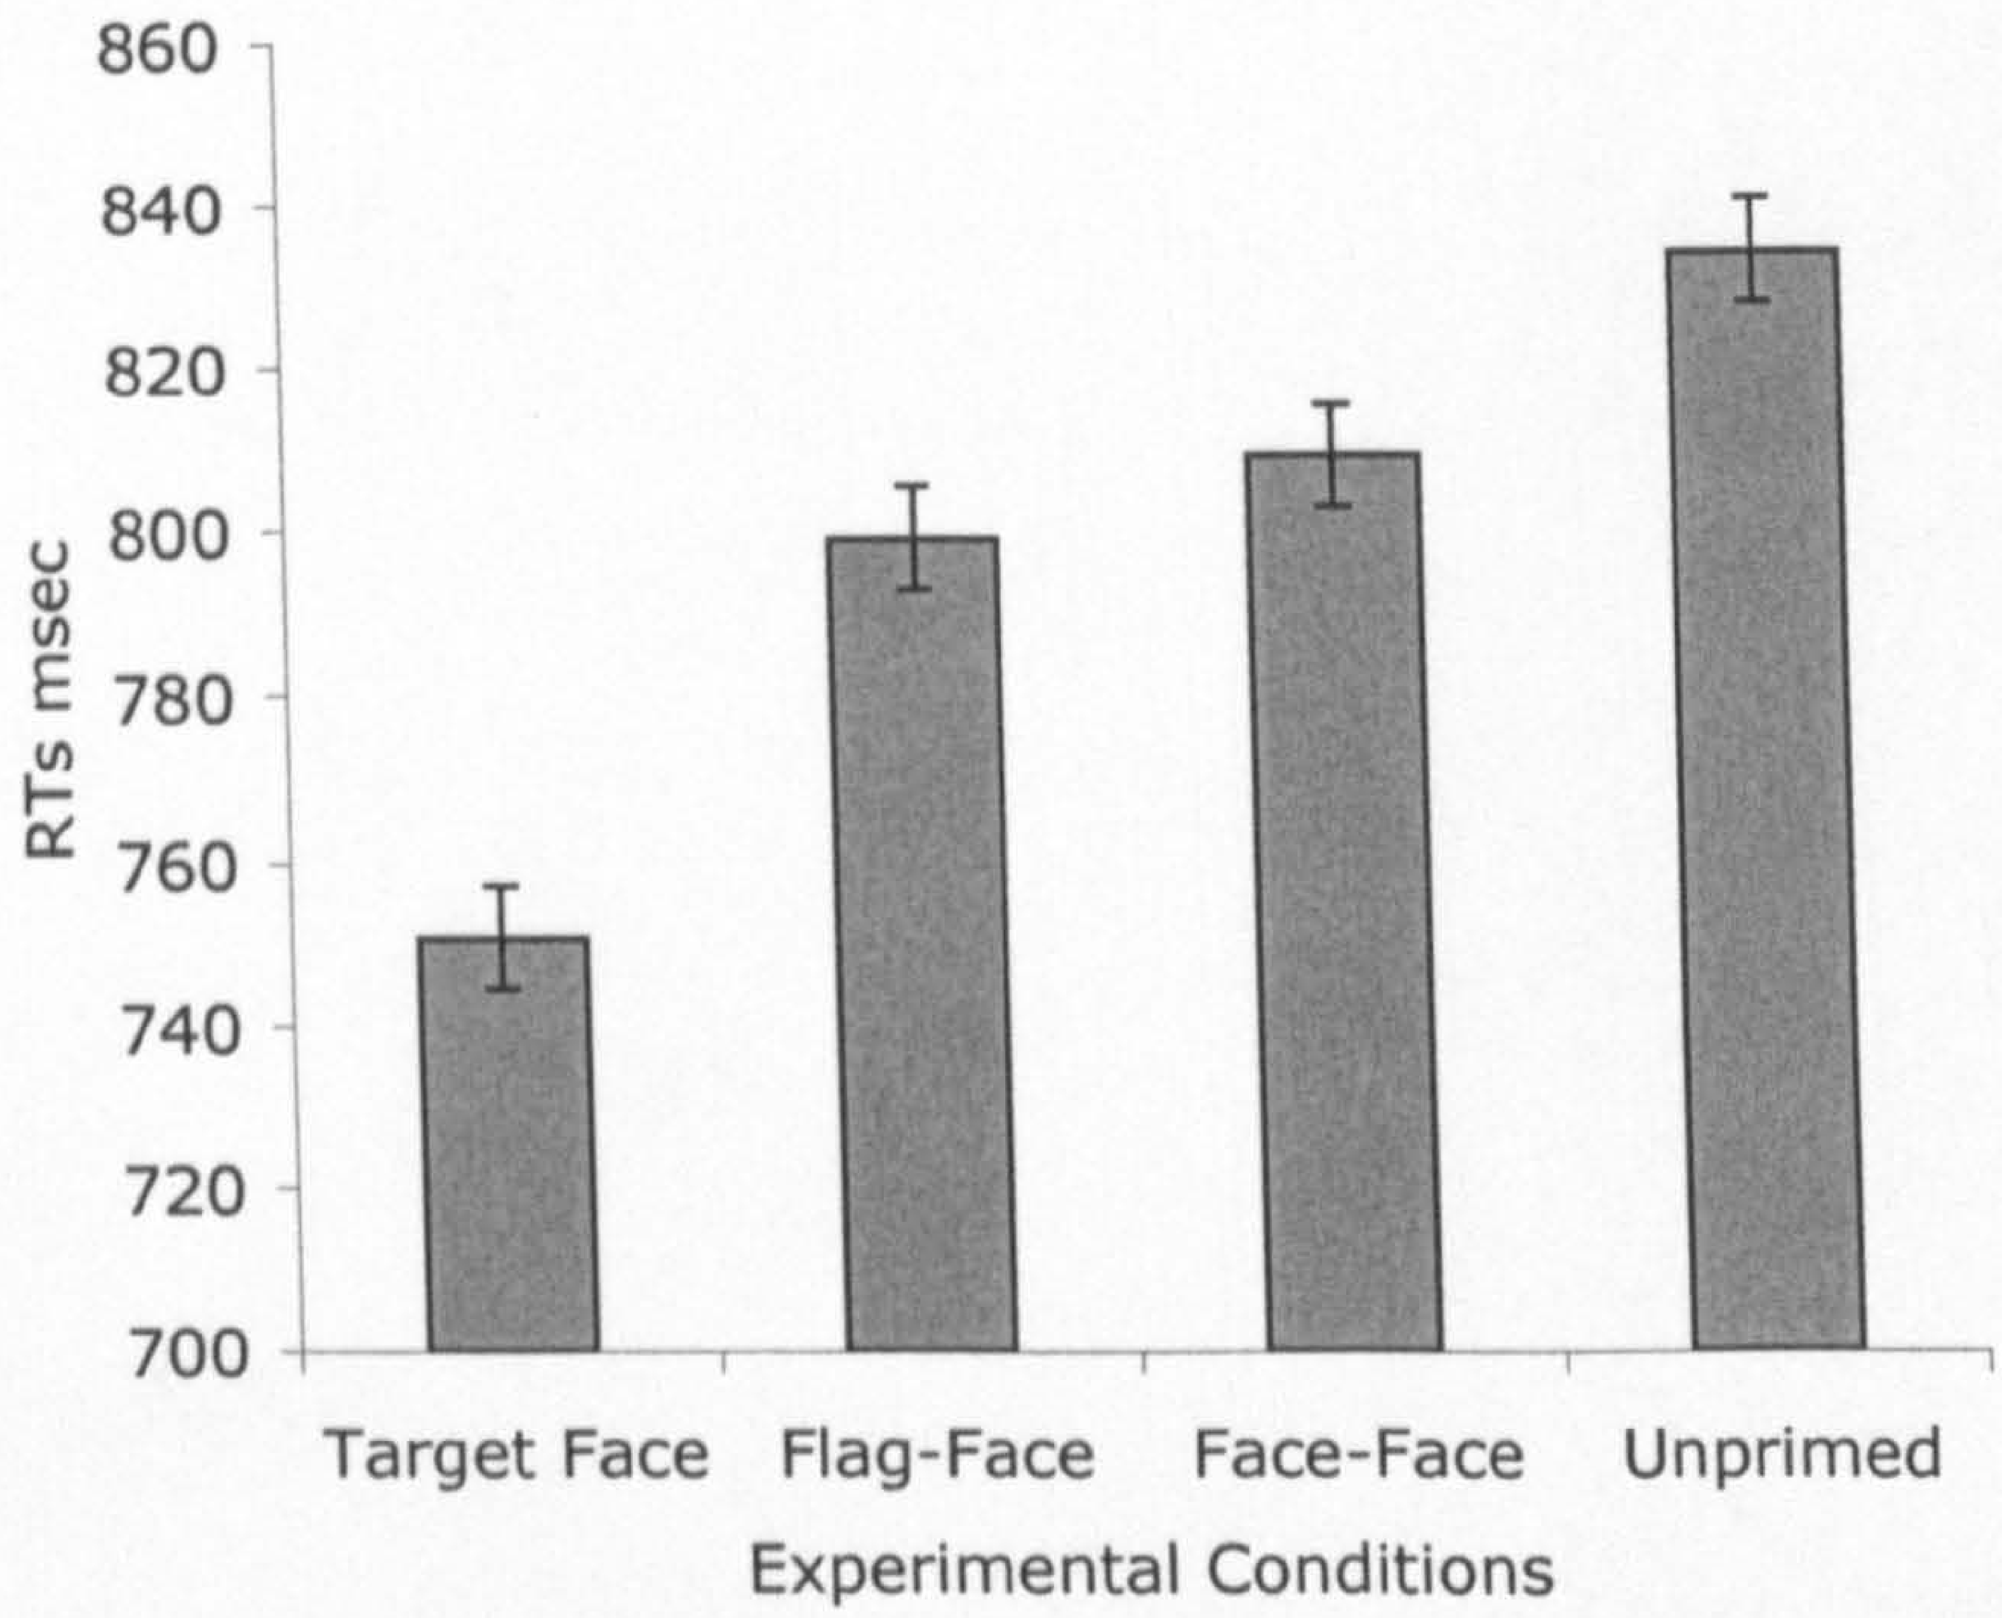

Figure 4.3 Mean correct responses (n=36) to famous faces in the test phase of Experiment 10. Performance is shown as a function of prime type; *Target Face*, *Flag-Face*, *Face-Face*, and *Unprimed*. Standard error bars are shown.

A one-factor within-subjects ANOVA (*Target Face* versus *Face-Face* versus *Flag-Face* versus *Unprimed*) showed an effect of prime type, $F(3,105)=30.46$, $p<.01$. As in Experiment 9, Tukey's HSD test ($p<.01$) showed more priming in the *Target Face* condition than in the *Face-Face* condition, the *Flag-Face* condition and the *Unprimed* condition. This time however, peripheral faces produced significant repetition priming whether presented with a flag target or a face target (*Flag-Face* versus *Unprimed*, $p<.01$; *Face-Face* versus *Unprimed*, $p<.05$). As expected, these priming effects were numerically larger in the *Flag-Face* condition, but this difference (of 10 msec) did not reach significance.

Error rates were analysed as for the RTs. Incorrect responses were made on 6.0% of *Target Face* trials, 9.9% of *Flag-Face* trials, 10.3% of *Face-Face* trials, and on 11.3% of *Unprimed* trials. ANOVA showed a main effect of prime type, $F(3,105)=6.05$, $p<.01$, reflecting fewer errors in the *Target Face* condition in comparison with each of the other conditions (versus *Flag-Face*, $p<.05$; versus *Face-Face*, $p<.01$; versus *Unprimed*, $p<.01$). No other comparisons were significant.

Discussion

Experiment 9 sought to determine whether famous face distractors are subject to cross-image repetition priming while a fixated target is classified. The results showed that distractor faces produce priming, but only during nonface target processing. By contrast, distractor priming was eliminated during face target processing. Experiment 10 repeated this design with identical stimuli at prime and test. As before, repetition priming was obtained for face targets. However, unlike Experiment 9 reliable repetition priming was now also observed for face distractors regardless of target type. These results have several implications. Although repetition priming is a well-researched area, these experiments demonstrate a dissociation between cross-image and same-image repetition priming of distractor faces. Moreover, despite converging claims from Experiments 4-9 that face processing is capacity limited to just a single face, from the present data it appears that a second face in a display, in the form of a task-irrelevant distractor, is nonetheless processed by the visual system during face target processing. How can these different results be reconciled? The finding that repetition priming survives changes in image has repeatedly been attributed to changes within the system that processes the perceptual identity of a face. This system is clearly capable of a great degree of generalization, since familiar faces are recognized even under difficult conditions (see e.g. Bruce, Henderson, Newman & Cowan, 2001). Therefore, the absence of cross-image distractor priming during face target processing in Experiment 9, suggests that the distractor faces were not recognized. As the same design was applied in the present experiment, one can assume that the face distractors of this condition were not recognized here either. However, since these faces did show same-image priming, the question remains to what extent they were processed.

Repetition priming is consistently greatest when identical face images are used. In the introduction to this experiment, it was suggested that a same-image priming advantage might at least partially arise from additional facilitation at a general processing stage prior to face recognition. However, the existence of such a stage is a contentious issue. On the one side, negative priming effects have been observed for unfamiliar faces (Khurana, 2000; Ellis et al, 1990; Young et al, 1986b), and for novel object shapes (DeSchepper & Treisman, 1996). Because subjects could not have maintained stored representations of these stimuli prior to the experiments, it seems unlikely that these effects occurred within the recognition system. Yet, if the face distractors of the *Face-Face* condition were processed similarly to unfamiliar faces and novel objects, then should they not also have produced negative rather than repetition priming? Crucially, Khurana (2000) found that unfamiliar faces produce repetition priming when the face stimuli are contrast-inverted or disrupted with high frequency noise between prime and test.

DeSchepper & Treisman (1996) also observed facilitation when the size of novel object shapes was changed between prime and test. Moreover, these priming effects were long-lived, indicating that they did not simply reflect transient activation patterns within early visual pathways, but that shape information was stored long term. Note that stimulus size was also manipulated in the present experiments (3.6 x 4.5 cm at prime versus 6.0 x 7.5 cm at test), and that the same-image priming effects survived across a considerable number of intervening trials between the presentation of the faces at prime and at test phase. The observation of repetition priming in the *Face-Face* condition in Experiment 10 therefore appears consistent with previous experiments using unfamiliar faces and novel shapes, and suggests that these effects might be located at a general processing stage, perhaps akin to a pictorial memory store outside the face recognition system.

A potential problem for this explanation is the observation that repetition priming is not observed onto sex or expression decisions when the same faces are used at prime and test (Ellis et al, 1990). However, since repetition priming enhances the speed at which stimuli are processed over successive presentations, visual stimuli that are frequently encountered should operate closer to the limit at which a stimulus can be processed. Sex and expression decisions to faces are categorical, and the number of these categories is extremely limited in comparison to the seemingly infinite number of faces encountered in everyday life (e.g. Etcoff & Magee, 1992). Moreover, due to the dynamic nature of faces, expression judgements may be made with much greater frequency than identity and sex decisions. Thus one might expect less priming for expression than for sex information, and less priming for both these types of information than for face identity. In line with this reasoning, sex and expression judgements are already relatively fast, and although repetition priming is usually not found onto these decisions (but see Goshen-Gottstein & Ganel, 2000; McNeill & Burton, 2003), more often than not sex decisions show small non-significant priming patterns (Ellis et al, 1990; McNeill & Burton, 2003). Consequently, the absence of repetition priming onto sex and expression decisions does not rule out the existence of a general processing component in same-image priming. The experiments in this chapter indicate that such same-image priming might be dissociable from priming within the face recognition system. However, at present Experiment 9 provides the only evidence that processing a face eliminates cross-image repetition priming of another face. Therefore, before reaching this strong conclusion it is necessary to replicate the key findings from Experiment 9, in particular the absence of cross-image repetition priming in the *Face-Face* condition. This is the purpose of the next experiment.

Experiment 11

The purpose of the final experiment of this chapter was two-fold. The first aim was to replicate the extinction of cross-image distractor priming that was observed during face target processing in Experiment 9. The second aim was to explore the visual properties that may be required to elicit face processing limits. A striking aspect of the preceding experiments, including Chapter 3, is that classifying nonface targets was never sufficient for eliminating distractor processing at and beyond the level at which faces are recognized. However, although the preceding experiments used nonface targets for which an equivalent face classification task could be made, none of these stimuli resembled any of the physical attributes of faces. Therefore, these findings do not imply that other 'face-like' visual stimuli are not capable of eliciting face processing limits.

A number of researchers have shown that even simple face-like schema may automatically activate face processing. For example, Suzuki & Cavanagh (1995) observed a cost in feature search for a single up-turned arc among down-turned arcs, when the arcs were arranged in sets of three in a face-like configuration. Similarly, Mack, Pappas, Silverman & Gay (2002) found that a happy face icon, consisting of an outer circle and two round dots above an upwards-arched line, is detected more frequently in a crowded display than a scrambled stimulus consisting of the same features. If these properties are sufficient for engaging face processes, then face-like stimuli might also be capable of eliminating cross-image face distractor priming. As a consequence, it might be possible to isolate some of the visual properties that determine face processing limits.

To examine this possibility, the current experiment used images of frontal and three-quarter rear and full rear view photographs of American and British cars as targets. Human faces have two horizontally separated eyes positioned vertically above a wide central mouth, in the context of a general left-right symmetry. Similarly, car fronts possess two horizontally separate and often circular headlamps above a wide central grille, again in the context of a general left-right symmetry. Like the schematic stimuli of previous studies (e.g. Mack et al, 2002; Suzuki & Cavanagh, 1995), these features occasionally provide a striking resemblance to faces. Therefore, similar to photographs of genuine face targets, cross-image priming of distractor faces could be greatly reduced or extinguished by photographs of car fronts. On the other side, three-quarter rear views and full back views of cars often possess similar features to car fronts, such as rear-lights, and full car backs also exhibit a great degree of symmetry, but these views usually do not appear to resemble faces to the same extent. Thus, car fronts might engage face processing limits when car backs do not. Experiment 11 addressed this possibility by measuring cross-image face priming under four conditions: i) from task-relevant face targets (the *Face Target* condition); ii) from face distractors presented alongside a relevant face target (the *Face-Face* condition); iii) from face distractors presented alongside images of car fronts (the *Front-Face* condition); and iv) from face distractors presented alongside images of three-quarter and full rear views of cars (the *Back-Face* condition). As before, an *Unprimed* condition was also included as a baseline.

Method

Pilot Study Before describing the main experiment, the perceived nationality and the perceived faceness of car images was examined directly in two rating tasks. One aim of this preliminary study was to ensure that the car images provided an intuitive analogue in terms of nationality to the famous face stimuli at prime phase, even for participants with limited knowledge about cars. The second more important aim was to verify the hypothesis that the perceived faceness of car fronts is greater than car backs. In the nationality rating tasks, 10 participants (all undergraduate students from the University of Glasgow) were instructed to sort decks of 40 car images into two piles. These images consisted of 20 typical American cars and 20 typical British cars, exactly half of which depicted car fronts with the other half depicting cars from three-quarter and full rear view (see Figure 4.4 overleaf).

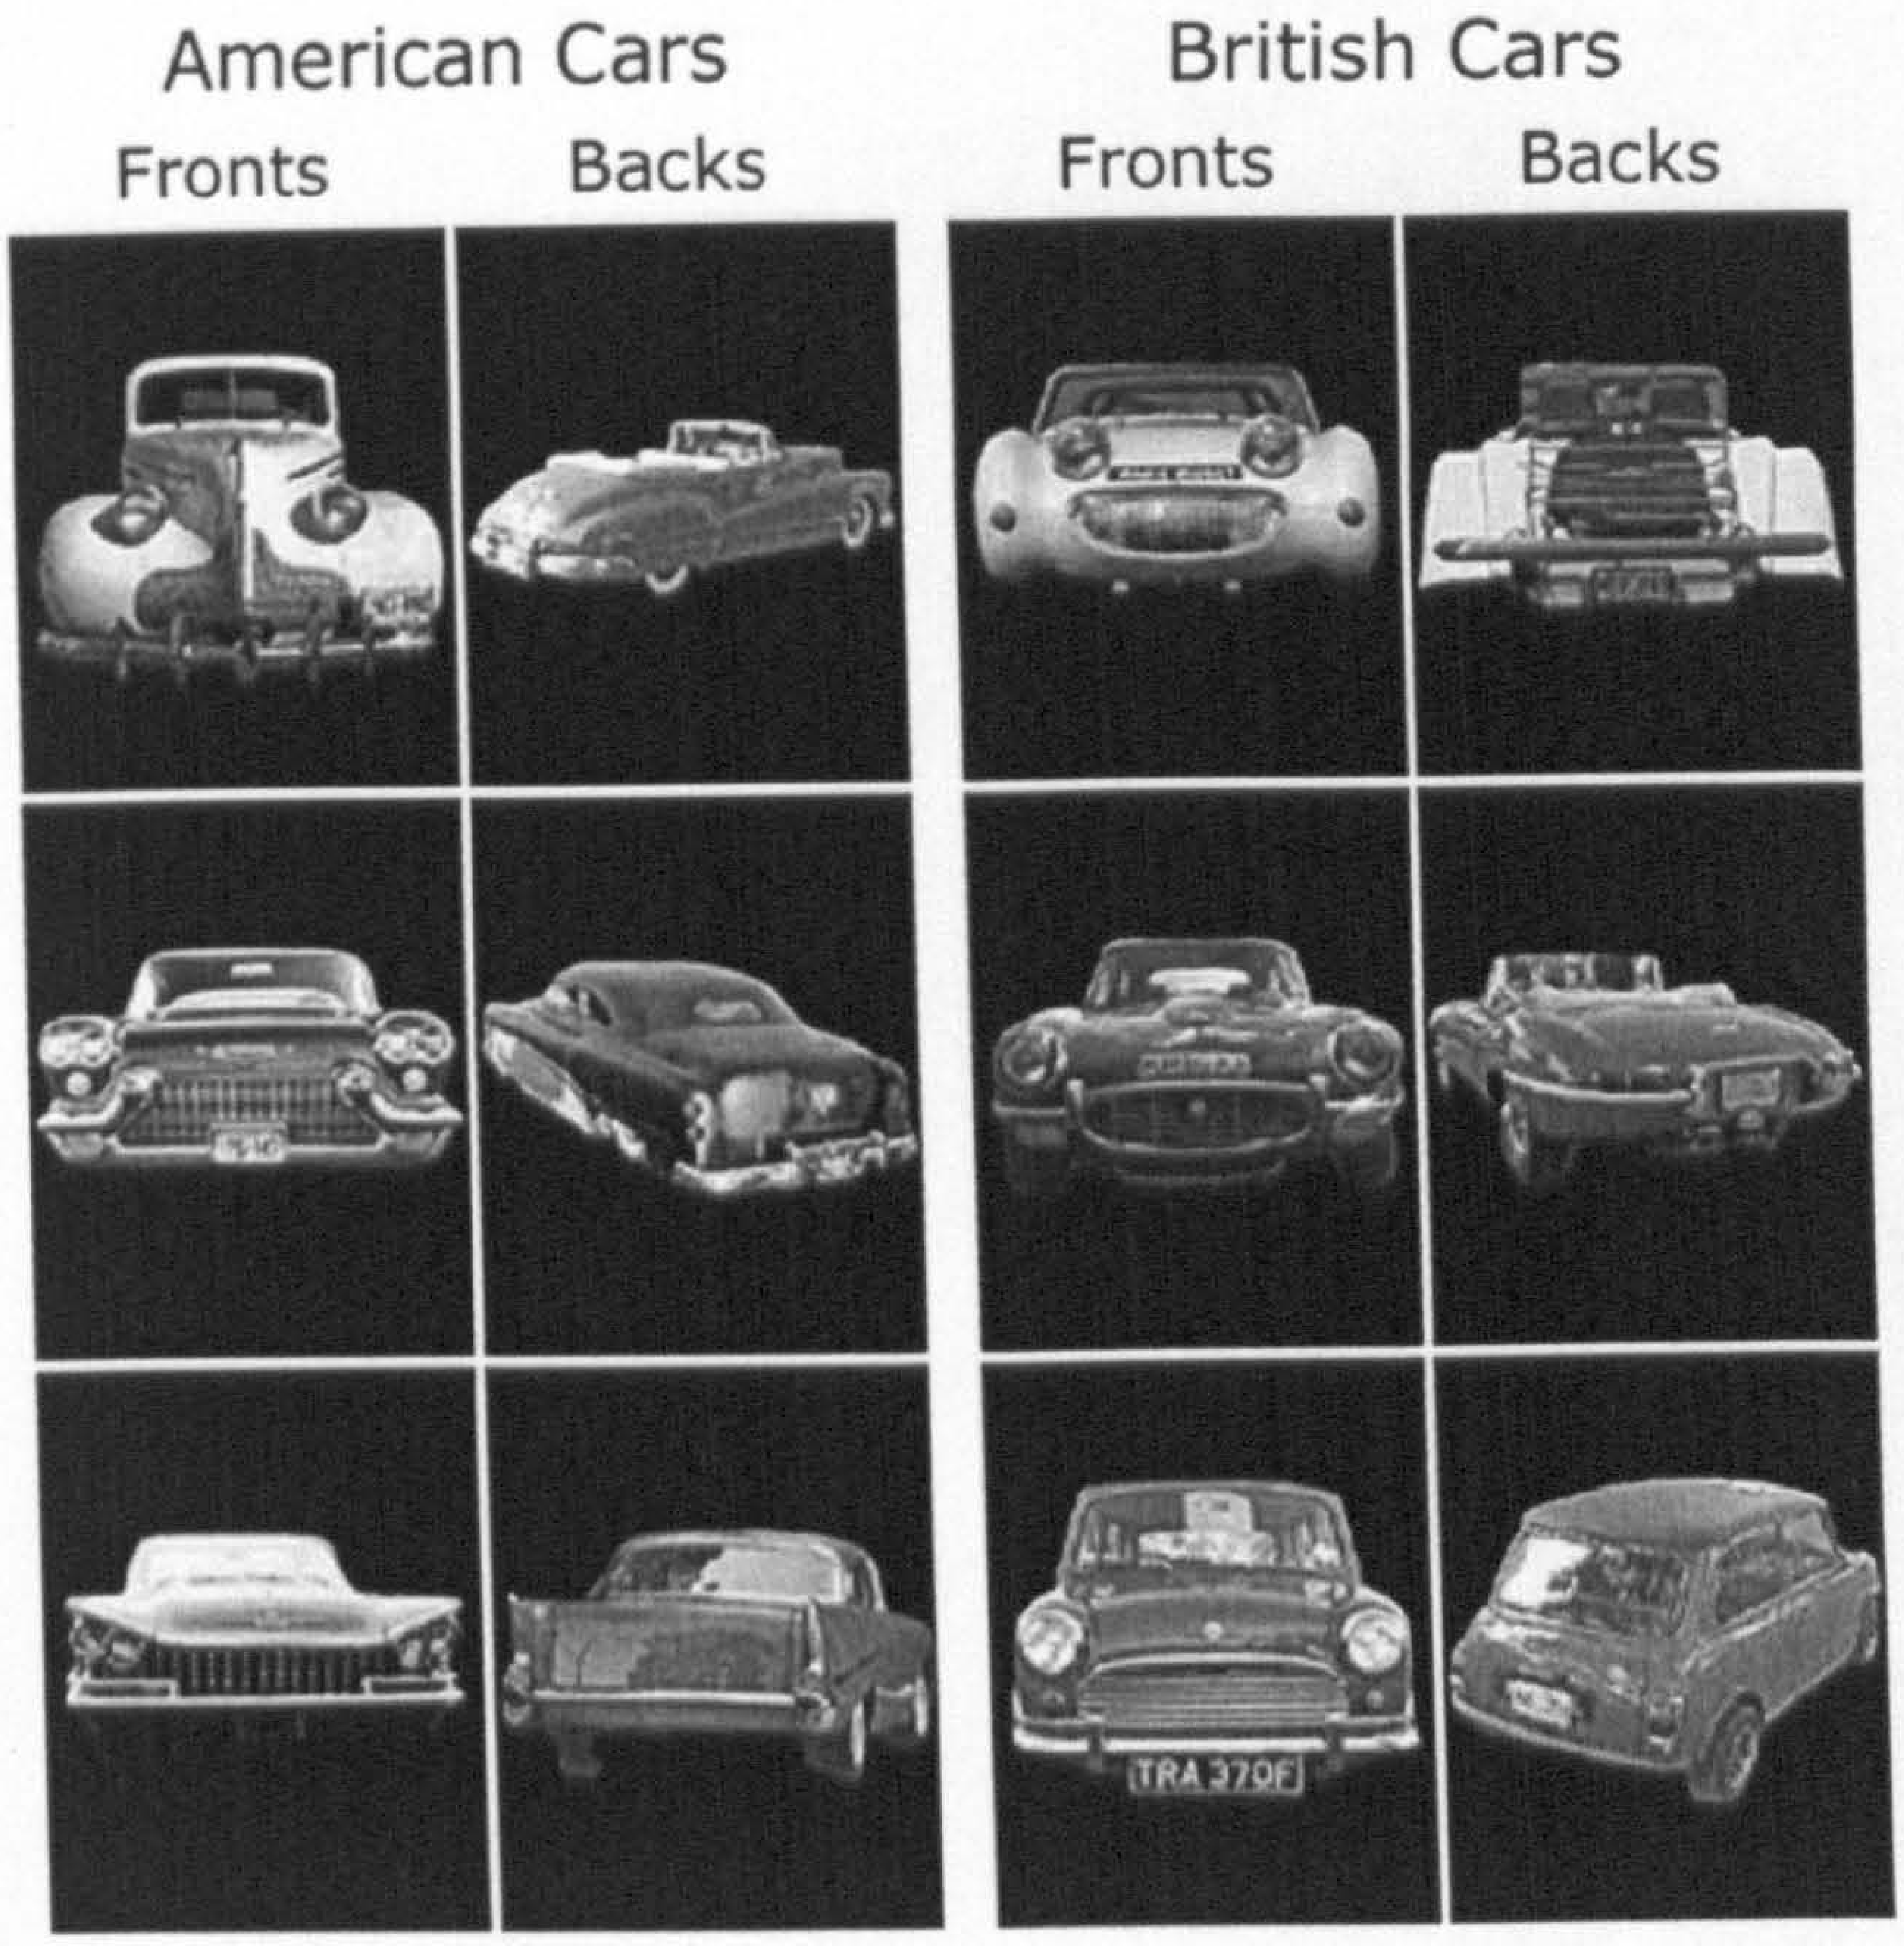

Figure 4.4 Example displays of the car stimuli used in Experiment 11. The stimuli consisted of photographs of face-like car fronts and photographs of face-unlike three-quarter and full back views of American and British automobiles.

The sorting criterion was whether cars looked American or British. As a classification strategy, participants were encouraged to guess when uncertain regarding the correct answer, but were told that American cars generally look bulkier and had big chrome grills and tail fins, in comparison to the more modest, sportier looking British cars. These instructions proved successful as 85% of all cars fronts (American, 85.3% versus British, 83.8%) and 73% of other car images (American 73.5% versus British, 73.3%) were categorized correctly. Subsequently, these images were ranked according to faceness, with a rank of 1 given to the most and 40 to the least face-like looking car. Results were consistent with the hypothesis that car fronts possess greater faceness than other car images, with average rankings of 10.6 for car fronts (std 3.9) and of 30.1 (std 3.9) for the remaining images.

Subjects In the main experiment, 45 British students from the University of Glasgow, aged 18-30 years, volunteered to participate for a small fee. All subjects reported normal or corrected vision and had not participated in any of the previous experiments.

Design & Stimuli In the prime phase, the 40 car images and photographs of 100 celebrities (50 American & 50 British, see Appendix C) were used to construct stimulus displays consisting of a central target and a flanking face distractor, presented 1.0 cm (1.0° of VA) to the left or right of the nearest target contours. All images were presented in greyscale on a black background at a size of 3.6 cm x 4.5 cm (3.4° x 4.3° of VA). Target-distractor pairings were counterbalanced so that they were equally likely to be of same or different nationality. Combining 20 car fronts, 20 car backs, and 20 famous faces with a famous face distractor resulted in 60 target-distractor displays. Overall, this involved 80 of the 100 face images. The remaining celebrities' faces were reserved as unprimed controls for the test phase. As before, all faces were rotated across these conditions over the course of the experiment, so that each face appeared in each of the conditions an equal number of times.

At test phase, 100 novel images of the same celebrities' faces and 100 unfamiliar faces were used in a familiarity task (famous vs. unfamiliar). Images were presented at a size of 6.0 cm x 7.5 cm (5.7° x 7.3° of VA), one at a time, on a black background. The experiment was run and responses were recorded on an Apple

Macintosh computer using PsyScope 1.2.5 software. Viewing distance was fixed at 60 cm by a chinrest.

Procedure The procedure for the prime phase was almost identical to the preceding experiments, with each trial consisting of a fixation cross for 750 ms, a target-distractor display for 200 ms, and a blank screen until a response was registered. Subjects again made speeded judgements concerning the targets' nationality by pressing the "D" key for American targets or the "L" key for British targets. If no response was made within 2.5 seconds of stimulus onset, the next trial was initiated. For car targets, participants were not expected to possess particular car expertise, but were instructed to classify bulky-looking models with large chrome grills or tail fins as American and more modest, sportier looking cars as British. Subjects were encouraged to guess if they were uncertain regarding the correct answer. All subjects underwent a short example block of 18 trials, made up from an additional 6 car and 12 face images, which were not encountered in the experimental block. This was followed by a critical block of 60 randomly intermixed trials.

The test phase was the same as for the preceding experiments, except that each subject now completed four blocks of 50 randomly intermixed trials. Thus, each block consisted of 10 more faces (5 famous, 5 unfamiliar), to accommodate the face stimuli from the additional priming condition.

Results

Priming Phase Incorrect responses were discarded and the mean correct RTs and error rates were then calculated for responses to face and car targets. The averages of these means across subjects for each target type were: face targets 946 msec (error rates 10.1%), car-fronts 882 msec (error rates 17.7%), and car-backs 906 msec (error rates 21.9%).

Test Phase Incorrect responses and RTs exceeding 2 seconds were excluded from the analysis (less than 1% of all correct responses). The mean RTs and error rates across subjects are shown in Figure 4.5 as a function of experimental condition.

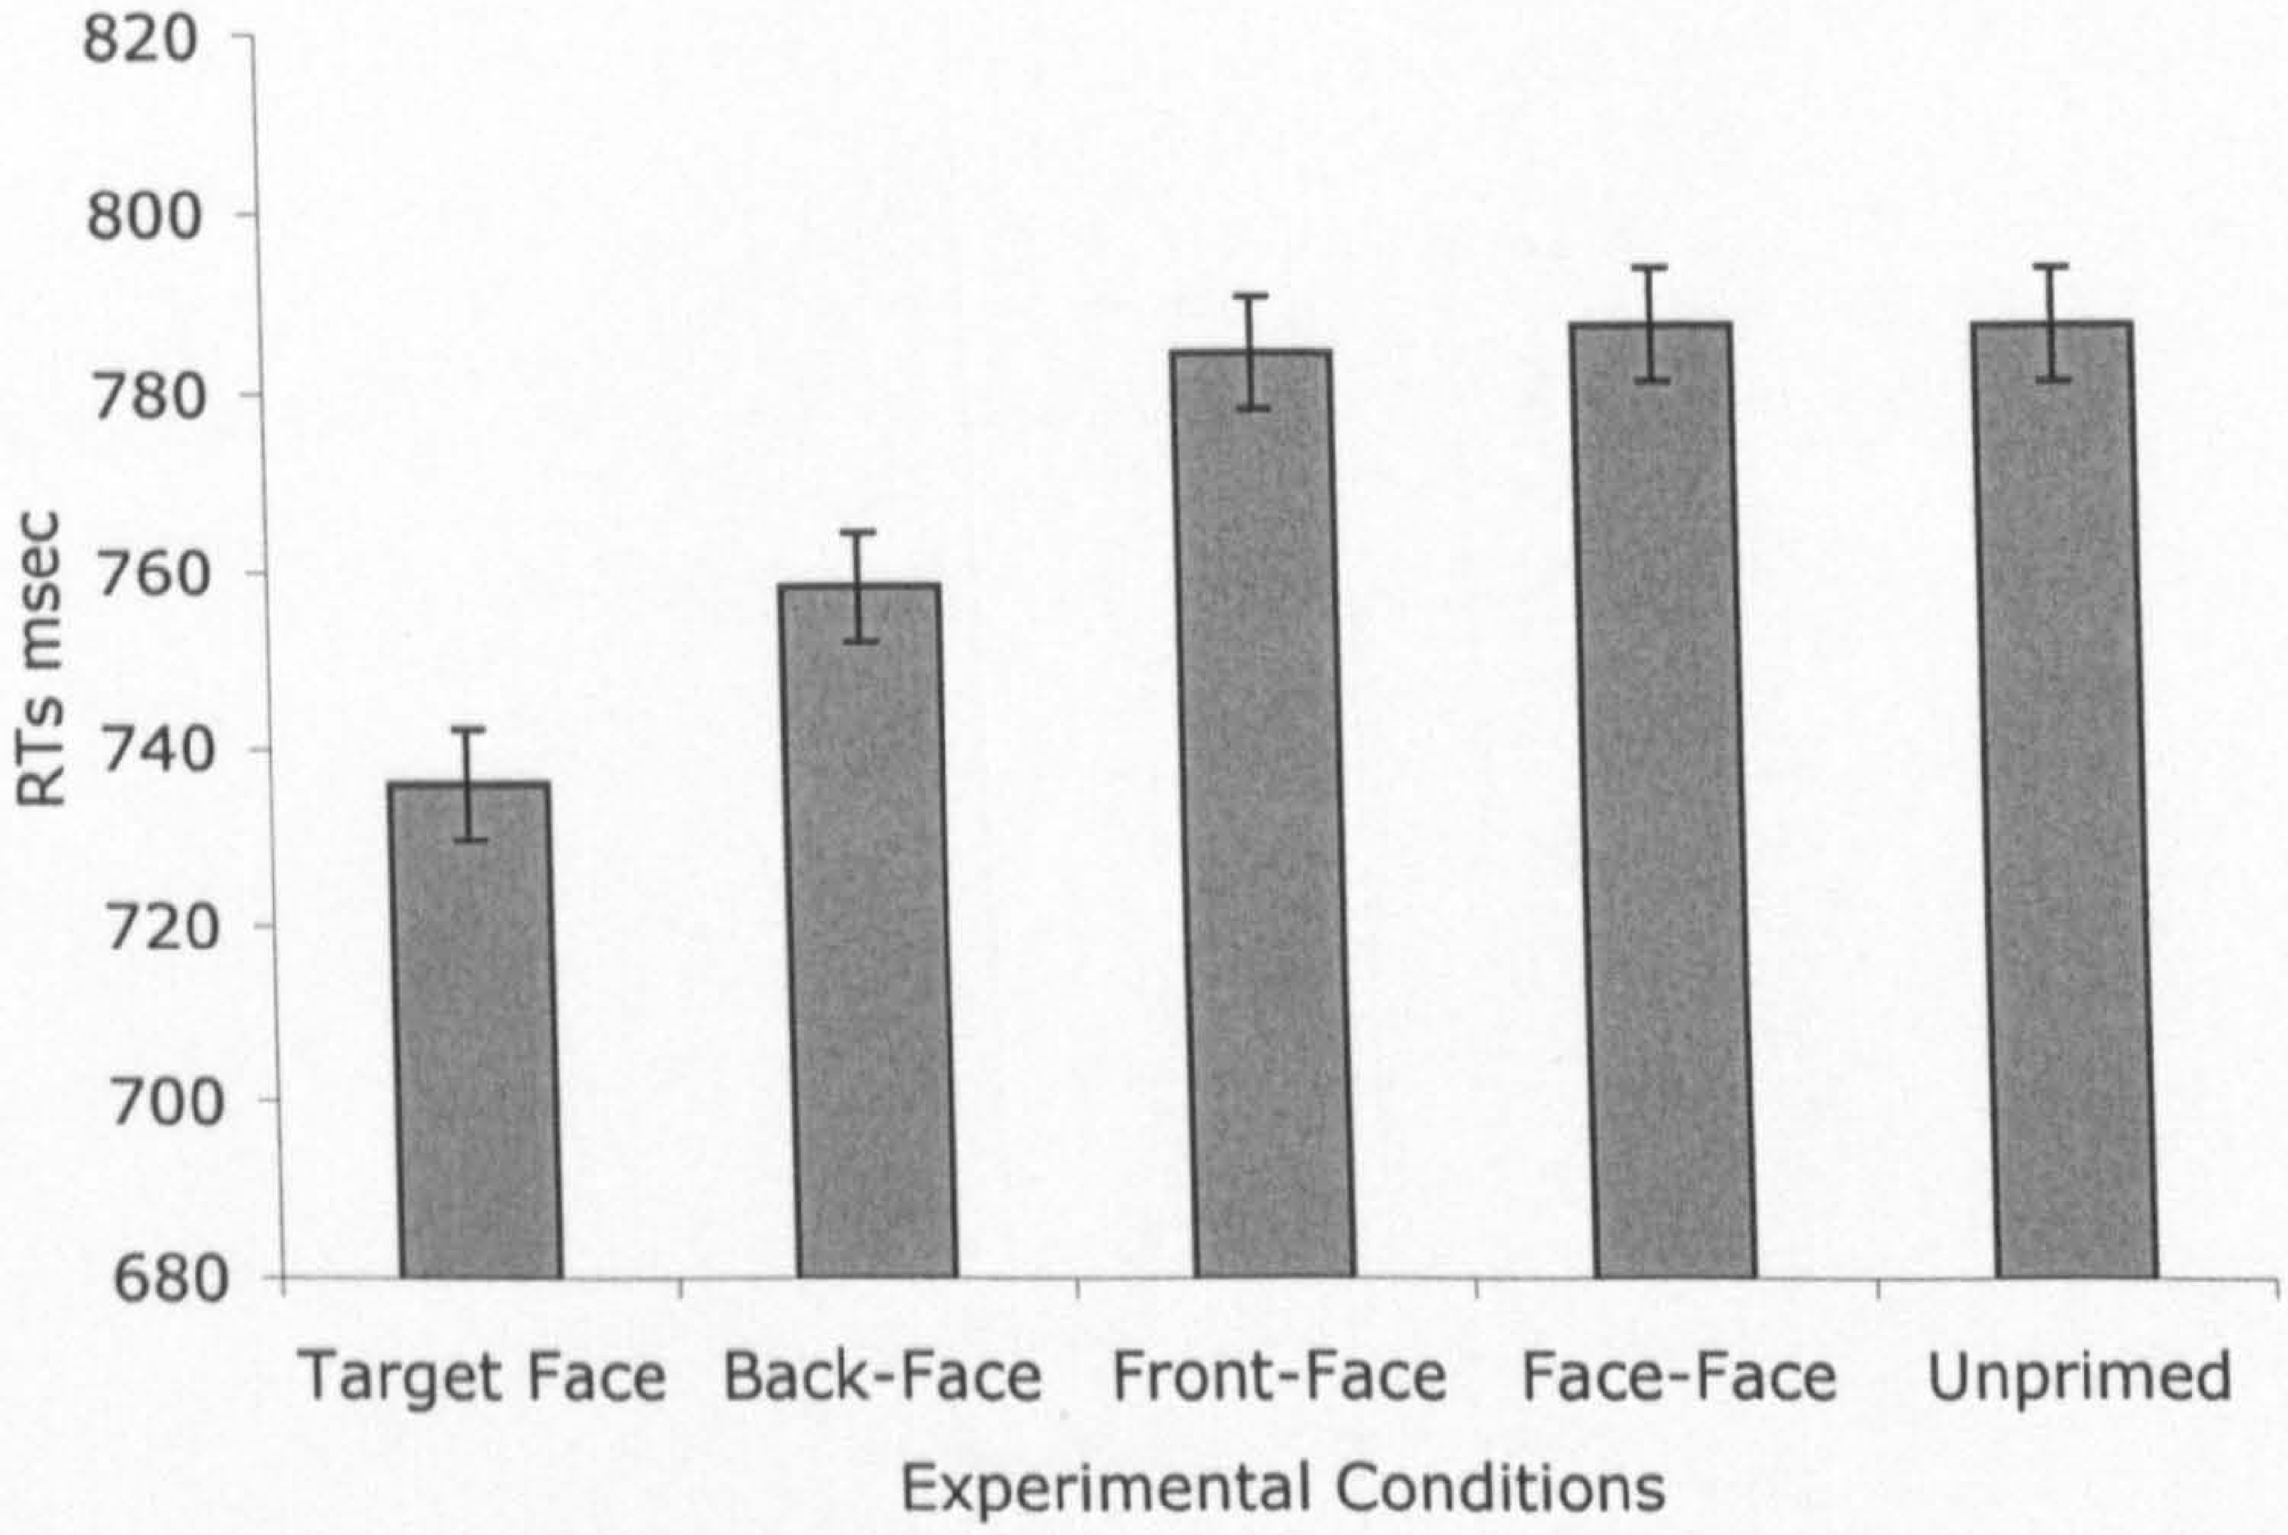

Figure 4.5 Mean correct responses (n=45) to famous faces in the test phase of Experiment 11. Performance is shown as a function of prime type; *Target Face*, *Back-Face*, *Front-Face*, *Face-Face*, and *Unprimed*. Standard error bars are shown.

A one-factor within-subjects ANOVA of these conditions (*Face Target* versus *Back-Face* versus *Front-Face* versus *Face-Face* versus *Unprimed*) showed an effect of prime type, $F(4,176)=13.60$, $p<.01$. *Target Faces* showed significant repetition priming in comparison with the *Front-Face* condition, the *Face-Face* condition, and the *Unprimed* condition (all Tukey's HSD, $p<.01$). Although RTs to

Target Faces were also faster than to *Back-Face* primes, this difference did not reach significance. Face distractors in the *Back-Face* condition also showed reliable priming in comparison with the *Front-Face* and the *Face-Face* condition and the *Unprimed* faces (Tukey's HSD, $p<.05$). However, the *Front-Face* condition and the *Face-Face* condition showed no repetition priming in comparison with the *Unprimed* condition.

Error rates were analyzed in the same way as the RTs. Errors were made on 10.8% of *Target Face* trials, 17.2% of *Back-Face* trials, 13.9% of *Front-Face* trials, 17.0% of *Face-Face* trials, and on 16.3% of *Unprimed* trials. A one-factor within-subjects ANOVA showed an effect of prime type, $F(4,176)=5.58$. $p<.01$, reflecting lower percentage errors to *Face Target* primes than to *Back-Face* primes, *Face-Face* primes and *Unprimed* faces (Tukey's HSD, $p<.05$). No other comparisons were significant.

Discussion

Significant repetition priming was obtained from target faces and from distractor faces that flanked a car back. Since the prime phase and the test phase presented different images of the famous faces, these effects imply that the faces were processed to the point of identification in these conditions. In contrast, peripheral faces presented alongside a face target did not give rise to repetition priming. This replicates the main finding of Experiment 9, again indicating that processing the target face blocks identification of the accompanying flanker face, and also provides a further conceptual replication of the face processing limits that were first specified in Chapter 3. However, the most important finding of this experiment is that priming from a peripheral face was also eliminated when the target was a car front. In other words, the car fronts, which were rated as face-like by observers, behaved like face targets in blocking identification of a peripheral face. In contrast, the car backs, which did not resemble faces, behaved like the flag targets in Experiment 9, allowing processing of the distractor faces to proceed unfettered. Thus this pattern appears consistent with the notion that any sufficiently face-like stimulus engages limited face processing resources, whether it is a genuine face or a 'false positive'.

However, there might be alternative explanations for these results. For example, although car fronts and car backs presumably possessed similar levels of visual complexity and homogeneity, they were not matched for low-level image characteristics, such as their component spatial frequencies. Yet, given the sheer number of car stimuli used, it would seem extraordinary if low-level features could account for this dissociation. Another possibility is that subjects perceive transient displays of car fronts as an oncoming threat. Threatening words, pictures, and faces have all been shown to capture processing resources in trait- and state-anxious individuals to the detriment of other stimuli, in comparison to their non-threatening counterparts (e.g. Amir, Elias, Klump & Przeworski, 2003; Fox, Russo, Bowles & Dutton, 2001; Yiend & Mathews, 2001). If images of oncoming cars are perceived as equally threatening, this might contribute to distractor face extinction. However, the car stimuli did not contain motion cues and it seems unlikely that the participants, all undergraduate volunteers, were clinically anxious. Moreover, there is now considerable evidence that task-relevant nonface processing does not affect the perception of irrelevant face distractors (e.g. Boutet et al, 2002; Jenkins et al, 2003; Lavie et al, 2003). Consequently, it seems most plausible that the present findings require an explanation in terms of the perceived faceness of the car stimuli. In addition to the interference data from Chapter 3, where analogous capacity limits were never observed for nonface comparisons (forenames, famous names, flags) or for any of the face-nonface pairings, Experiment 11 thus provides some further, innovative support for a unique role of faceness in the processing limits described in this thesis.

Overall, the results of Chapter 4 are consistent with the notion of a limited face processing capacity, such that only a single face can be processed. When face processing resources are loaded by a task-relevant face target, identification of a task-irrelevant distractor face, as indicated by cross-image repetition priming, is prevented. Access to these limited resources is only granted to distractor faces during the processing of nonface targets, provided that they are sufficiently face-unlike to leave face processing resources relatively intact. In fact, the present findings suggest that even seemingly artificial face stimuli, such as car-fronts possess the ability to elucidate face processing limits. Although the data provide no clues to the specific nature of the face information portrayed by car fronts, it seems implausible that it consists of more than a general "eye-mouth" configuration. Future variations of this paradigm, such as manipulations of the visual attributes of the target stimulus, have the potential to further clarify the fate of face-like stimuli in the visual system, and hence the selectivity of the face processing system. Finally, Chapter 4 also found same-image priming of distractors faces when they were alongside a face target, which suggests that these distractors were still subject to early image processing. Note that this could still include processes such as face detection, which appears to be dissociable from face recognition (de Gelder & Rouw, 2001). This possibility is not ruled out by the present experiments and is another subject for future research. For now, I turn to an issue that may be implicated in the results of Chapters 3 & 4, namely attention biases to faces.

Chapter 5 Disengagement and Engagement of Attention from Faces and Nonface Objects

Introduction

This final empirical chapter continues from face processing limits to a related theme, namely factors that might influence face processing in multiple-item displays. The experiments in Chapters 3 & 4 found that while distractor face processing was eliminated during the classification of a face target, the same face distractors seemed to be unaffected by nonface targets. In fact, whenever a solitary distractor was used in Chapter 3, the largest interference effects were consistently observed in the condition in which a nonface target was flanked by a face distractor (Experiments 4-6). Moreover, although nonface distractors interfered with face and nonface targets alike, these effects were generally smaller in the face-nonface condition. Thus, faces seemed to interfere more with task-relevant processing than other stimuli in these experiments, but were also subject to less interference. The question addressed in the current chapter is whether these differences could be attributed to an attentional bias for faces, which might allow faces to maintain a processing advantage in comparison with other visual stimuli.

Attention researchers have identified at least two processes during which a face bias could arise, disengagement of attention from one stimulus and engagement of attention by another (see e.g. Posner, 1980; Posner & Petersen, 1990; Posner, Snyder & Davidson, 1990; Theeuwes, 1991; Theeuwes, de Vries & Godijn, 2003). A disengagement bias refers to the ability of a focused-on stimulus to retain or to hold attentive resources. This occurs after the initial orienting to that stimulus and is characterized by less efficient processing of other stimuli within the visual field. Engagement, on the other hand, increases alertness to the spatial location of another stimulus and enhances its processing. This distinction between the ability to retain and to engage attention certainly seems consistent with Experiments 4-6, where faces appeared to have an advantage as both relevant targets and irrelevant distractors whenever they were paired with a nonface object, but has not always been applied to the face domain. Studies that have ignored this distinction generally argue that faces may have a propensity to engage or to capture attention (e.g. Mack, Pappas, Silverman & Gay, 2002; Ro, Russell & Lavie, 2001; Vuilleumier, 2000), whereas others have attributed attention biases to a difficulty in disengagement, although this may only occur under specific conditions (e.g. Fox, Russo & Dutton, 2002; Fox, Russo, Bowles & Dutton, 2001).

A variety of sources have suggested that faces may have an advantage in capturing or engaging attention. For example, Vuilleumier (2000) reports three neuropsychological patients suffering from left-sided spatial neglect following brain damage to the right hemisphere, who were presented with line-drawn faces, scrambled faces, names, and meaningless shapes in either the left or the right or both visual hemifields. Remarkably, patients were less likely to report left-sided shapes, scrambles, and names when they were accompanied by a face stimulus in the right hemifield. Moreover, left-sided faces were reported more frequently than any other type of stimuli, suggesting that faces not only have an advantage in capturing attention when competing with other stimuli, but also in overcoming visual extinction. However, because Vuilleumier's (2000) findings are based on just three visual neglect patients and on fairly artificial face stimuli, the generalizability of these results is somewhat limited. In a study of neurologically normal patients but with similarly unrealistic stimuli, Mack et al (2002) also found that happy face icons are reported more often from a stream of visual stimuli, presented at a rate of 75ms/item, than outlines of Christmas trees and inverted faces (see also Mack & Rock, 1998). Notably, Ro et al (2001) also observed a detection advantage for photographs of *real* faces, with a flicker paradigm. In such tasks, participants are typically shown displays containing several stimuli, alternating with blank screens, and are asked to detect a change in one of them. Ro et al (2001) found that changes were detected more quickly in faces than in nonface objects from a range of categories (food, clothes, musical instruments, appliances and plants), which indicates that real faces may also have a special capacity to attract attention.

However, these findings have not gone unchallenged. Palermo & Rhodes (2003) re-examined Ro et al's (2001) claims, to determine whether their results could be explained in terms of an "odd-one-out" rather than a face processing advantage. This was based on the observation that Ro et al (2001) only ever presented one face among a range of nonface objects, which could have differed from faces in a number of ways (e.g. living versus non-living stimuli). In support of this line of reasoning, Palermo & Rhodes (2003) found a similar change-detection advantage for a single nonface target when embedded among several faces, leading them to suggest that uniqueness may be more important than faceness in change detection. However, Palermo & Rhodes (2003) were unable to produce a clear replication of the change detection advantage for faces reported by Ro et al (2001), even under seemingly identical conditions. Either way, these results suggest that faces are no more able to attract attentional resources in a change detection task than other objects.

An "odd-one-out" explanation could also account for Mack et al's (2002) findings with happy face schema, as these were presented in a stream of nonface objects with little social or biological meaning (e.g. line-drawn sailing boats, heart shapes, telephones). Indeed, in a separate experiment Mack et al (2002) found that observers were also more likely to detect their own name in a stream of letter strings than words, leading them to suggest that the ability to capture attention might actually depend on meaningfulness, rather than faceness. In line with this reasoning, several other studies indicate that people may be particularly slow to disengage from a range of visual stimuli, including faces, depending on their emotional connotation and the emotional state of the observer. In one such study, Fox et al (2001) asked subjects to detect a dot probe, the onscreen location of which could be correctly or incorrectly cued by emotionally threatening or neutral stimuli, consisting of either faces or words. While cue validity influenced responses in normal subjects, with slower responses on incorrectly-cued trials, threat value had no effect on dot probe detection times. However, responses were markedly slowed in highly anxious people on invalidly-cued threatening trials, suggesting a disengagement bias to threatening material. Others have shown similar attention biases for negatively charged emotional faces in normal individuals (e.g. Eastwood, Smilek & Merikle, 2003) and anxious individuals (see e.g. Bradley, Mogg, Falla & Hamilton, 1998; Fox et al, 2002; Van Honk, Tuiten, Van den Hout, De Haan & Stam, 2001), but also for threatening pictures (Yiend & Mathews, 2001) and threatening words (Amir, Elias, Klumpp & Przeworski, 2003). Overall, these findings are at least partially consistent with claims that an attentional bias might depend on the meaning of a stimulus. However, it should be noted that none of these studies can address whether a *general* attentional bias for faces still exists, independent of any emotional connotations, as faces were never compared with other classes of stimuli within the same experiment.

In all, existing evidence that faces may be particularly strong competitors for attention is not convincing. Moreover, it is unclear whether any advantage might reflect an ability to engage or to maintain attentional resources, or even both. In order to address these issues, I devised a simple classification task in which subjects were required to attend to the colour of a central go/no-go signal before responding to the location of a vertical line target (i.e. left versus right of fixation) within the display. In the first three studies (Experiments 12-14) the go/no-go signal was either presented on a blank background, or superimposed on to-be-ignored face images or nonface distractors, and the line targets were positioned in the periphery of the displays, clearly separated from fixation. Note that identifying the colour of the go/no-go signal is a task thought to place minimal demands on attention (e.g. Treisman, 1993), which should make it impossible not to process other information presented at fixation (e.g. Lavie, 1995, 2000). Therefore, if faces are capable of retaining attention then this should result in an increase in target classification times on trials in which a face is presented in comparison with trials on which other stimuli are shown. In addition, if faces are also proficient at engaging attention then they should increase classification times even when they never appear in a task-relevant location. This was examined in a fourth study (Experiment 15), in which the line targets were moved within close proximity of fixation, and the to-be-ignored face and nonface objects now appeared clearly separated from the targets within the periphery of the displays.

Experiment 12

In Experiment 12, subjects were shown displays consisting of a coloured central go/no-go signal and two black lines, a vertical and a horizontal line, which were presented in the periphery, one to the left and one to the right of fixation. The subjects' task was to classify the position of the vertical line (left or right) on go trials, in which a green fixation dot appeared. These trials were complemented with occasional no-go trials, designated by a red fixation dot, in which a target-neutral response was required (to initiate the next trial). This go/no-go distinction was included to confirm that subjects were attending to the centre of the displays. Response times were measured under four conditions. In the *Blank* condition, the lines and fixation dot were presented alone on an otherwise blank background. In addition, these displays could contain to-be-ignored photographs of (1) upright faces, presented in the centre of screen behind the go/no-go signal (the *Upright Face* condition), (2) inverted faces (the *Inverted Face* condition), or (3) meaningful nonface objects, in this case items of fruits (the *Object* condition). The latter two conditions were intended to serve as nonface controls for the upright face stimuli. Inverted faces provide a perfect match for their upright equivalents in terms of spatial frequency, complexity, and stimulus homogeneity, but are perceived and recognized so poorly in comparison that it has repeatedly been suggested that they may be processed more similarly to objects than to face stimuli (see e.g. Farah, Wilson, Drain & Tanaka, 1995; Moscovitch, Winocur & Behrmann, 1997). In contrast to inverted faces, the fruit stimuli were not equated to the low-level visual properties of the face stimuli, but were included to provide a meaningful object comparison. If faces are particularly proficient at retaining attention, orienting attentive resources from the central go/no-go signal to the peripheral line targets should be less efficient on *Upright Face* trials, as should be seen from an increase in target classification times. On the other hand, if it is no more difficult to disengage from faces than from other stimuli, target classification times should be equivalent regardless of stimulus type.

Method

Subjects Ten postgraduates from the University of Utrecht, the Netherlands, volunteered to participate for free in this experiment, and a further ten students from the University of Glasgow took part in exchange for course credits or volunteered for free. Participants' ages ranged from 18-27 years, and all had normal or corrected to normal vision.

Design & Stimuli The experimental displays contained a central go/no-go signal in the shape of fixation dot of either green or red colour, which measured 0.2 cm in diameter (0.2° of VA at a viewing distance of 60 cm) and was flanked by a vertical and a horizontal line. These lines were presented in black at a size of 0.1 cm x 0.4 cm (0.1° x 0.4° of VA), and were positioned 4.8 cm (4.6° of VA) to the left and right of fixation (see Figure 5.1 overleaf). The position of these lines was counterbalanced throughout the experiment, so that each line occurred equally often in each location. Apart from the fixation dot and the line targets, these displays either remained blank or could contain one of three types of stimuli within the centre: upright faces, inverted faces, and images of fruits. The upright face stimuli consisted of high quality digital photographs of three female celebrities (Pamela Anderson, Marilyn Monroe, and Britney Spears), which were cropped to remove any extraneous background, rendered in greyscale and sized to 2.4 cm x 3.0 cm (2.3° x 2.9° of VA) using graphics software. These images were then duplicated and turned upside-down to produce a matching set of inverted faces. The fruit images consisted of three photographs of an apple, grapes, and a plum, which were manipulated in the same way as the upright face stimuli (see Figure 5.1).

Including the blank displays, combining each of these nine stimuli under each level of target location (left or right of fixation) and for go/no-go signals resulted in a total of 40 different stimulus displays. These displays were presented on an Apple Macintosh on a white background and responses were recorded using PsyScope 1.2.5 software.

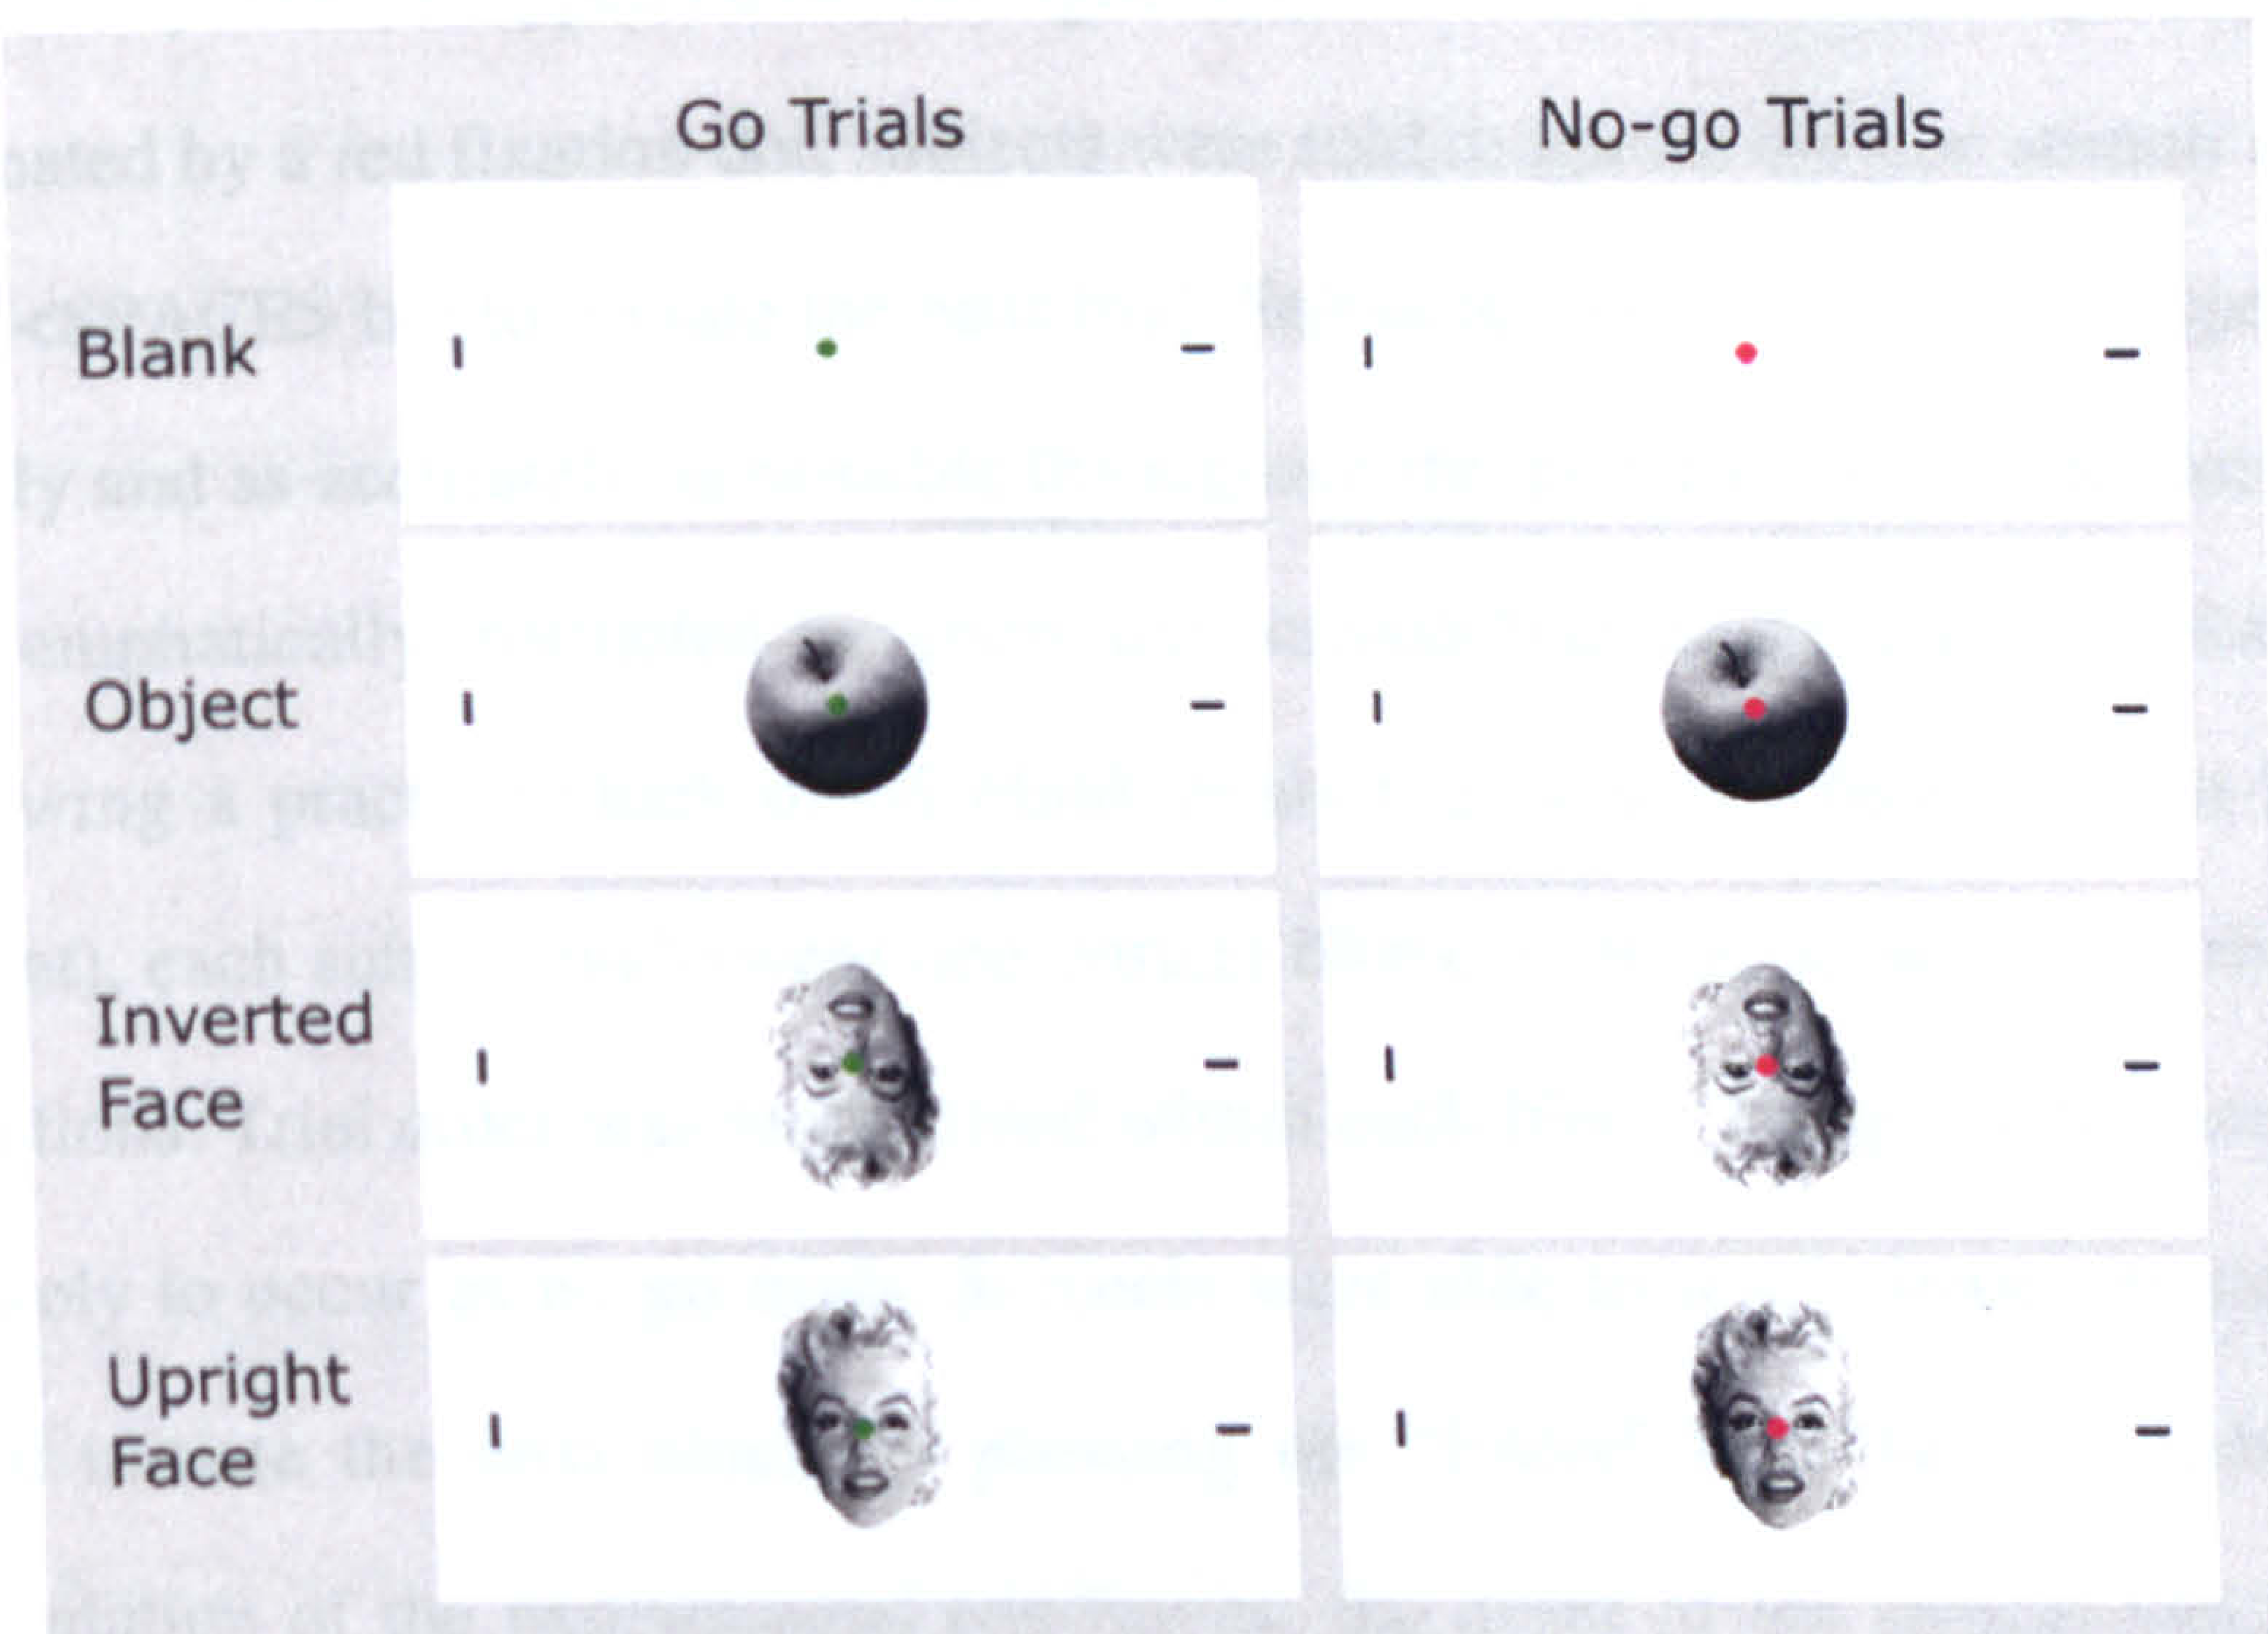

Figure 5.1 Example displays from the conditions of Experiment 1. On go trials, designated by a green fixation dot, the task was to classify whether the vertical line target appeared on the left or right side of the display. On no-go trials (red fixation dot), subjects were to ignore the target and press <space> to initiate the next trial. To assess attention biases from faces, the go/no-go displays could be *Blank*, or could contain a meaningful nonface item (the *Object* condition), an upside-down face (the *Inverted Face* condition), or a matching upright face (the *Upright Face* condition).

Procedure Subjects were seated at a viewing distance of 60 cm from the screen, which was kept constant by means of a chinrest. Each trial began with a black fixation dot for 750 ms, followed by a stimulus display, which appeared for 200 ms (i.e. too briefly to permit a stimulus-responsive eye-saccade), and a blank screen until a response registered. Subjects were instructed to focus on the centre of the screen at the start of each trial, and if the experimental display contained a green fixation dot (on go trials) to make speeded judgements regarding the location of the vertical line. Responses were made by pressing the "D" key on a standard computer keyboard to indicate when the target appeared on the left and the "K" key when it appeared on the right side of the display. For no-go displays, designated by a red fixation dot, subjects were told to ignore the line stimuli and to press <SPACE> bar to initiate the next trial. Subjects were requested to respond as quickly and as accurately as possible throughout the experiment. In addition, they were emphatically instructed to ignore any stimuli that might appear at fixation. Following a practice block of 36 blank trials (i.e. with no face/nonface image present), each subject underwent one critical block of 36 trials of each of the four conditions. Trial order was randomized within each block, but go trials were twice as likely to occur as no-go trials. Subjects were able to rest between blocks, and could initiate the next block by pressing the "Enter" key. To counterbalance presentation of the experimental conditions, the order of the critical blocks was rotated across subjects over the course of the experiment.

Results

Incorrect responses were discarded and the median correct response times and error rates were calculated for all subjects. Performance for go/no-go trials was important to confirm that subjects were attending to the centre of the displays.

Overall accuracy was high (97.0% for go trials versus 92.9% for no-go trials), indicating that subjects had complied with these task demands. Data from no-go trials was not analyzed further. For go trials, the median correct RTs and error rates were computed as a function of the experimental conditions (*Blank*, *Object*, *Inverted Face*, and *Upright Face*). The intersubject means of these RTs are shown in Figure 5.2.

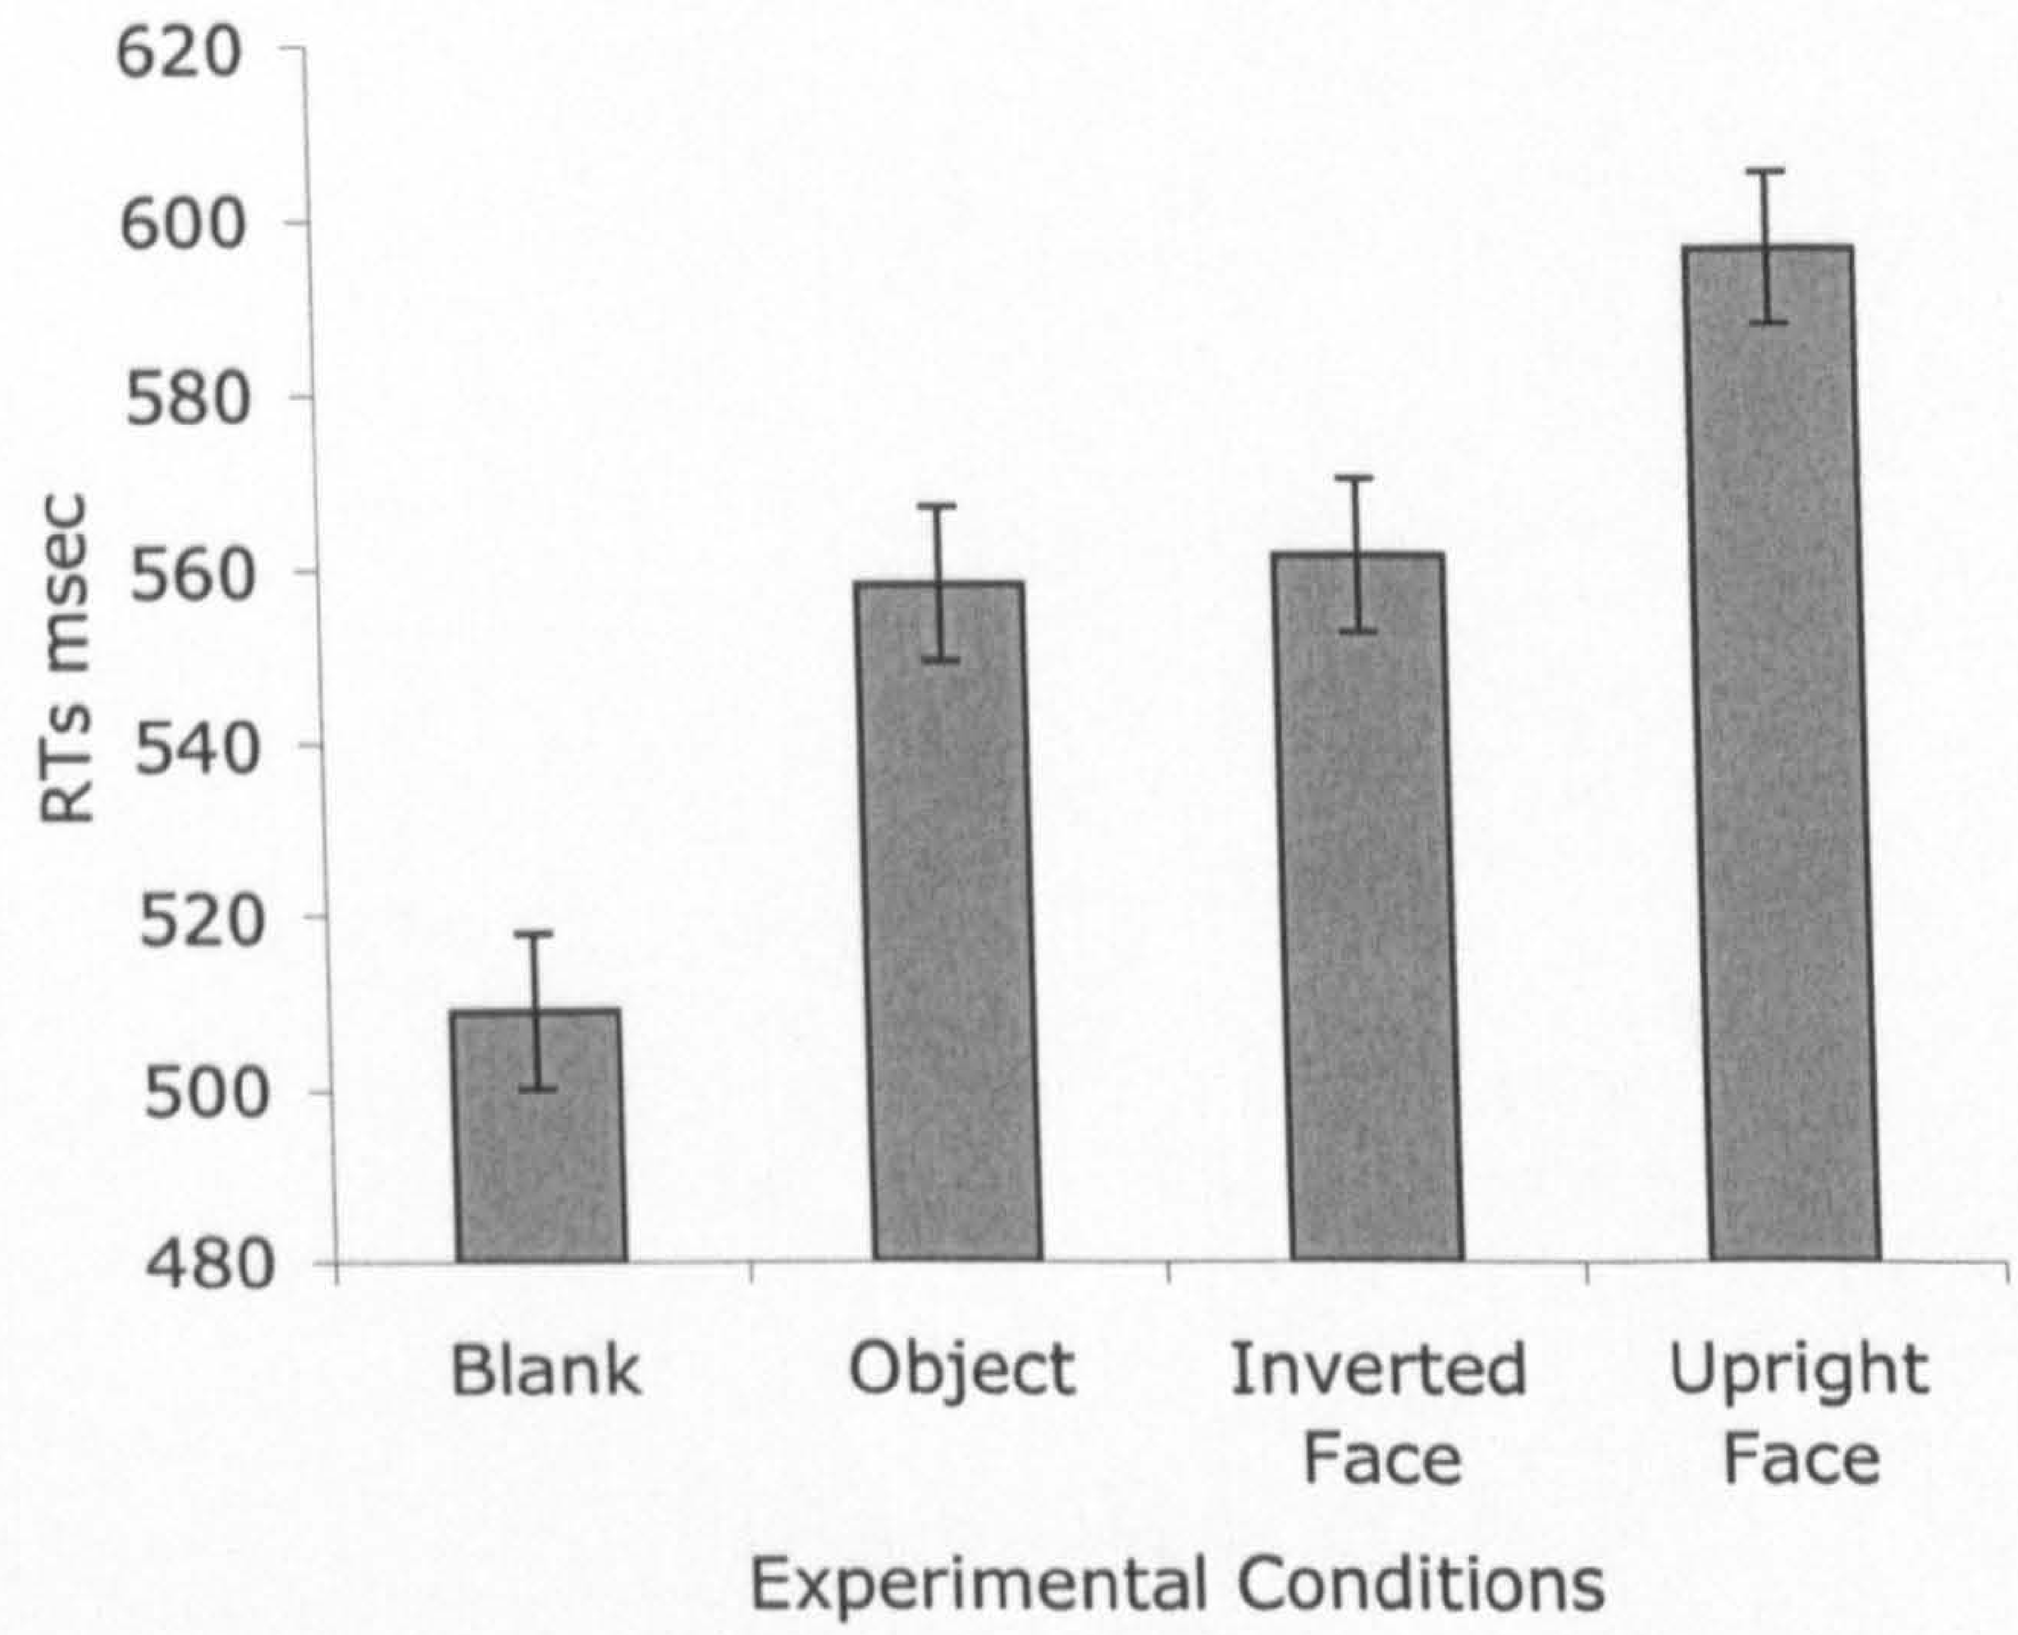

Figure 5.2 Mean correct responses (n=20) to line targets on go trials in Experiment 12. Performance is shown as a function of experimental condition; *Blank*, *Object*, *Inverted Face*, and *Upright Face*. Standard error bars are shown.

A one-factor within-subjects ANOVA (*Blank* versus *Object* versus *Inverted Face* versus *Upright Face*) revealed a main effect of experimental condition, $F(3,57)=16.58$, $p<.01$. Tukey's HSD test showed that RTs were significantly faster in the *Blank* condition, in which no additional stimulus was present, in comparison with each of the other conditions ($p<.01$). More importantly, RTs were also significantly faster in the *Object* condition and in the *Inverted Face* condition than in the *Upright Face* condition ($p<.05$). In addition, no difference between the *Object* condition and the *Inverted Face* condition was found. Thus, upright faces delayed target classification most, followed by inverted faces and objects.

Errors were made on 1.3% of *Blank* trials, 3.5% of *Object* trials, 4.2% of *Inverted Face* trials, and 3.1% of *Upright Face* trials. In contrast to the RT data, a one-factor within-subjects ANOVA of the percentage errors showed that these were evenly matched across the experimental conditions, $F(3,57)=2.16$. Errors were not analyzed further.

Discussion

Responses to the peripheral targets were fastest in the *Blank* condition in comparison with each of the other conditions. More important, responses were evenly matched on trials on which inverted faces and objects (in this case images of fruits) were shown, but were slowest when upright faces were displayed. This pattern does not appear to reflect a speed-accuracy trade off, as error rates were evenly matched across all conditions. Indeed, the error pattern indicates that presenting upright faces, inverted faces, and fruits did not affect whether the line targets were processed (see e.g. Lavie, 1995, 2000), but only delayed responses to these targets. There are two potential explanations for these results. First, it is conceivable that the increase in RTs might reflect the reduced salience of the coloured go/no-go signal, when it is superimposed on a more complex visual image than a plain background. Although this is possible, it cannot explain why responses were slowest for upright faces, which matched their inverted counterparts in almost every aspect except orientation. A more plausible explanation is that subjects took longer to disengage attention from these stimuli, which in turn delayed target classification, but particularly so from images of upright faces. This latter interpretation would not only be consistent with previous findings suggesting an attentional bias to faces (Mack et al, 2002; Ro et al, 2001; Vuilleumier, 2000), but also with claims that inverted faces are subject to object rather than typical face processes (e.g. Farah et al, 1995; Haxby et al, 1999; Moscovitch et al, 1997). However, at present these findings provide only initial evidence of a *general* face bias, as only famous faces were used. In addition to low-level physiognomic face information, known faces also provide higher-level information such as identity and semantic knowledge about a person to an observer. If this contributes to attentional retainment then the present face bias might only reflect face familiarity. This possibility was examined in the next experiment.

Experiment 13

The main finding of Experiment 12 is that response times to a peripheral target increase when faces are presented at fixation. This increase cannot simply be due to the low-level properties of face stimuli, as equivalent increases were not found when these faces were inverted. Similarly, target classification was less affected by meaningful objects than by upright faces. In addition to existing claims that threatening face stimuli hold attention in highly anxious people (e.g. Fox et al, 2001, 2002), these findings provide some initial evidence that it might also be *generally* difficult to disengage from faces. To provide a further test for this claim, the next experiment examined whether these effects are still observed when unfamiliar faces are used, which, unlike the famous face stimuli of Experiment 12, rule out any possible influence of higher-level identity and semantic information.

Method

Subjects, Stimuli & Procedure Eight postgraduate students from the University of Utrecht, The Netherlands, and twenty-one students from the University of Glasgow, whose ages ranged from 19 to 31 years, participated in the experiment on a voluntary basis or for course credits. All had normal or corrected to normal vision. The famous face stimuli from Experiment 11 were replaced with equivalently prepared face stimuli of three unfamiliar female models (see Figure 5.3). In all other respects, the design and procedure were identical to the previous experiment.

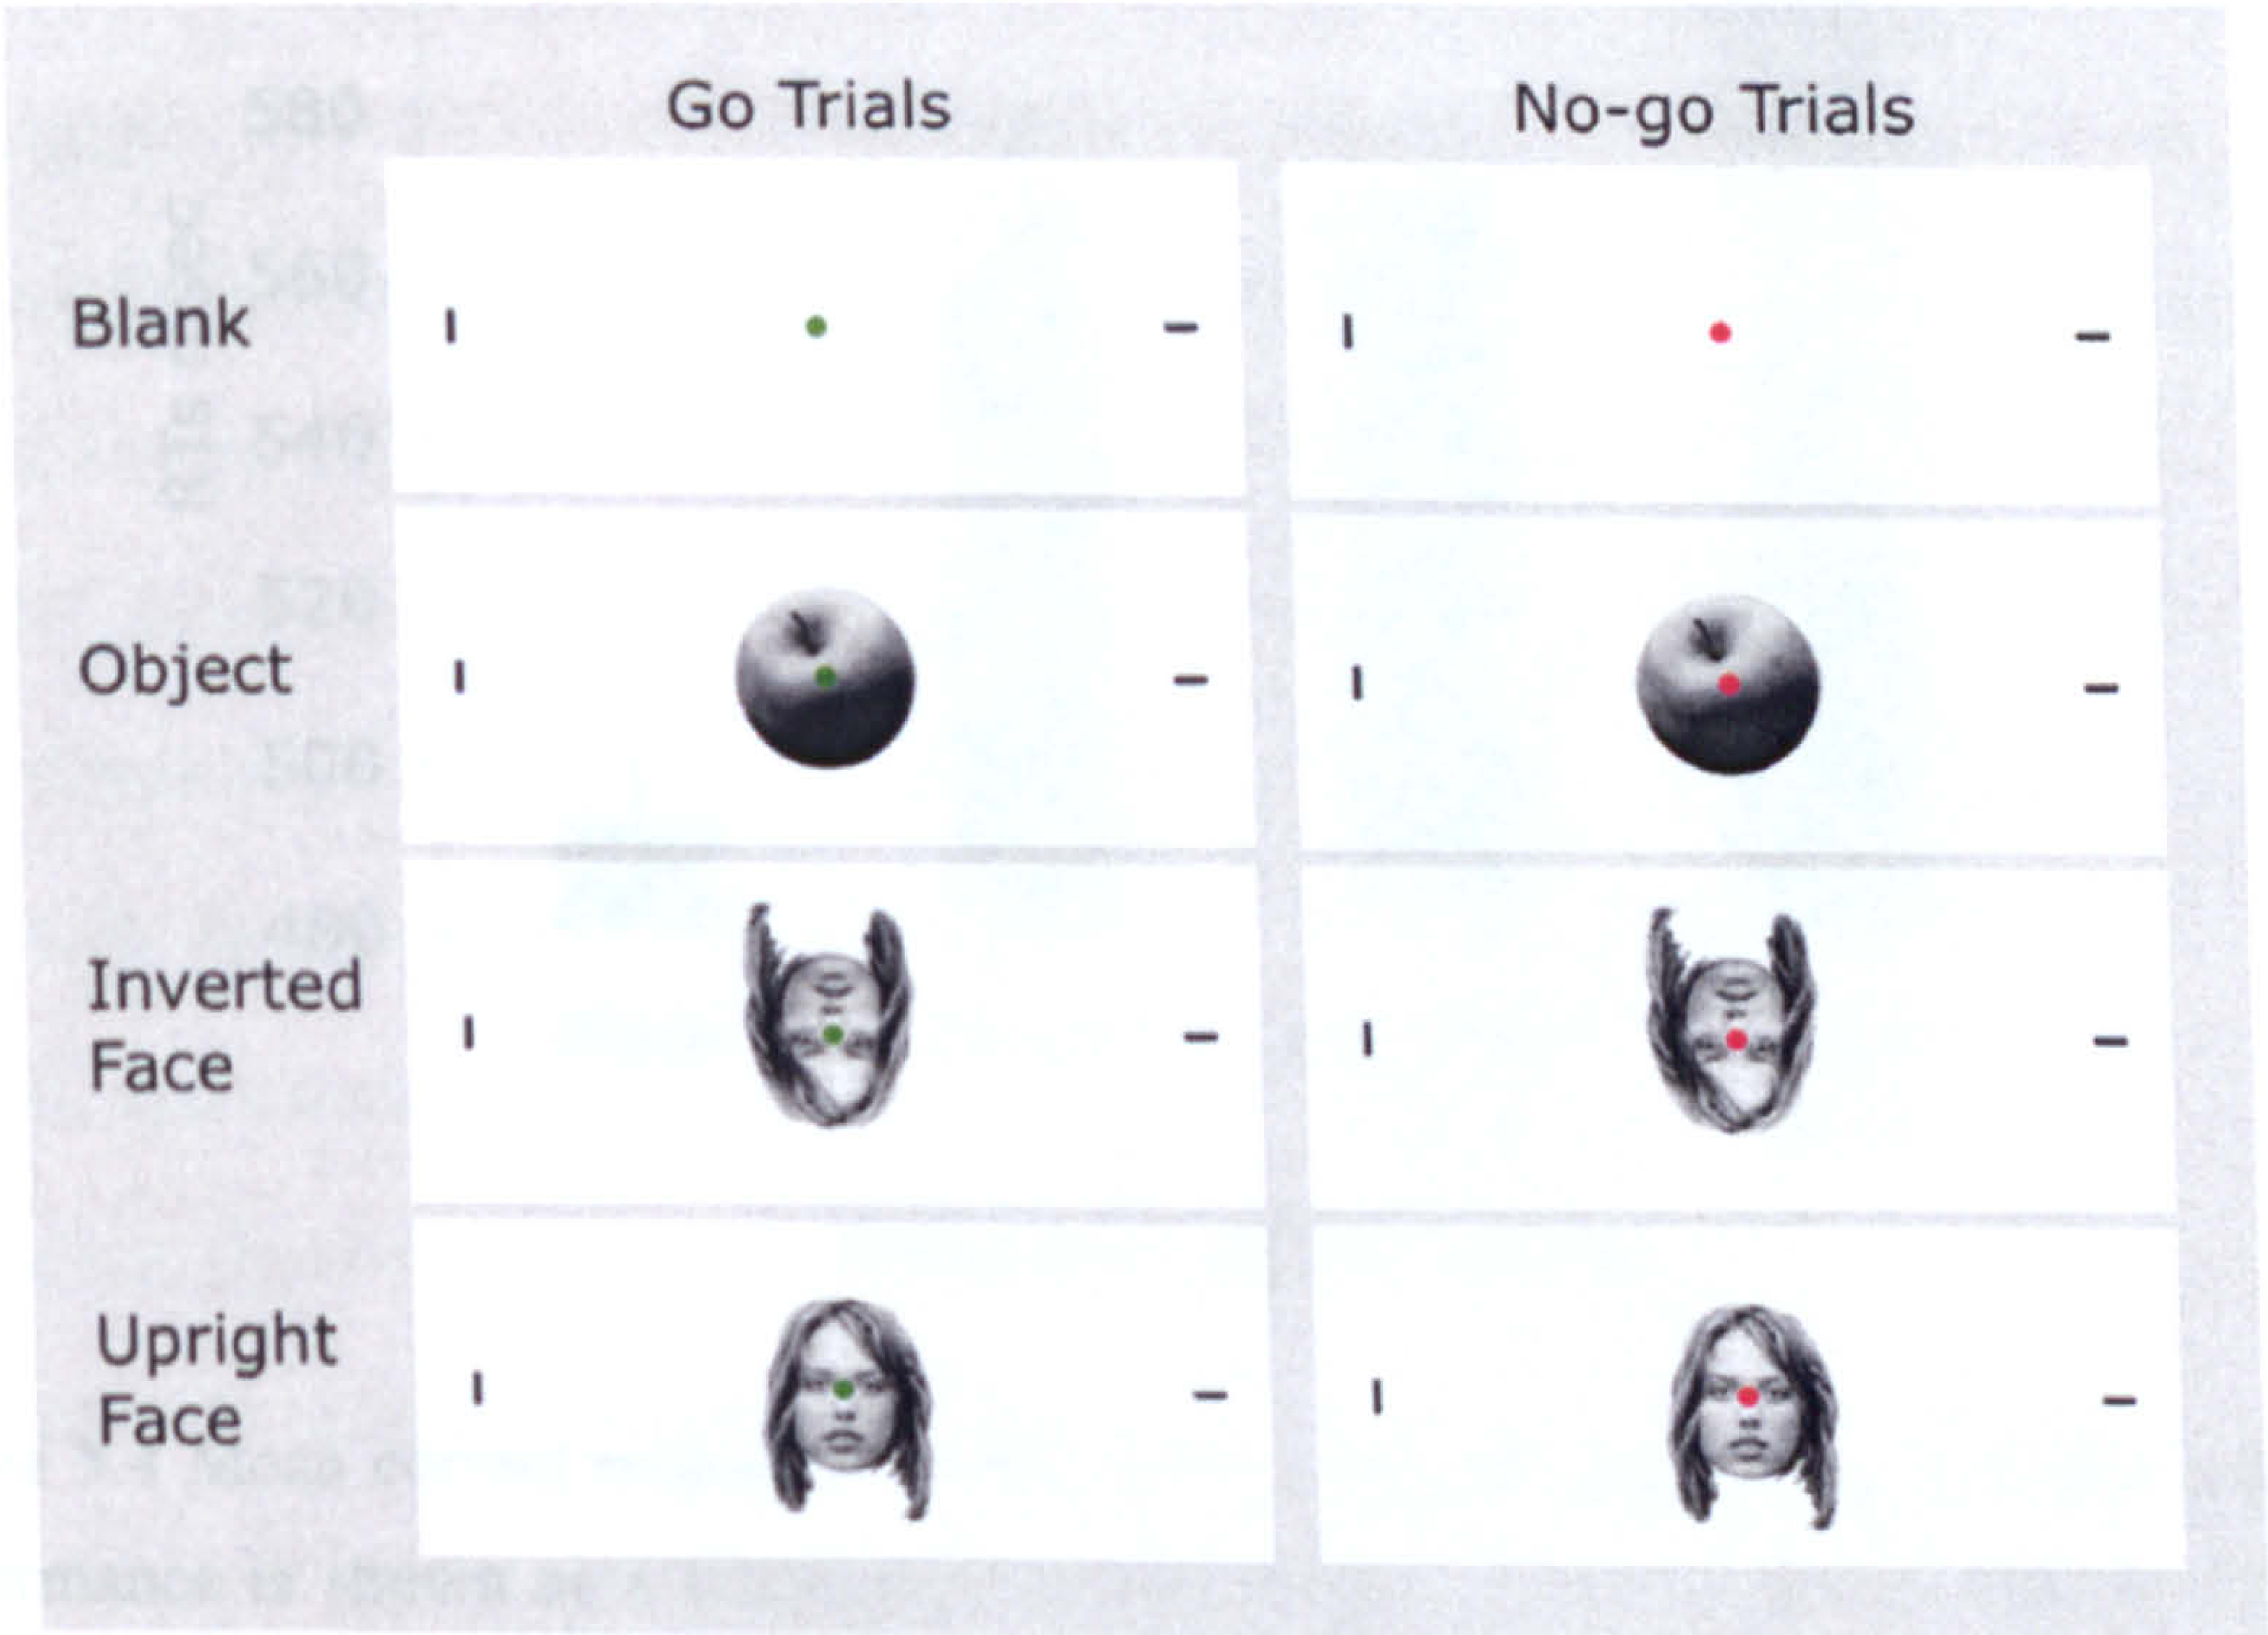

Figure 5.3 Example displays from the four experimental conditions (*Blank*, *Object*, *Inverted Face*, and *Upright Face*) and for go and no-go trials in Experiment 13.

Results

Incorrect responses and data from one subject whose response latencies were exceptionally slow (3.3 SD from the group mean) were discarded from analysis. Go/no-go accuracy was high (93.8% for go trials vs. 90.0% for no-go trials), which confirms that subjects were attending to the centre of the displays. For go trials, the median correct RTs and error rates were computed for each of the experimental conditions (*Blank*, *Object*, *Inverted Face*, and *Upright Face*). The averages of the median RTs across subjects are shown in Figure 5.4.

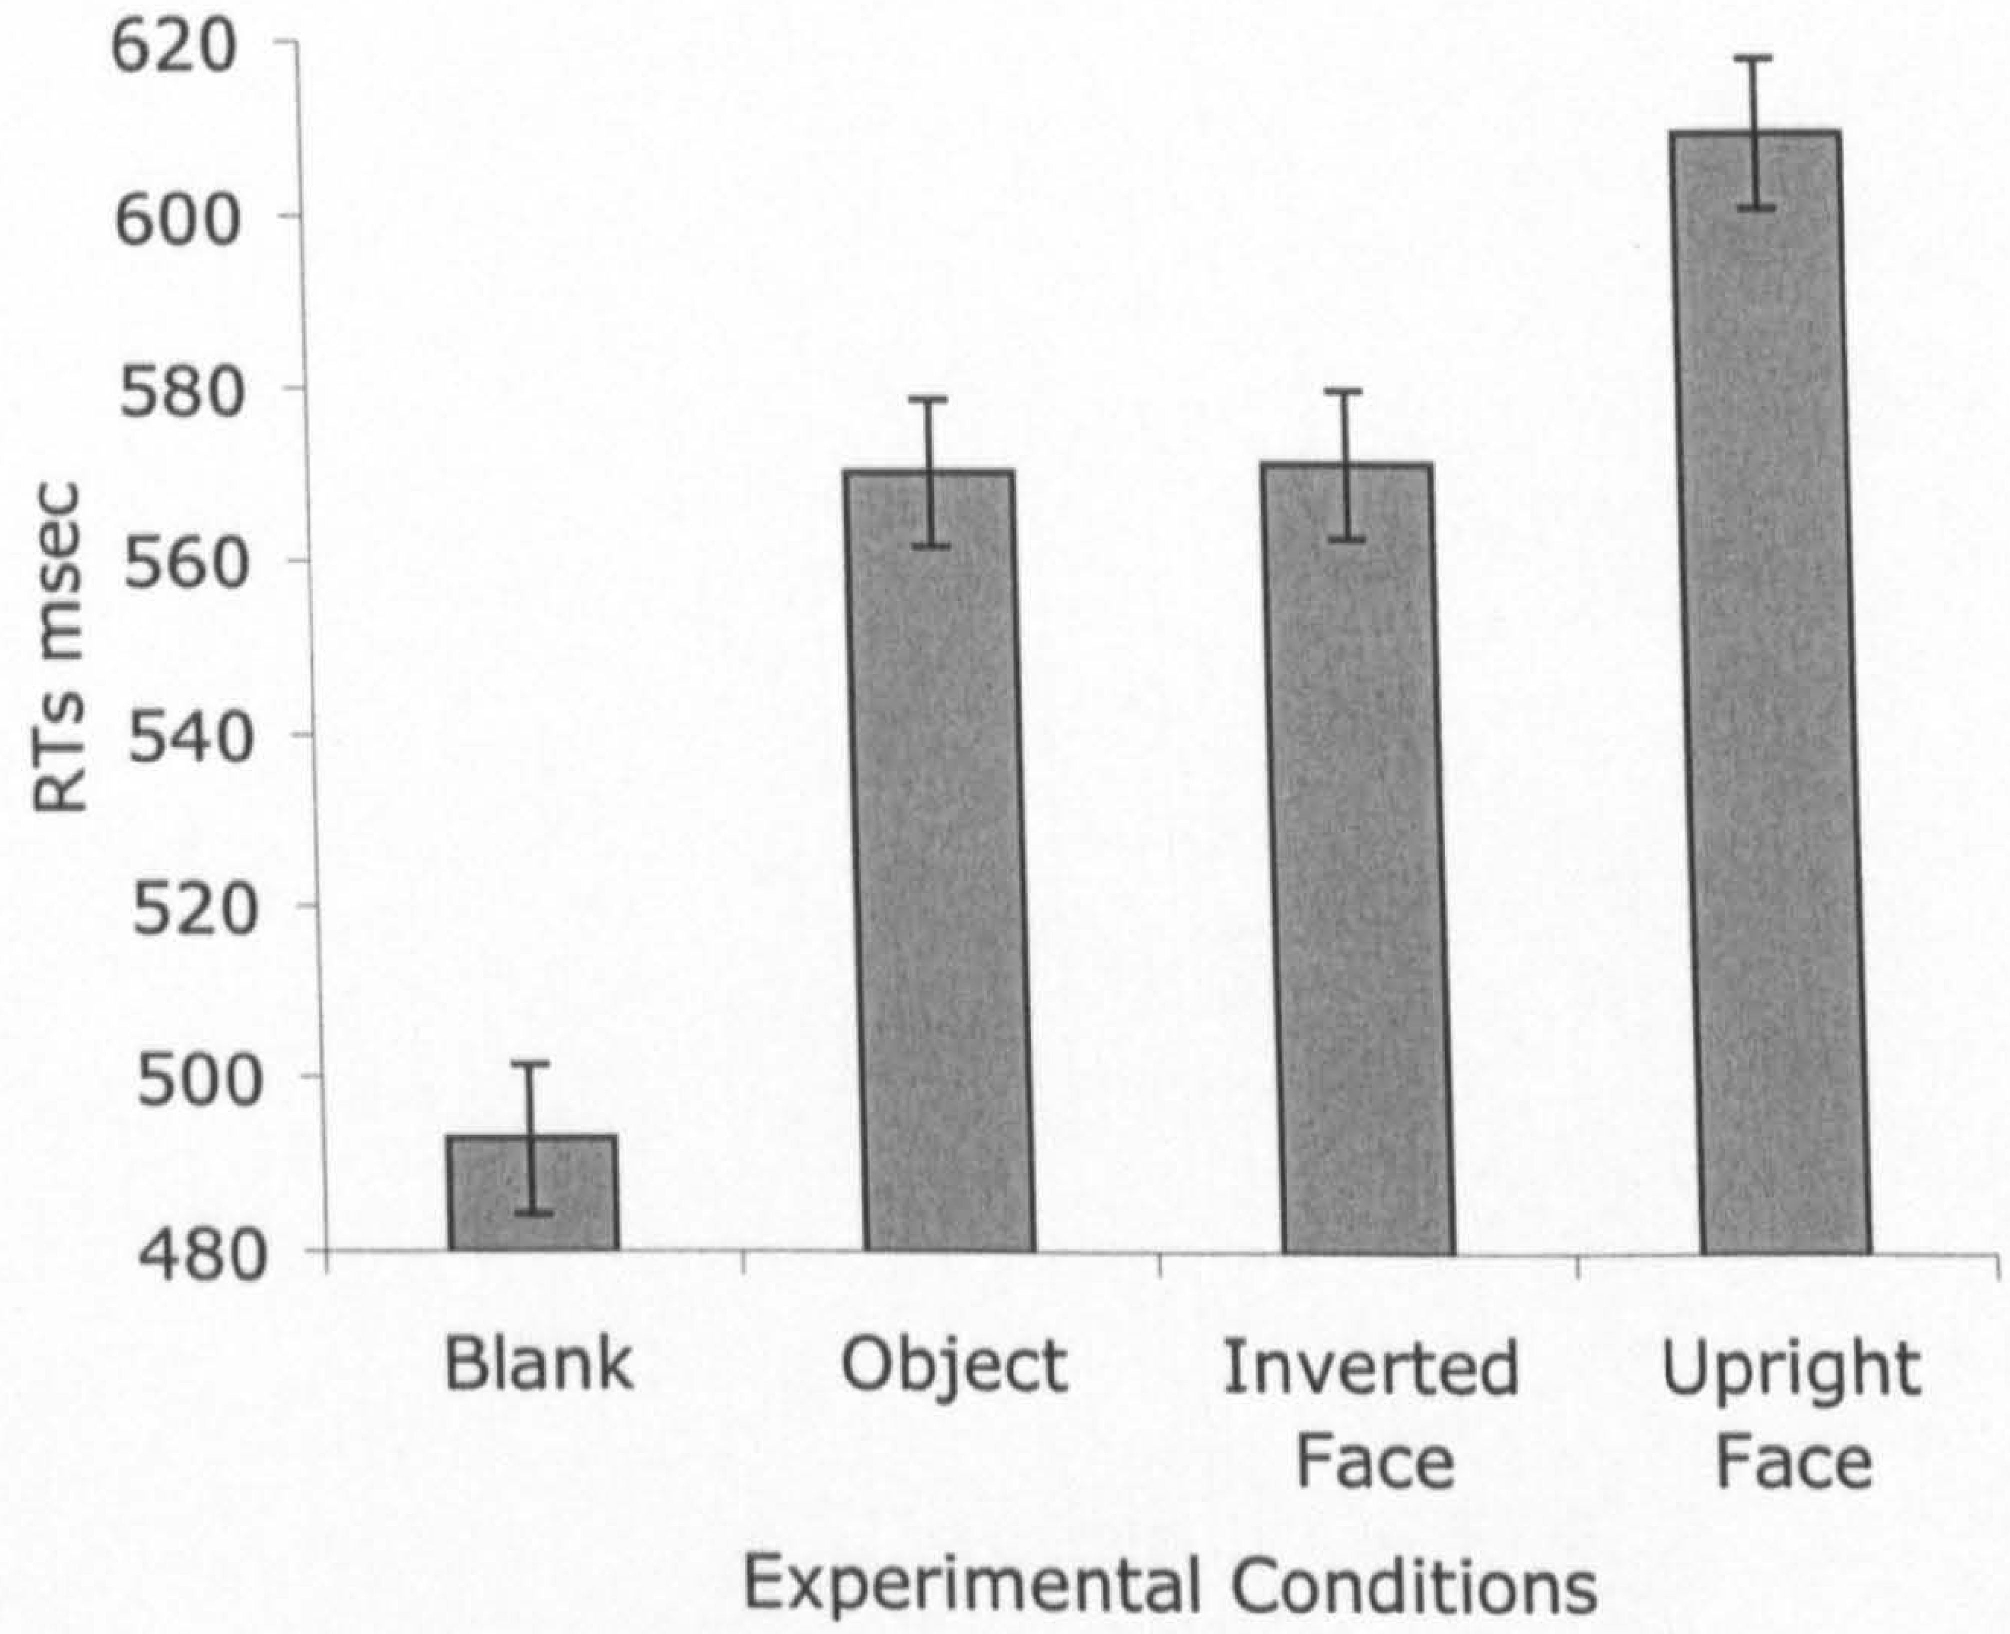

Figure 5.4 Mean correct responses (n=28) to line targets on go trials in Experiment 13. Performance is shown as a function of experimental condition; *Blank*, *Object*, *Inverted Face*, and *Upright Face*. Standard error bars are shown.

A one-factor within-subjects ANOVA on the RT data revealed an effect of condition, $F(3,81)=33.16$, $p<.01$. As before, Tukey's HSD test showed significantly faster RTs in the *Blank* condition than in any of the other conditions

($p<.01$). More important, target RTs were again significantly faster in both the *Object* condition and in the *Inverted Face* condition than in the *Upright Face* condition ($p<.05$), but there was no difference between the *Object* and the *Inverted Face* condition.

Errors were made on 3.7% of *Blank* trials, 9.4% of *Object* trial, 5.7% of *Inverted Face* trials, and 6.3% of *Upright Face* trials. A one-factor ANOVA of the error data revealed a main effect of condition, $F(3,81)=4.06$, $p<.05$, reflecting higher percentage errors in the *Object* condition than in the *Blank* condition. No other differences were significant.

Discussion

The results are strikingly similar to Experiment 12. As before, response times were fastest in the condition in which a blank background was used, slower when inverted faces or objects were added at fixation, but slowest when upright faces were displayed. This pattern is consistent with the notion of a general attentional bias such that it may be particularly difficult to disengage from upright face stimuli, and extends the results of Experiment 12 from famous to unfamiliar faces. Therefore, these findings rule out the possibility that a face bias reflects processes that can only be engaged by familiar faces.

Experiment 14

To strengthen the claim that it may be particularly difficult to disengage from faces, Experiment 14 examined whether a face bias is still observed when the inverted famous faces and fruits of Experiment 12 are replaced with famous names and images of flags. Peoples' names provide virtually no resemblance to their faces but just like them are recognised with little difficulty. Moreover, names give the same status as persons as faces. If famous faces are more proficient at maintaining attention than the same person's names, this would therefore provide further evidence that faces maintain attention more effectively than other classes of stimuli. In addition to substituting famous names for inverted famous faces, the fruit distractors were also replaced with images of national flags. One particular reason for choosing flags in addition to names was to provide an analogue to the face-nonface conditions of Experiments 4-6, where faces were subject to less distractor interference than flag and name targets. If faces' ability to hold attention surpasses that of names and flags, then this could provide one conceivable explanation for these effects.

Method

Subjects Eight postgraduates from the University of Utrecht and a further sixteen undergraduate students from the University of Glasgow, aged 18-28, participated in the experiment on a voluntary basis or for course credits. All had normal or corrected to normal vision.

Stimuli & Procedure The stimuli and procedure were identical to Experiment 1, except as follows. The inverted famous faces were replaced by the names of the same celebrities (Pamela Anderson, Marilyn Monroe, and Britney Spears). These were printed in black 18-point Arial font, with forenames printed above and surnames below the go/no-go signal, and measured between 1.3 cm and 1.8 cm (1.3°–1.7° of VA) in width. In addition, the images of fruits were replaced with a set of national flags (the Greek flag, the South African flag, and the Swiss flag). The flag stimuli were derived from photographs, which were rendered to greyscale, cropped to remove extraneous background, and sized to 2.4 cm x 3.0 cm (2.3° x 2.9° of VA). As before, the names and flags were then copied onto the go/no-go displays, behind the fixation dot, for the *Name* and *Object* conditions (see Figure 5.5).

Figure 5.5 Example displays from the four experimental conditions (*Blank*, *Object*, *Name*, and *Face*) and for go and no-go displays in Experiment 14.

<u>Results</u> As for the previous experiments, incorrect responses were discarded from analysis. Accuracy for go/no-go trials was high once again, averaging at 93.7% for go trials and 86.7% for no-go trials. The means of the median correct RTs for the go conditions are shown in Figure 5.6.

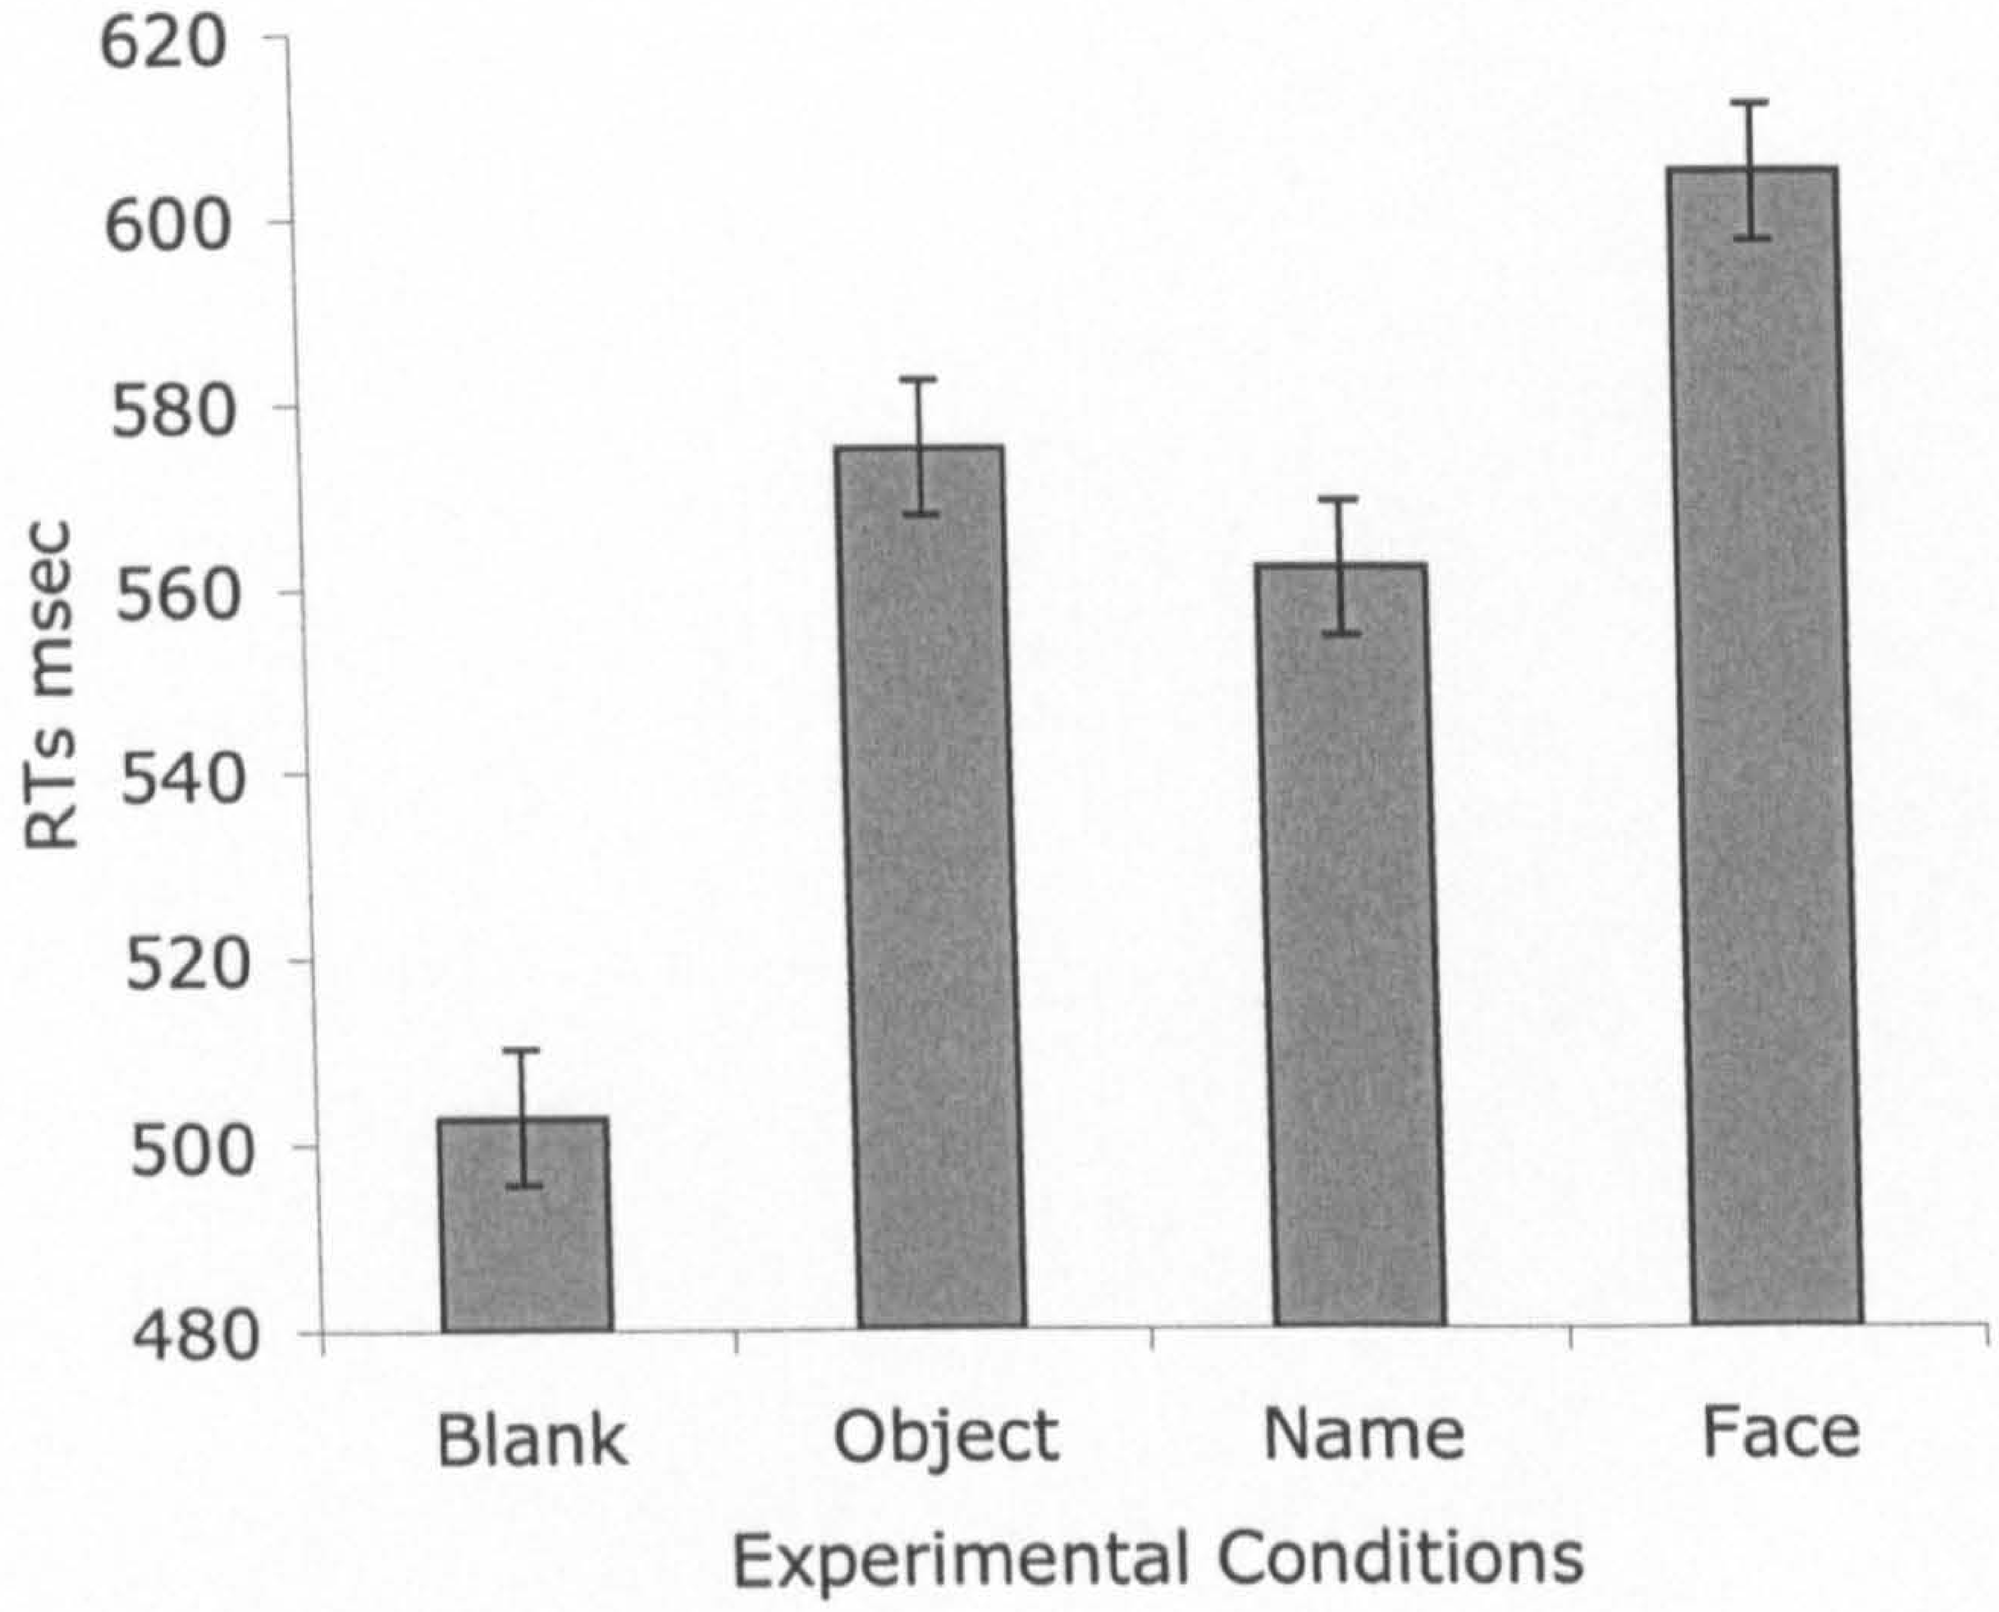

Figure 5.6 Mean correct responses (n=24) to line targets on go trials in Experiment 14. Performance is shown as a function of experimental condition; *Blank*, *Object*, *Name*, and *Face*. Standard error bars are shown.

A one-factor within-subjects ANOVA (*Blank* versus *Object* versus *Name* versus *Face*) of the mean RT data revealed a main effect of condition, $F(3,69)=33.48$, $p<.01$. Comparisons on this effect with Tukey's HSD test showed significantly faster classification times for the *Blank* condition in comparison with each of the other three conditions ($p<.01$). In addition, RTs were significantly faster for the *Object* condition and the *Name* condition than for the *Face* condition ($p<.05$). However, the *Object* and the *Name* condition did not differ from each other.

Errors were made on 3.1% of *Blank* trials, 8.0% of *Object* trial, 6.9% of *Name* trials, and 7.3% of *Face* trials. An analogous one-factor ANOVA of the error data revealed an effect of condition, $F(3,69)=4.76$, $p<.01$, reflecting fewer errors in the

Blank condition than in the *Object*, the *Name*, and the *Face* conditions ($p<.05$). No other comparisons were significant.

Discussion

Experiment 14 reveals several important findings. First, a person's face again appeared more potent in delaying target classification than other classes of stimuli, in this case famous names and images of flags. This provides further evidence that faces are particularly persistent in holding attention. In addition to the findings from Experiments 12 & 13, these results also suggest that the interference effects of Experiments 4-6, specifically the observation that faces were subject to less interference than name and flag targets, might at least partially reflect an attentional bias for faces. There are, of course, a number of differences between those experiments and the present studies. Perhaps most important, the target stimuli in Experiments 4-6 were always task-relevant, whereas the faces and object comparisons here were presented in a task-relevant location but were not explicitly implicated in target classification. To determine the relationship between attention and face processing, however, this has the advantage of diminishing the potential influence of any ongoing task-demands.

The present results may also serve to resolve one recurring aspect of Experiments 12-14, namely that target RTs were still significantly slowed by inverted faces and nonface stimuli, albeit less so than by upright faces, in comparison with trials in which a blank background was presented. In Experiment 12 it was already suggested that this might reflect the reduced salience of the coloured go/no-go signal, when it is superimposed on complex visual stimuli. However, the present effects with name stimuli do not favour this explanation, as the names were presented above and below the go/no-go signal, thus leaving its salience relatively intact (see Figure 5.5). Rather, these findings suggest that names, inverted faces, and images of fruits and flags also retain attention in this design, but less so than faces.

Experiment 15

The previous experiments show that participants are slower to classify a peripheral target, when a face is displayed at fixation in the location of a task-relevant go/no-go signal, but are less affected by inverted faces, names, or meaningful objects. These findings appear consistent with the hypothesis that it may be particularly difficult to disengage attention from faces. The purpose of the next experiment was to examine whether a similar face bias still exists when the targets are presented in close proximity to the go/no-go signal and the faces and nonface distractors appear clearly separated, to the side of the target. Thus, the face and nonface stimuli always appeared in a task-irrelevant location in this design. If faces disrupt target classification more than printed names and images of flags here, then this would support the idea that they are also particularly adept at engaging attentive resources (e.g. Mack et al, 2002; Ro et al, 2001; Vuilleumier, 2000).

Method

Subjects The thirty-two subjects were postgraduates and undergraduates from the University of Glasgow whose ages ranged from 18-25 years. All reported normal or corrected to normal vision and volunteered to participate for free or for course credits.

Stimuli & Procedure The stimuli and procedure were the same as for Experiment 14, except that the line targets were now positioned just 0.7 cm (0.7° of VA) from the centre of the display, and the face and nonface stimuli were presented in the periphery to the left or right of fixation, with a horizontal distance of 1.0 cm (1.0° of VA) between the nearest line (horizontal or vertical) and the face/nonface stimuli (see Figure 5.7). As before, subjects completed a practice block of *Blank* trials, followed by an experimental blocks of each of the four conditions (*Blank*, *Name*, *Object*, and *Face*). Each of these blocks consisted of 36 randomly ordered trials, and block order was counterbalanced across the experiment.

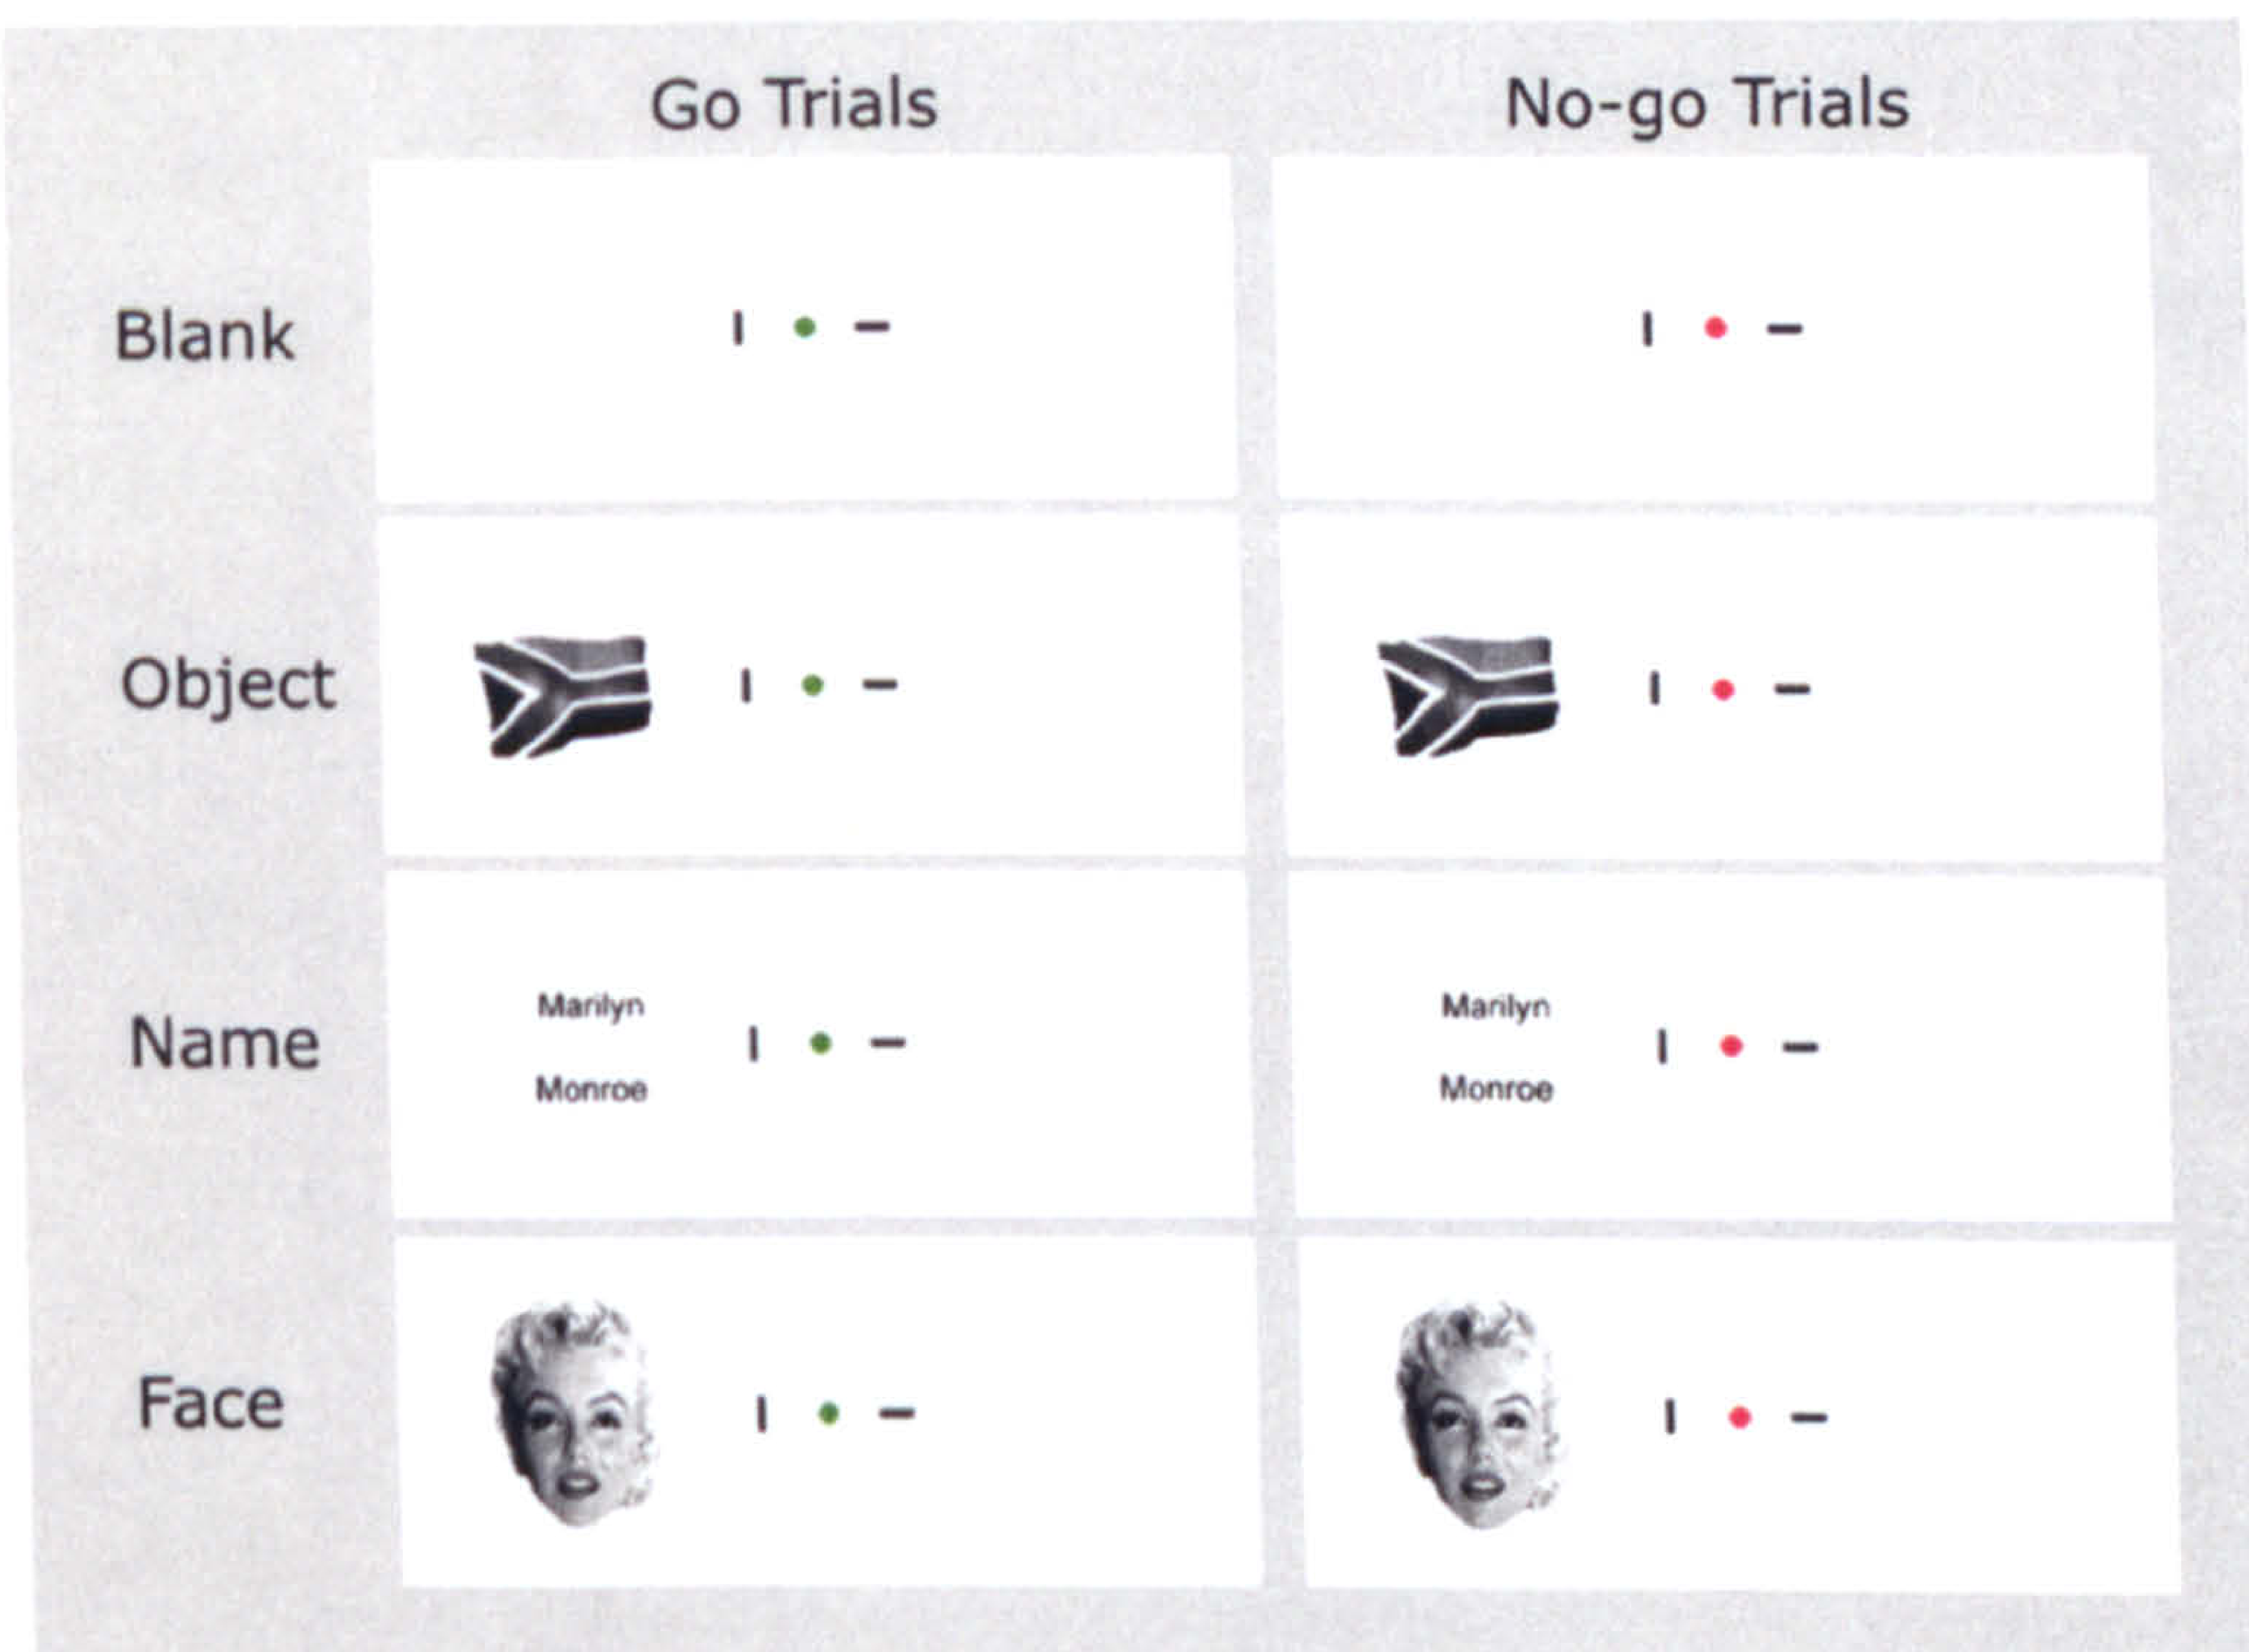

Figure 5.7 Example display from Experiment 15. The vertical line (the target) and the horizontal line were now presented in the centre of the display, and the faces and nonface comparisons were moved to a task-irrelevant, peripheral location.

Results

Incorrect responses were excluded from analysis and overall accuracy for go/no-go trials was computed by pooling the number of correct responses across the experimental conditions. As in all of the preceding experiments, accuracy was high for go trials (97.2% correct) and no-go trials (88.3% correct), indicating that subjects were attending to the go/no-go signals. For go trials, the median correct RTs and error rates were computed separately for the experimental conditions (*Blank*, *Name*, *Object*, and *Face*). The intersubject means of the medians RTs are displayed in Figure 5.8.

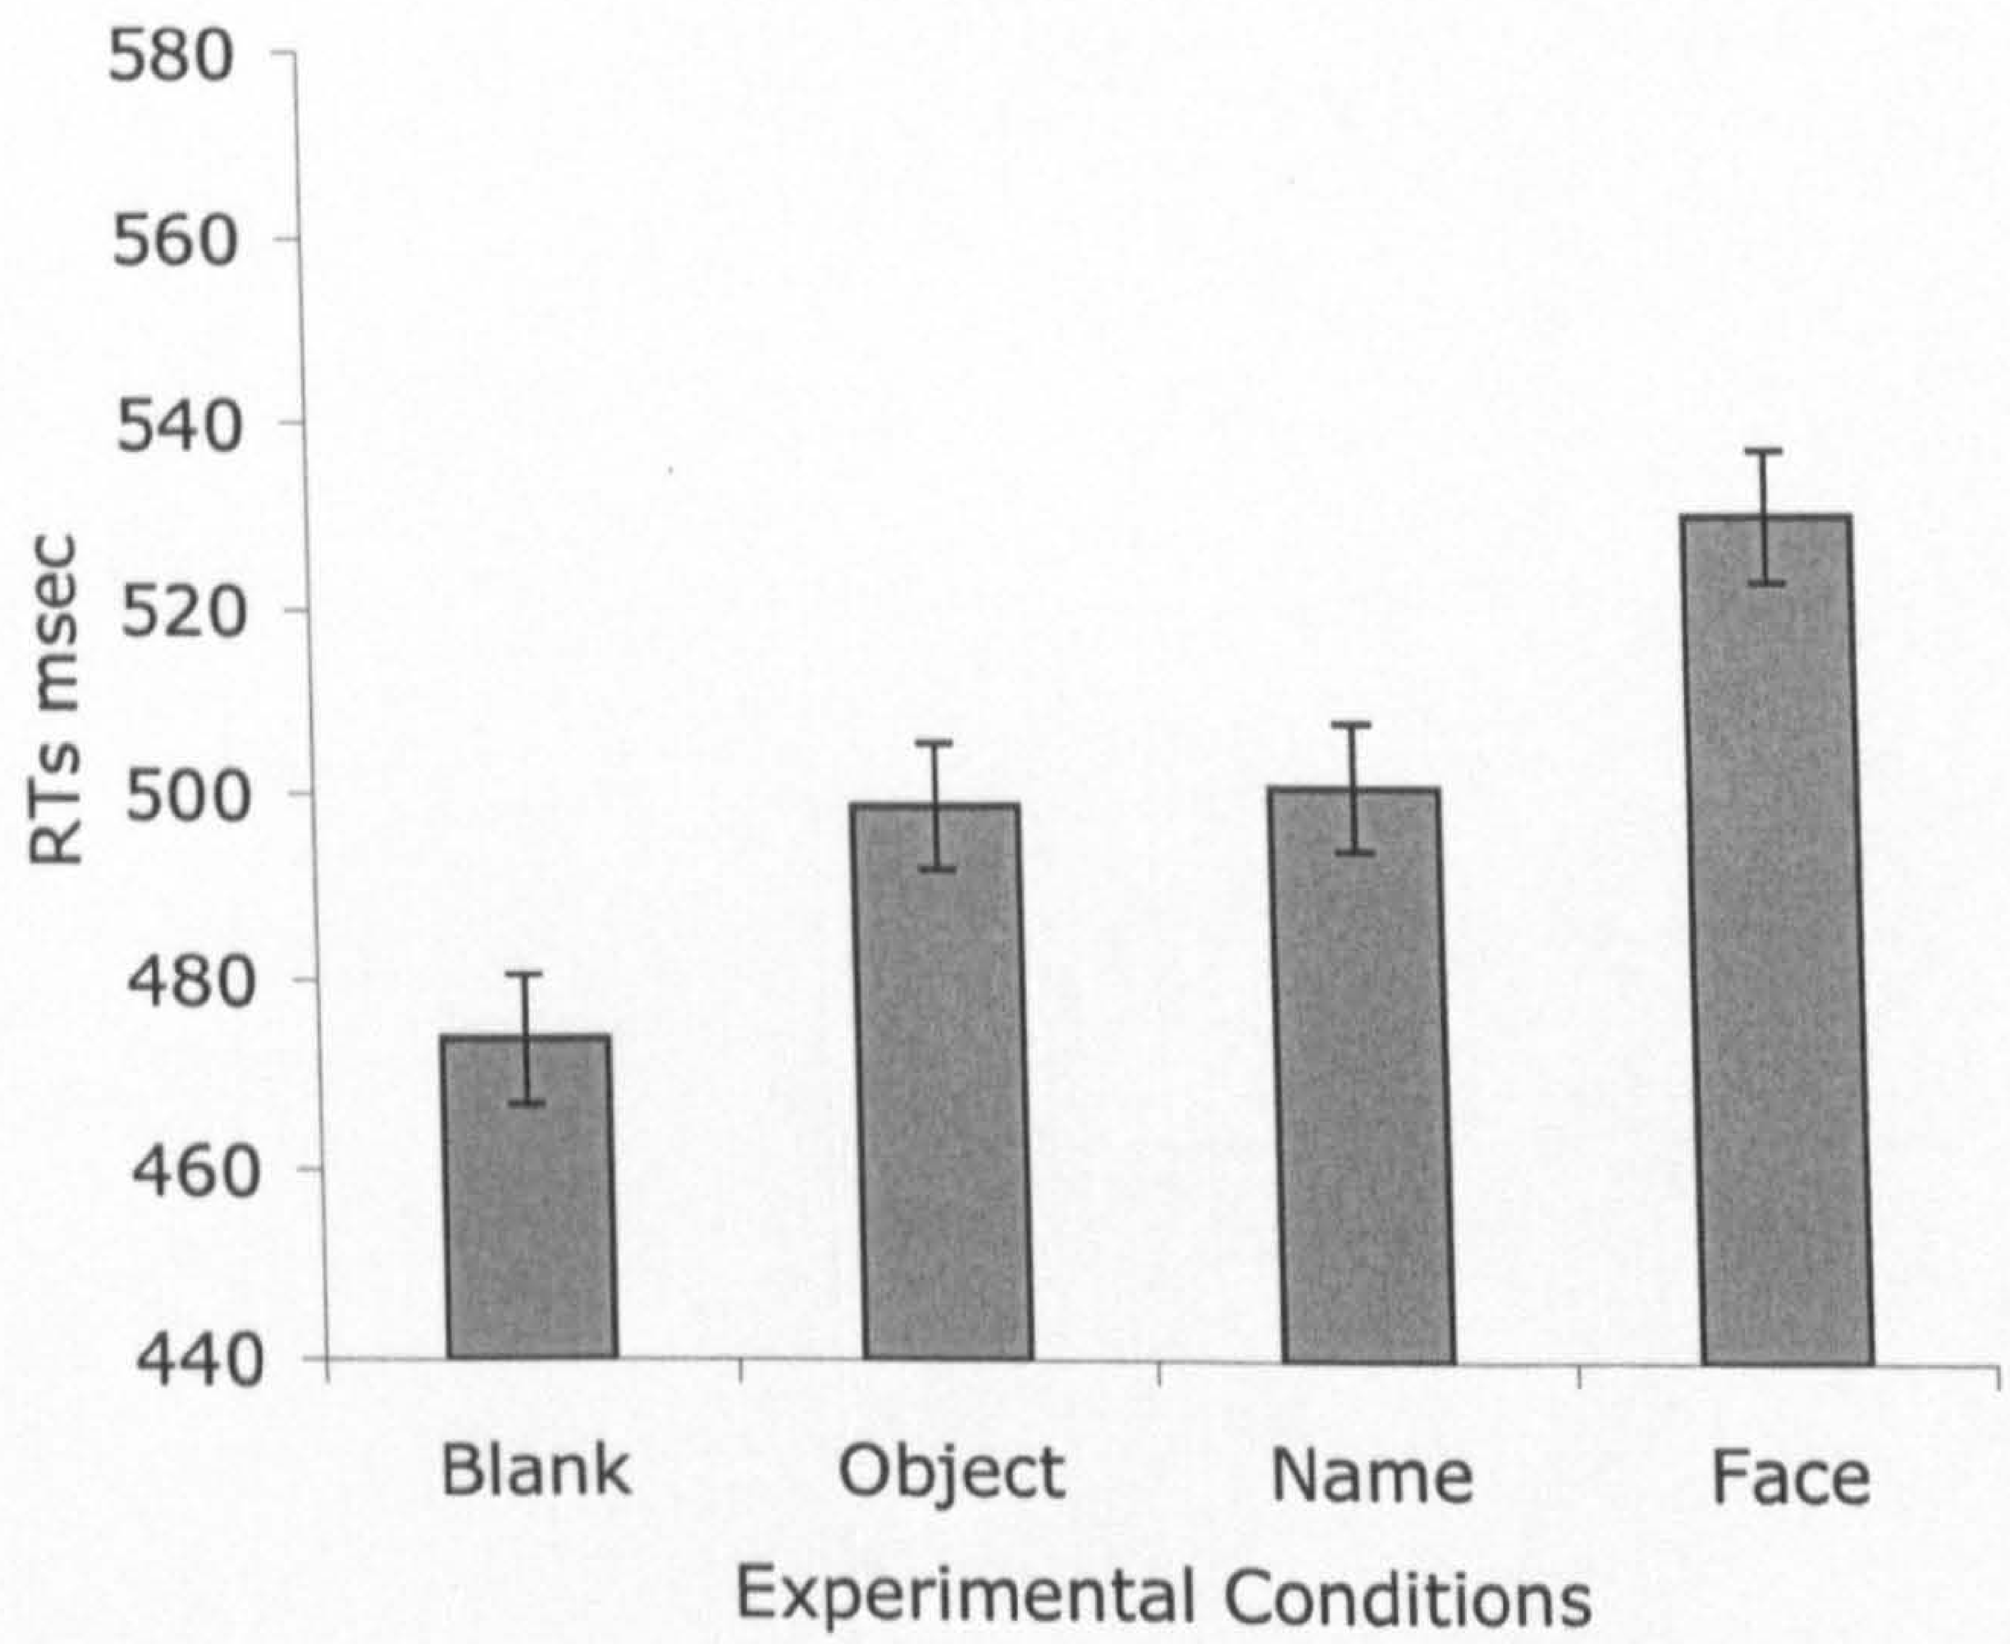

Figure 5.8 Mean correct responses (n=32) to line targets in the go displays of Experiment 15. Performance is shown as a function of experimental condition; *Blank*, *Object*, *Name*, and *Face*. Standard error bars are shown.

Comparisons across the experimental conditions A one-factor within-subjects ANOVA (*Blank* versus *Object* versus *Name* versus *Face*) on the mean RT data showed the usual main effect of condition, $F(3,93)=11.48$, $p<.01$. Comparisons between each of these conditions with Tukey's HSD test revealed significantly faster RTs in the *Blank* condition than in each of the other conditions ($p<.05$). More important, RTs were evenly matched in the *Name* condition and the *Object* condition, but significantly slower in the *Face* condition than in the *Blank*, *Name*, and *Object* conditions ($p<.01$). Errors were made on 2.9% of *Blank* trials, 3.3% of *Name* trials, 3.0% of *Object* trials, and 2.9% of Face trials. ANOVA showed no effect of condition, $F(3,93)<1$, and errors were not analyzed further.

Comparisons of spatial congruency effects within conditions A separate 3 (*Object* versus *Name* versus *Face*) x 2 (congruent versus incongruent) ANOVA was conducted to compare performance on trials on which face and nonface stimuli were presented on the same side as the vertical line target (spatially congruent trials) in comparison to when these images were presented on the opposite side of the display (on spatially incongruent trials), closer to the horizontal line and further from the vertical line target. The RTs for these conditions are shown in Figure 5.9 (see overleaf). If face and nonface images engage visual attention, as was indicated by significant comparisons with the *Blank* condition, then RTs should also vary as a function of spatial congruence, with slower responses on incongruent compared to congruent trials. This was confirmed by the statistical analysis, which showed a main effect of condition, $F(2,62)=9.027$, $p<.01$, again reflecting slower RTs for the *Face* condition than the *Name* and *Object* conditions, and a main effect of spatial congruency, $F(1,31)=25.74$, $p<.01$, reflecting slower responses on incongruent than congruent trials.

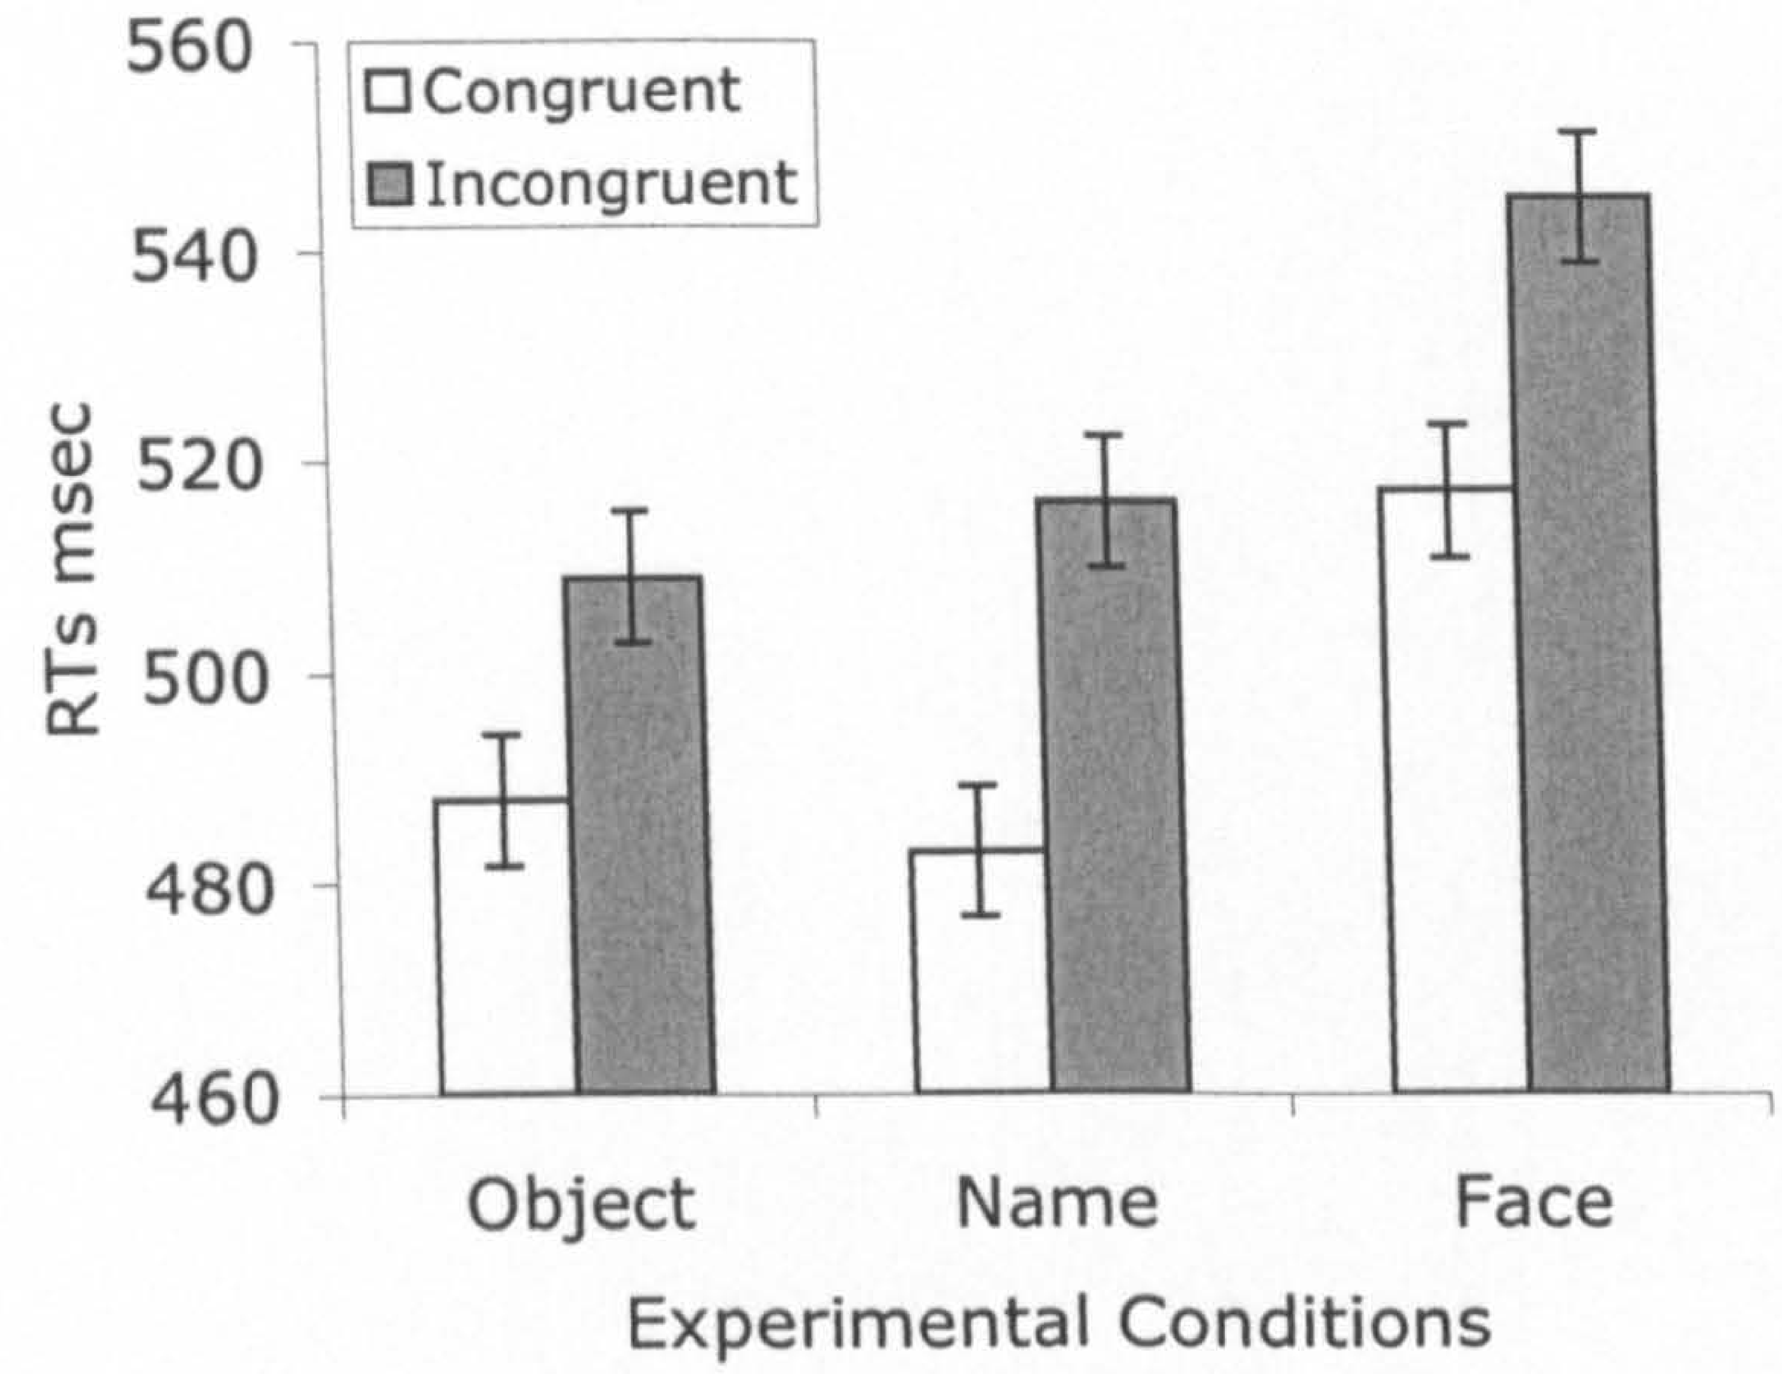

Figure 5.9 Means of the median reaction times for the *Object*, *Name* and *Face* conditions as a function of spatial congruency with the line target. Standard error bars are shown.

As Figure 5.9 suggests, analysis of simple main effects revealed significant congruency effects for each of the three conditions: *Name* condition, $F(1,32)=12.37$, $p<.01$, *Object* condition, $F(1,32)=5.01$, $p<.05$, and *Face* condition, $F(1,32)=9.19$, $p<.01$. However, no interaction between condition and spatial congruency was found, $F(2,62)<1$, indicating that these congruency effects were equivalent across the experimental conditions.

Discussion

The results of Experiment 15 show that classification of a centrally presented target, which occurred either just to the left or right of fixation, was affected by the peripheral presentation of simultaneously presented names and images of flags, but even more so by images of faces. These results converge with those of the preceding experiments, in which target classification was also more affected by faces than a range of other stimuli. However, in contrast to those experiments, where participants were required to attend to the location of these stimuli, the face and nonface objects here were never presented in a task-relevant location. Thus, whereas Experiments 12-14 provide a measure of disengagement from one spatial location, the present study suggests that faces are also particularly efficient at engaging attention at another.

Both of these processes may of course be influenced by the same mechanisms. Given the relative simplicity of the target classification task, it seemed inevitable that any spare processing capacity would spill over to the peripheral face and nonface objects (e.g. Lavie, 1995, 2000). In addition, because of their sudden onset these stimuli may have acted as exogenous cues, immune to any higher cognitive influences and capable of reflexively attracting attentive resources (e.g. Briand & Klein, 1987; Posner, 1980), and other stimuli might have done so too. The engagement bias for faces that was observed in this experiment thus most likely reflects an increase in attentional dwell time, not unlike that observed in Experiments 12-14, rather than an initial orienting bias towards faces. Whether faces are particularly likely to attract attentive resources to themselves in the first place or can do so very rapidly needs to be resolved in future research. To this point, it should be noted that there is some evidence that threat-related information from emotional stimuli, such as faces, might be detected pre-attentively through a process implicating the amygdala, and could contribute to an orienting bias for some emotional faces (Vuilleumier, Armony, Driver & Dolan, 2001). Nonetheless, Vuilleumier and associates also report that subsequent face processing requires attention (Vuilleumier et al, 2001; Holmes, Vuilleumier & Eimer, 2003). This is consistent with the results of Chapters 3 & 4, which indicate that faces are not processed automatically or even mandatory (for such claims see e.g. Farah, 1995;

Lavie, Ro & Russell, 2003), but are subject to capacity limits just like other classes of stimuli. Any attention bias for faces might therefore depend on similar variables as the processing of other objects. In line with this reasoning, it is worth recording that attentional biases have now been observed for many types of stimuli, including substance related pictorial cues in smokers and cannabis and alcohol users (e.g. Jones, Jones, Smith & Copley, 2003; Waters, Shiffman, Bradley & Mogg, 2003).

However, those findings do by no means contradict the current evidence for a general attention bias for faces. The present results indicate that this bias reflects both faces' ability to retain attention at locations of task-relevance but also to engage attentional resources at irrelevant locations, and when it might actually be beneficial to ignore them. Although the comparison stimuli displayed similar patterns, none of them were capable of matching or outperforming faces in any of the experiments. This is particularly striking as the comparisons included inverted faces with identical low-level image characteristics to these faces (Experiments 12 & 13), famous names giving the same status as persons as faces (Experiments 14 & 15), and meaningful nonface objects (Experiments 12-15). In addition, virtually identical results were obtained for unfamiliar faces, for which participants could not possess any higher-level identity or semantic knowledge, and for famous faces (cf. Experiments 12 & 13). Note also that previous studies investigating attention biases to faces may be criticized for using artificial face stimuli (e.g. Vuilleumier, 2000; Mack et al, 2002), and for presenting only a single face alongside several nonface objects (Ro et al, 2001) or within a visual stream of nonface objects (Mack et al, 2002), which may have produced an "odd-one-out" effect rather than a face advantage (see Palermo & Rhodes, 2003). By contrast, the current experiments used high quality photographs of real faces and were designed to eliminate odd-one-out effects by grouping stimuli from the same category (i.e. faces, names, flags) into blocks. The present findings thus provide perhaps the strongest evidence yet of a general attentional bias for faces.

Such a bias might provide a plausible explanation for some of the interference effects that were observed in Experiments 4-6. In these experiments, faces interfered more with nonface targets, which could be names (Experiment 4), famous names (Experiment 5), or flags (Experiment 6), than nonface distractors, but were generally also subject to less interference as targets. This could be explained in terms of a disengagement bias for face targets and an engagement bias for face distractors, whereby the processing of these stimuli was enhanced in each case to the detriment of the nonface comparisons. However, target-distractor interference tasks provide a more complex scenario than the current experiments, as the explicit processing demands of these stimuli must also be considered. Thus, Young et al (1986) found that names actually interfere *more* with faces than faces with names in target naming tasks. Although this accentuates the potentially intricate relationship between attention and other ongoing task demands, it also illustrates the benefit of the relative simplicity of the present paradigm, where subjects did not have to produce a response to the face and nonface stimuli. Finally, in preceding chapters it has already been discussed whether faces and other objects might have their own specific processing resources. Although some attention biases appear strikingly similar across different stimulus categories, including faces (cf. Ro et al, 2001 versus Palermo & Rhodes; Fox et al, 2001 versus Yiend & Mathews, 2001), a face bias might arise at least partly from such a division. All these points are considered further in the concluding chapter.

Chapter 6 Summary and Conclusions

The research carried out in this thesis investigated the relation of attention and face processing with emphasis on face encoding, capacity limitation, and attention biases. The introduction reviewed previous studies that are of relevance to these topics and identified a number of shortcomings. The first of these was a failure to examine the influence of attention on the encoding of different types of facial information, such as identity *and* expression. Although several studies have investigated the role of attention in face encoding, these have focused exclusively on identity information (Boutet & Chaudhuri, 2002; Palermo & Rhodes, 2002; Reinitz, Bartlett & Searcy, 1997; Reinitz, Morrissey & Demb, 1994). This issue was complicated further by controversial reports of an asymmetric processing dependency between identity and expression (Schweinberger, Burton & Kelly, 1999; Schweinberger & Soukup, 1998). Another shortcoming concerned the number of faces that can be processed simultaneously. There is good evidence that faces are processed even under conditions that should make this difficult, provided that only a single face is presented at a time. Consequently, it has been argued that face processing proceeds independent of general capacity limits, but that it may be subject to its own processing limits (e.g. Jenkins, Lavie & Driver, 2003; Lavie, Ro & Russell, 2003). However, even though a considerable number of studies have examined face processing in multiple-face displays, none of these apply a direct test for such limits (e.g. Boutet & Chaudhuri, 2001; Palermo & Rhodes, 2002; Jenkins et al, 2003). The final shortcoming that was targeted in this thesis concerned attention biases in visual processing. It has repeatedly been suggested that faces may have an advantage in capturing processing resources over other stimulus categories. This evidence is not completely compelling as it is based on artificial stimuli (e.g. Mack, Pappas, Silverman & Gay, 2002; Vuilleumier, 2000) and paradigms that are open to alternative explanations (e.g. Ro, Russell & Lavie, 2001; see Palermo & Rhodes, 2003). Moreover, although some classes of stimuli, including faces, appear particularly adept at retaining attention, often depending on their emotional connotation and the emotional state of the observer (e.g. Amir, Elias, Klump & Przeworski, 2003; Fox, Russo, Bowles & Dutton, 2001; Yiend & Mathews, 2001), it is presently unknown whether a general disengagement bias for faces exists relative to other classes of stimuli. This thesis offered novel approaches to each of these topics over a series of 15 experiments, by measuring task-relevant and irrelevant face processing in response-competition and priming tasks.

Chapter 2 began by examining the functional independence of facial identity and expression information. Previous work provides substantial evidence for this independence (e.g. Bruce, 1986; Calder, Young, Keane & Dean, 2000; Campbell, Brooks, de Haan & Roberts, 1996; Etcoff, 1984; Young, McWeeney, Hay & Ellis, 1986a). Recently however, this idea has been challenged by two studies reporting an asymmetric dependency in a Garner paradigm, such that expression decisions are influenced by task-irrelevant variations in identity information but not vice versa (Schweinberger & Soukup, 1998; Schweinberger et al,1999). In Chapter 2, a number of asymmetric treatment effects were identified that could have accounted for this pattern. These included picture-based strategies and the use of unreliable facial cues for decision-making, and effects of face familiarity.

Picture-based effects could have resulted from superficial image similarities such as brightness or colour contrasts, perhaps reflecting the conditions under which different stimuli were produced. Unreliable facial cues, on the other hand, could reflect the use of external features, such as hairstyle, which in many instances provide only a vague means of person identification. Both of these are less likely to correlate with expression than identity, but could have affected expression decisions even if these can usually be made without interference from identity. The influence of such cues might have been enforced particularly by the use of an extremely small stimulus set and a great number of trial repetitions (as in Schweinberger & Soukup, 1998; Schweinberger et al, 1999). Face familiarity, not explicitly specified by Schweinberger and associates (1998, 1999), might also relate to the potential influence of external features, which tend to contribute more to the recognition of unfamiliar than familiar faces. Alternatively, identity might not interfere with expression classification of unfamiliar faces if participants are incapable of distinguishing different stimulus identities. Thus, an asymmetric relationship might reflect only specific processes associated with either familiar or unfamiliar face processing.

To address these concerns, I used a variation of Schweinberger et al's (1998, 1999) task but with a substantially larger and more varied stimulus set, and by using participants who were unfamiliar (Experiment 1) and personally familiar (Experiment 2) with the stimulus identities. Both experiments generated virtually identical results: Identity decisions were completely unaffected by facial expression. Although expression decisions were generally slower than identity decisions, they were also unaffected by facial identity information. These results therefore contradict claims of an asymmetric dependency between expression and identity processing. Furthermore, expression classification times were faster when identity and expression information was correlated, than when both dimensions varied independently of each other. This suggests that, when possible, participants were making use of task-irrelevant identity information to facilitate the expression task. The results are therefore not only consistent with the idea that identity processing and expression processing are dissociable cognitive functions, but also support the view that they are processed in parallel (see e.g. Bruce & Young, 1986).

While Experiments 1 & 2 examined how expression and identity are deciphered from faces, rather little is known about how these different types of facial information are integrated within the same face percept during visual encoding. A few studies have implicated visual attention in the holistic encoding of facial identity (Palermo & Rhodes, 2002; Reinitz et al, 1994). Experiment 3 therefore examined whether attention is required to encode expression and identity within the same face. This was done by measuring response-competition from two distractor faces during the categorization of a face target, which was classified according to particular expression-identity conjunctions. The distractors were composed so that they contained expression and identity information from the same or the opposite response category as the target. These two types of facial information were either combined in one of the distractors (with the second distractor composed of response-neutral expression and identity information) or distributed across both. It was hypothesized that distractor interference should be equivalent for these conditions if information about the correct expression-identity conjunctions is unavailable from the distractors. Additionally, a few conditions were included in which the distractors only contained one type of response-critical information (i.e. identity *or* expression) or only response-neutral information. This was done to examine whether any target-distractor interference effects were due to expression *and* identity, or just one of these types of facial information.

The results were unexpected as Experiment 3 failed to produce reliable response-competition effects in all conditions. Of course, this means that this experiment fell short of its objective to investigate the encoding of expression and identity information. However, these results were particularly perplexing as reliable response-competition effects have been obtained previously from nonface conjunction objects in a similar design (Lavie, 1997), and also in letter-letter (e.g. Eriksen & Hoffman, 1972, 1973; Eriksen & Eriksen, 1974), picture-word (e.g. Smith & Magee, 1980) and face-name response-competition tasks (e.g. Jenkins et al, 2003; Lavie et al, 2003; Young et al, 1986). This contrast between previous studies and the present result suggested that it might be worthwhile to investigate the absence of distractor interference in Experiment 3 further.

Consequently, Chapter 3 examined response-competition between face targets and face distractors under conditions that normally allow for distractor interference. Thus the studies here compared speeded responses to target faces and nonface comparisons, which could be flanked by a face distractor or a nonface stimulus. Experiment 4 showed that face distractors interfere with forenames as targets, and that forenames also interfere with face targets and forename targets in a sex classification task. By contrast, no interference effects were found when a face distractor flanked a target face. Subsequent experiments generally replicated this pattern of results with famous faces and famous names (Experiment 5) and with famous faces and national flags (Experiment 6) in semantic classification tasks.

This remarkable absence of face-face interference was explored further in Experiments 7 & 8. In Experiment 7 participants again performed a semantic classification task on famous faces and national flags but the number of task-irrelevant distractors was now increased from one to four to boost the total of the potentially distracting information in each display. As before, the results showed response-competition from flag distractors with face and flag targets. However, distractor faces were still unable to influence responses to face targets. In fact, multiple distractors even failed to produce interference with nonface targets.

Experiment 8 then explored the temporal conditions under which interference from task-irrelevant face distractors is eliminated. The specific aim was to determine whether the ongoing processing of a face is sufficient to eliminate interference from a face distractor in a subsequent interference display. To this end, participants made occupational decisions to a name target (i.e. pop-star vs. politician) while ignoring a flanking face distractor, but only if the target-distractor display was preceded by a British (contrary to a non-British) face or flag cue. The results showed that distractor interference was not eliminated by processing a preceding flag cue, even when the flags were presented very briefly and immediately prior to the name-face displays. Name-face interference was also found following the relatively long presentation of a face cue. However, this interference was completely eliminated following a short-lived face cue. Chapter 3 thus provides considerable evidence that distractor *face* processing is eliminated in interference tasks by another *face*. This other face may take the form of a concurrently presented target (Experiments 4-6), or of additional face distractors (Experiment 7), or of a face that immediately precedes an interference display (Experiment 8). By contrast, the same faces were subject to reliable interference effects whenever they were paired with a nonface item, and none of the nonface comparisons were subject to analogous within-category effects.

Despite the absence of face-face interference in Chapter 3, visual stimuli may still undergo substantial processing, as indicated by priming effects, even when they do not give rise to response-competition (Driver & Tipper, 1989). Moreover, Chapter 3 only examined distractor processing with sex decisions (Experiment 4) and semantic decisions (Experiments 5-8). However, sex and identity information can be extracted independently from faces, and the retrieval of personal semantic information is relatively deep and follows face recognition. Thus, the possibility remained that distractor faces were processed at some level, perhaps even involving access to facial identity. For these reasons, Chapter 4 employed repetition priming to examine the processing of face distractors in two-item displays. In Experiments 9 & 10, subjects were shown displays of either a face or a flag target and a face distractor in a priming phase, followed by a familiarity task to primed and unprimed faces. Experiment 11 then presented face distractors alongside face targets and face-like and -unlike nonface targets in a similar task, to explore the influence of 'faceness' on distractor priming.

Experiment 9 found repetition priming for face distractors when they were presented alongside flag targets. This priming effect evidently involved access to the distractors' identities, surviving a change in image between prime and test phase. By contrast, processing a target face eliminated face distractor priming. Experiment 10 then found that distractor priming can be obtained in this condition when identical images are used at prime and test. Thus Experiment 10 demonstrated, for the first time, that face distractors undergo some processing even when they appear alongside a task-relevant face. Experiment 11 then showed that cross-image face distractor priming is not only eliminated during face target processing, but also by nonface objects that were rated as face-like prior to the experiment. Conversely, distractor priming was still obtained alongside face-unlike stimuli. Chapter 4 therefore provided the first instance of face distractor extinction in a nonface-face display (Experiment 11), but importantly, this appeared to reflect the perceived faceness of the target stimuli.

The findings of Chapters 3 & 4 may be explained if face processing is subject to capacity limits, such that only a single face can be processed at a time. If demand of processing resources cannot be met, for example, during the classification of a task-relevant face target, processing of another (distractor) face suffers as a result. This conclusion converges with a number of recent reports hinting at face processing limits, although none of these studies have examined this limit directly. Boutet & Chaudhuri (2001) used displays consisting of two overlapping face stimuli, a rather artificial situation that our perceptual system is not usually confronted with, and found that participants have difficulty in delineating more than one face at exposure times of less than 2 seconds. Palermo & Rhodes (2002) also used relatively long exposure times of ≥ 1.5 seconds for three-face displays in a divided attention paradigm. This may have permitted the serial processing of these faces, again making it difficult to specify an exact limit. On the contrary, Jenkins et al (2003) used short display times, but only measured task-irrelevant face processing. Without taking task-relevant resources into account, this also makes it difficult to draw a direct inference about face processing limits from this study. Consequently, the present data provides arguably the most direct evidence for face processing limits yet.

Although the studies reported in Chapters 3 & 4 were not explicitly designed to examine whether faces may have their own dedicated processing capacity, they also bear some relevance to the controversial issue of face-specificity. Several researchers have argued for the existence of face-specific neural mechanisms, which are held to operate to some extent independently of a more general object recognition system (e.g. Farah, 1995; Farah, Levinson & Klein, 1995; Farah, Wilson, Drain & Tanaka, 1995; Haxby et al, 1999; Kanwisher, McDermott & Chun, 1997; McNeill & Warrington, 1993). This dichotomy between face processing and nonface processing also receives support from response-competition and distractor priming studies. Lavie et al (2003) found that face distractors even interfere with name targets under an attentional load that is more than sufficient of eliminating response-competition from nonface distractors. Jenkins, Burton & Ellis (2002) obtained conceptually similar results in a repetition priming task. Jenkins et al (2003) also showed that interference from face distractors with name classification could be diluted by the presence of another face, but not by general competition from other stimulus classes. By the same token, face distractor processing proceeded seemingly unaffected by nonface targets in Chapters 3 & 4. The only exception comes from Experiment 11, in which face-like nonface targets behaved similarly to photographs of real faces. Moreover, nonface stimuli generally also interfered with face targets (Experiments 4, 6 & 7), and none of the nonface comparisons produced analogous within-category processing limits (Experiments 4-8). This could be interpreted as further evidence that face processing involves it own specific resources.

Alternatively, one could argue that faces constitute a visually more complex and homogeneous category of stimuli that is simply more demanding of general processing resources than printed names, and images of flags and (face-unlike) cars. In that case, comparable processing limits could be obtained for more complex nonface stimuli. However, the target RTs and error rates in Experiments 4-6 suggest that the processing of faces was generally no more difficult than that of other stimuli. Furthermore, the same-image priming effects of Experiment 10 showed that face distractors undergo some processing even when they appear alongside a task-relevant face. On its own, this finding might contradict the notion of face processing limits and also of a face-specific capacity. However, in the absence of cross-image priming in Experiment 9, this same-image face priming effect suggests that the distractors did not gain access to any putative face recognition system in Experiment 10, but perhaps only registered at a more general processing stage. This conclusion receives some support from same-image priming effects with unfamiliar faces (Khurana, 2000) and novel shapes (DeSchepper & Treisman, 1996), for which no recognizable descriptions could have existed prior to the priming phase. In addition, Experiment 11 showed that face distractor processing was only extinguished by face-like stimuli, the car fronts, but not their face-unlike opposites. These two categories presumably possess similar levels of visual complexity and homogeneity, although they were not matched for low-level characteristics, such as their component spatial frequencies. Yet, given the sheer number of car stimuli used, it would be remarkable if low-level features could account for this dissociation. Until further work is conducted, it thus seems most plausible that the processing limits from Chapters 3 & 4 were elicited by the face-ness of the stimuli.

Experiment 11 also suggests that face processing limits may be elicited by stimuli bearing only a remote, albeit measurable, resemblance to real faces. This is not unusual amid abundant claims that even artificial face stimuli, sometimes consisting of only a few lines, are capable of engaging face processes (e.g. Eastwood, Smilek & Merikle, 2001, 2003; Mack et al, 2002; Suzuki & Cavanagh, 1995; Vuilleumier, 2000). Indeed, it has been argued that any dedicated face mechanism may have no choice but to process other stimuli with geometrical features that are characteristically face-like (e.g. Pinker, 1997; see also Sperber, 1994). Nonetheless, face-like stimuli should at least be excluded from some face processing stages, such as identity analysis. To this point, DeGelder & Rouw (2001) recently proposed a dual-route model of face recognition. According to this account, face recognition entails a face detection system and a functionally distinct identification system. Whereas the latter system requires extensive learning to distinguish between different face exemplars, face detection is a cruder, innate mechanism. But are face-processing limits then located at a face detection stage or do they only affect subsequent processing stages?

On the one hand, the present face processing limits are of surprising severity. This is particularly so as face-face interference was even eliminated in a sex classification task, in which the faces preserved external cues such as hairstyle (Experiment 4). These salient cues could have been used to categorize the distractors without processing actual face information. The absence of cross-image distractor priming in Experiments 9 & 11, a seemingly highly sensitive measure of face processing (e.g. Bruce, Carson, Burton & Kelly, 1997; Brunas, Young & Ellis, 1990; Jenkins et al, 2002), also serves to strengthen this impression. Of course, similar limits might not apply in tasks that measure the ability to process multiple stimuli, rather than the ability to ignore a distractor. However, Lavie's (1995, 2000) perceptual load theory of selective attention states that spare capacity not consumed by relevant processing should automatically spill over to irrelevant distractor processing. As this was the case in all but the face-face conditions, the target-distractor paradigms of Chapters 3 & 4 should provide an accurate measure of face processing capacity, at least subsequent to face detection.

On the other hand, there are indications that several faces might be subject to some parallel processing. Jenkins et al (2003) found that an additional, unfamiliar face distractor reduced interference from a famous face distractor during the occupational classification of a target name. These dilution effects were face-specific, which suggests that both distractors were registered as faces at some level. This notion also receives some tentative support from the finding that four face distractors not only failed to interfere with the classification of face targets but also the nonface targets (images of flags) in Experiment 7. At first sight, it is hard to see why this might have happened. After all, when available, face-processing capacity always seemed to spill over to the distractors in two-item displays (Experiments 4-6, 8-11). Moreover, if participants had recognized just a single of the four distractors, then they might have produced the same strong nonface-face interference effects of other experiments. However, it may be that the four face distractors were competing for limited processing resources. This competition could have remained unresolved due to their equally task-irrelevant status, thereby preventing the recognition and necessary semantic analysis of any of the distractors. Intriguingly, this would imply that the present processing limits do not extend to face detection. Of course, this reasoning is clearly speculative, being largely post-hoc. Future work should therefore assess capacity limits in face detection, and in particular whether several task-irrelevant faces are detected alongside a nonface target.

Nonetheless, Experiments 4-6 provide some indication that faces really may be strong competitors for processing resources, if only in displays that contain a single face stimulus. In each of these experiments, faces interfered more with nonface targets than nonface distractors interfered with faces. This pattern could be explained if faces are encoded into a form that particularly suits sex and semantic categorization tasks. Names, for instance, are notoriously difficult to retrieve from faces, and previous studies that have shown this face advantage in semantic tasks have also found the reverse pattern in naming tasks (Young et al, 1986). However, this explanation is not entirely satisfactory as solitary face distractors still interfered more with national flags during nationality decisions than flags interfered with faces (Experiment 6). Surely, flags should be coded more readily into nationality than faces, which provide much more information than a person's nationality. Consequently, the final empirical chapter examined an alternative explanation for these interference patterns, namely whether faces are particularly adept at engaging and retaining attentional resources in comparison with other stimuli.

In Experiments 12-14, participants performed a simple classification task, in which they were required to focus on a central go/no-go signal before responding to the onscreen location of a peripheral line target (i.e. left versus right). The task thus required an attentional shift from the location of the go/no-go signal to that of the target. The results showed that target classification was delayed by the presence of a visual stimulus in the location of the go/no-go signal, such as an upside-down face or an image of a fruit. Importantly however, this effect was most pronounced when upright faces were used and occurred regardless of face familiarity (cf. Experiment 12 & 13). Experiment 14 replicated this face disadvantage in comparison with another class of meaningful nonface stimuli, in this case images of national flags, and showed that these results also cannot be attributed to the personhood of faces, as performance with the same person's names was indistinguishable from the flag condition. A variation of this paradigm was then carried out in which the target was always presented at fixation, directly beside the go/no-go signal, and the face and nonface stimuli appeared in the periphery (Experiment 15). In this study, the face and nonface stimuli thus never appeared in a task-relevant location. As in the three preceding experiments, it was found that faces delayed target RTs more than other classes of stimuli.

It is clear from these findings that faces are particularly efficient in retaining (Experiments 12-14) and engaging attention (Experiment 15) in comparison to other stimuli. This may have contributed to response-competition effects in single-distractor displays (Experiments 4-6), and also with multiple distractors (Experiment 7). However, it should be noted that these effects are qualitatively different from those of the response-competition tasks, in which distractor *congruency* was manipulated to assess any task-irrelevant processing. Thus, to the extent that the distractors were processed, these effects arose from conflicting information during response production. In Chapter 5, on the other hand, the faces (and nonface stimuli) were always response-neutral and presented in the context of a simple perceptual task. These tasks therefore assessed faces' ability to vie for outright control with the subjects' intentions, regardless of the nature of the required response.

According to one theory already highlighted in this discussion, faces may draw on a specific capacity with its own limits. The findings from Chapter 5 also raise the question whether faces have their own dedicated attentional resources. If faces draw on a separate processing capacity, then this might produce something akin to a capture advantage in multi-item displays, which contain several nonface items and only a single face (as in Ro et al, 2001). However, this cannot explain the face biases in Chapter 5, in which the face and nonface comparisons were never presented within the same display. This also cannot account for the interference patterns that were obtained when only a single face and a single nonface item were paired (Experiments 4-6). Moreover, there is some evidence that faces and other types of visual stimuli are subject to a common selection mechanism at some stage, as Jenkins et al (2002) observed impaired explicit memory for face distractors under high compared to low task-relevant nonface load. Crucially, this occurred in a context in which face distractors showed equivalent repetition priming, regardless of task load. Similarly, the error patterns in Chapter 5 indicate that presenting the faces did not affect the extent to which the line targets were processed, but rather, only the timing of the responses to these targets. Perhaps then face processing itself occurs independent of a general capacity, but there are some stages involved in controlling responses to and awareness of perceived stimuli, and hence explicit memory, that are not.

So far, this discussion has largely been concerned with face processing limits and attention biases to faces. I return now to the encoding of facial information, which is another question that this work originally set out to explore (see Experiment 3). The specific aim was to determine how functionally dissociable types of facial information, such as expression and identity, are combined within one face without being confused with those from another face. Although Experiment 3 produced little apparent success, the results of subsequent experiments are perhaps more decisive, despite not being motivated directly by this issue. Under the processing limits of Chapters 3 & 4, correctly integrating different types of information would be a natural consequence of available face processing resources. To this point, Experiment 4 demonstrates that face-face interference was even absent when face distractors preserved salient face-related response cues for a sex categorization task. In view of this particular result, but also the findings of Experiments 3-11 in general, it seems unlikely that facial information relating to expression, sex or identity was registered from distractors in face-face displays. Note that this reasoning does not contradict earlier claims that the face processing limits of Chapters 3 & 4 might not extend to face detection. Some prosopagnosics, for example, show impaired face identification but have normal face detection (de Gelder & Rouw, 2000). Moreover, Lewis & Edmonds (2003) showed that the eyes appear to form the single most important part of face detection. Lewis & Edmonds (2003) suggest that the eyes might constitute a characteristic pair of equally sized luminance regions that can be extracted quickly from an image. It is conceivable that this could occur even when processing limits make finer types of facial information unavailable.

In conclusion, this thesis applies a range of attentional paradigms to establish several new facts about face processing. Most importantly, the research carried out here provides evidence that face processing is subject to capacity limits, such that only a single face can be processed at a time. It also demonstrates attention retention and engagement biases for faces, in comparison with other stimulus classes. In doing so, the present findings raise several theoretical issues that need to be confirmed by future research. For instance, it may be that the current face processing limits do not extend to face detection. It would also be interesting to find out whether attention biases to a limited set of faces (as in Chapter 5) are long-lived or diminish with increasing repetition. Finally, on a more integrative level, one way to extend this research could be to explore possible interactions between processing limits and attention biases. It is conceivable, for example, that a face bias could be aimed at resolving competition between several simultaneously-presented faces, at facilitating the retrieval of identity and semantic information from a fixated face, or at monitoring faces for changes in transient information such as expression and facial speech.

References

Amir, N., Elias, J., Klumpp, H., & Przeworski, A. (2003). Attentional bias to threat in social phobia: facilitated processing of threat or difficulty disengaging attention from threat? *Behaviour Research & Therapy, 41*, 1325-1335.

Anstis, S. (1974). A chart demonstrating variation in acuity with retinal position. *Vision Research, 14*, 589-592.

Baddeley, A., & Weiskrantz, L. (1993). *Attention: Selection, Awareness, & Control.* UK: Oxford University Press.

Bentin, S., Deouell, L.Y., & Soroker, N. (1999). Selective visual streaming in face recognition: evidence from developmental prosopagnosia. *Neuroreport, 10*, 823-827.

Blaxton, T.A. (1989). Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. *Journal of Experimental Psychology: Learning, Memory & Cognition, 15*, 657-668.

Boutet, I., & Chaudhuri, A. (2001). Multistability of overlapped face stimuli is dependent upon orientation. *Perception, 30*, 743-753.

Boutet, I., Gentes-Hawn, A., & Chaudhuri, A. (2002). The influence of attention on holistic face encoding. *Cognition, 84*, 321-341.

Bradley, B.P., Mogg, K., Falla, S.J., & Hamilton, L.R. (1998). Attentional bias for threatening facial expressions in anxiety: Manipulation of stimulus duration. *Cognition & Emotion, 12*, 737-753.

Briand, K., & Klein, R.M. (1987). Is Posner's "beam" the same as Treisman's "glue"? On the relationship between visual orienting and feature integration theory.

Journal of Experimental Psychology: Human Perception & Performance, 13, 228-241.

Broadbent, D.E. (1958). *Perception and communication.* Oxford: Pergamon.

Broadbent, D., & Broadbent, M. (1988). Anxiety and attentional bias: State and trait. *Cognition & Emotion, 2*, 165-183.

Brown, V., Huey, D., & Findlay, J.M. (1997). Face detection in peripheral vision: do faces pop out? *Perception, 26*, 1555-1570.

Bruce, V. (1986). Influences of familiarity on the processing of faces. *Perception, 15*, 387-397.

Bruce, V. (1988). *Recognizing faces.* UK: Lawrence Erlbaum Associates Ltd.

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

Bruce, V., Ellis, H.D., Gibling, F., & Young, A.W. (1987). Parallel processing of the sex and familiarity of faces. *Canadian Journal of Psychology, 41*, 510-520.

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

Bruce, V., Henderson, Z., Newman, C., & Burton, A.M. (2001). Matching identities of familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental Psychology: Applied, 7*, 207-218.

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

Bruce, V., & Young, A. (1998). *In the eye of the beholder: The science of face perception*. UK: Oxford University Press.

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

Brunas, J., Young, A.W., & Ellis, A.W. (1990). Repetition priming from incomplete faces: Evidence for part to whole completion. *British Journal of Psychology, 81*, 43-56.

Bundesen, C., Kyllingsbaek, S., Houmann, K.J., & Jensen, R.M. (1997). Is visual attention automatically attracted by one's own name? *Perception & Psychophysics, 59*, 714-720.

Burton, A.M. (1998). A model of human face recognition. In J. Grainger & A.M. Jacobs (Eds.), *Localist connectionist approaches to human cognition* (pp. 75-100). Hillsdale, NJ: Erlbaum.

Burton, A.M., & Bruce, V. (1993). Naming faces and naming names: exploring an interactive activation model of person recognition. *Memory, 1*, 457-480.

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

Burton, A.M., Bruce, V., & Hancock, P.J.B. (1999). From pixels to people: a model of familiar face recognition. *Cognitive Science, 23*, 1-31.

Burton, A.M., Kelly, S.W., & Bruce, V. (1998). Cross-domain repetition priming in person recognition. *The Quarterly Journal of Experimental Psychology, 51A*, 515-529.

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

Burton, A.M., Young, A.W., Bruce, V., Johnston, R.A., & Ellis, A.W. (1991). Understanding covert recognition. *Cognition, 39*, 129-166.

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

Campbell, R., Brooks, B., De Haan, E.H.F., & Roberts, T. (1996). Dissociating face processing skills: Decisions about lip-read speech, expression, and identity. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 49A*, 295-314.

Campbell, R., De Gelder, B., & De Haan E.H.F. (1996). The lateralization of lipreading: A second look. *Neuropsychologia, 34*, 1235-1240.

Campbell, R., Landis, T., & Regard, M. (1986). Face recognition and lipreading. A neurological dissociation. *Brain, 109*, 509-521.

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

Curcio, C.A., & Allen, K.A. (1990). Topography of ganglion cells in human retina. *Journal of Comparative Neurology, 300*, 5-25.

De Gelder, B., & Rouw, R. (2000). Configural processes in acquired and developmental prosopagnosia: evidence for two separate systems? *Neuroreport, 11*, 3145-3150.

De Gelder, B., & Rouw, R. (2001). Beyond localization: a dynamical dual route account of face recognition. *Acta Psychologica, 107*, 183-207.

De Haan, E.H.F., Young, A.W., & Newcombe, F. (1987). Faces interfere with name classification in a prosopagnosic patient. *Cortex, 23*, 309-316.

De Schepper, B., & Treisman, A. (1996). Visual memory for novel shapes: implicit coding without attention. *Journal of Experimental Psychology: Learning, Memory & Cognition, 22*, 27-47.

Deutsch, J.A., & Deutsch, D. (1963). Attention, some theoretical considerations. *Psychological Review, 70*, 80-90.

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

Driver, J. (2001). A review of selective attention research from the past century. *British Journal of Psychology, 92*, 53-78.

Driver, J., & Tipper, S.P. (1989). On the Nonselectivity of "Selective" Seeing: Contrasts between Interference and Priming in Selective Attention. *Journal of Experimental Psychology: Human Perception & Performance, 15*, 304-314.

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review, 37*, 272-300.

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

Eastwood, J.D., Smilek, D., & Merikle, P.M. (2001). Differential attentional guidance by unattended faces expressing positive and negative emotion. *Perception & Psychophysics, 63*, 1004-1013.

Eastwood, J.D., Smilek, D., & Merikle, P.M. (2003). Negative facial expression captures attention and disrupts performance. *Perception & Psychophysics, 65*, 352-358.

Ekman, P., & Friesen, W.V. (1976). *Pictures of facial affect*. Consulting Psychologists Press, Palo Alto, CA.

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

Ellis, A.W., Flude, B.M., Young, A.W., & Burton, A.M. (1996). Two loci of repetition priming in the recognition of familiar faces. *Journal of Experimental Psychology: Learning, Memory & Cognition, 22*, 295-308.

Ellis, A.W., Young, A.W., & Flude, B.M. (1990). Repetition priming and face processing: Priming occurs within the system that responds to the identity of a face. *Quarterly Journal of Experimental Psychology, 42A*, 495-512.

Ellis, A.W., Young, A.W., Flude, B.M., & Hay, D.C. (1987). Repetition priming of face recognition. *The Quarterly Journal of Experimental Psychology, 39A*, 193-210.

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

Ellis, H.D., Young, A.W., & Koenken, G. (1993). Covert face recognition without prosopagnosia. *Behavioural Neurology, 6*, 27-32.

Eriksen, B.A., & Eriksen, C.W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics, 16*, 143-149.

Eriksen, C.W., & Hoffman, J.E. (1972). Temporal and spatial characteristics of spatial encoding from visual displays. *Perception & Psychophysics, 12*, 201-204.

Eriksen, C.W., & Hoffman, J.E. (1973). The extent of processing of noise elements during selective encoding from visual displays. *Perception & Psychophysics, 14*, 155-160.

Eriksen, C.W., & Schultz, D.W. (1979). Information processing in visual search: A continuous flow conception and experimental results. *Perception & Psychophysics, 25*, 249-263.

Etcoff, N.L. (1984). Selective attention to facial identity and facial emotion. *Neuropsychologia, 22*, 281-295.

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

Farah, M.J. (1995). Dissociable systems for recognition: A cognitive neuropsychology approach. In S.M. Kosslyn & D.N. Osherson (Eds.), *An invitation to cognitive science: Visual cognition* (Vol.2). Cambridge, MA: MIT Press.

Farah, M.J., Levinson, K.L., & Klein, K.L. (1995). Face perception and within-category discrimination in prospagnosia. *Neuropsychologia, 33*, 661-674.

Farah, M.J., Wilson, K.D., Drain, H.M., & Tanaka, J.R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision Research, 35*, 2089-2093.

Fox, E., Lester, V., Russo, R., Bowles, R.J., Pichler, A., & Dutton, K. (2000). Facial expressions of emotion: Are angry faces detected more efficiently? *Cognition & Emotion, 14*, 61-92.

Fox, E., Russo, R., Bowles, R., & Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety? *Journal of Experimental Psychology: General, 130*, 681-700.

Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition & Emotion, 16*, 355-379.

Friedman-Hill, S.R., Robertson, L.C., & Treisman, A. (1995). Parietal contributions to visual feature binding: Evidence from a patient with bilateral lesions. *Science, 269*, 853-855.

Garner, W.R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.

Garner, W.R. (1976). Interaction of stimulus dimensions in concept and choice processes. *Cognitive Psychology, 8*, 98-123.

Gatti, S.V., & Egeth, H.E. (1978). Failure of spatial selectivity in vision. *Bulletin of the Psychonomic Society, 11*, 181-184.

Gauthier, I., & Logothetis, N.K. (2000). Is face recognition not so unique after all? *Cognitive Neuropsychology, 17*, 125-142.

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

Gauthier, I., Williams, P., Tarr, M.J., & Tanaka, J. (1998). Training "greeble" experts: a framework for studying expert object recognition processes. *Vision Research, 38*, 2401-2428.

George, M.S., Ketter, T.A., Gill, D.S., Haxby, J.V., Ungerleider, L.G., Herscovitch, P., & Post, R.M. (1993). Brain regions involved in recognizing facial emotion or identity: An oxygen-15 PET study. *Journal of Neuropsychiatry, 5*, 384-394.

Goshen-Gottstein, Y., & Ganel, T. (2000). Repetition priming for familiar and unfamiliar faces in a sex-judgement task: evidence for a common route for the processing of sex and identity. *Journal of Experimental Psychology: Learning, Memory & Cognition, 26*, 1198-1214.

Green, K.P., Kuhl, P.K., Meltzoff, A.N., & Stevens, E.B. (1991). Integrating speech information across talkers, gender, and sensory modality: Female faces and male voices in the McGurk effect. *Perception & Psychophysics, 17*, 278-288.

Hagenaar, R., & Van der Heijden, A.H.C. (1986). Target-noise separation in visual selective attention. *Acta Psychologica, 62*, 161-176.

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

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

Harmon, L.D. (1973). The recognition of faces. *Scientific American, 227*, 71-82.

Harris, C.R., & Pashler, H.E. (2004). Attention and the Processing of Emotional Words and Names. *Psychological Science, 15*, 171-178.

Harris, C.R., Pashler, H.E., & Coburn, N. (2004). Moray revisited: High-priority affective stimuli and visual search. *The Quarterly Journal of Experimental Psychology, 57A*, 1-31.

Haxby, J.V., Hoffman, E.A., & Gobbini, M.I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*, 223-233.

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

Heathcote, A., & Mewhort, D.J.K. (1993). Representation and selection of relative position. *Journal of Experimental Psychology: Human Perception & Performance, 19*, 488-516.

Holmes, A., Vuilleumier, P., & Eimer, M. (2003). The processing of emotional facial expression is gated by spatial attention: evidence from event-related brain potentials. *Cognitive Brain Research, 16*, 174-184.

Humphreys, G.W., Cinel, C., Wolfe, J., Olson, A., & Klempen, N. (2000). Fractionating the binding process. *Vision Research, 40*, 1569-1596.

Humphreys, G.W., Donnelly, N., & Riddoch, M.J. (1993). Expression is computed separately from facial identity, and is computed separately for moving and static faces: Neuropsychological evidence. *Neuropsychologia, 31*, 173-181.

Humphreys, G.W., & Rumiati, R.I. (1998). Agnosia without prosopagnosia or alexia: Evidence for stored visual memories specific to objects. *Cognitive Neuropsychology, 15*, 243-277.

Humphreys, G.W., Quinlan, P.T., & Riddoch, M.J. (1989). Grouping processes in visual search: Effects with single- and combined-feature targets. *Journal of Experimental Psychology: General, 118*, 258-279.

Jacoby, L.L. (1983). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory & Cognition, 9*, 21-38.

Jacoby, L.L., & Brooks, L.R. (1984). Nonanalytic cognition: Memory, perception and conceptual learning. In G.H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 18 (pp. 1-47). New York: Academic Press.

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

Jenkins, R., Lavie, N., & Driver, J. (2003). Ignoring famous faces: Category-specific dilution of distractor interference. *Perception & Psychophysics, 65*, 298-309.

Johnson, M.H., Dziurawiec, S., Ellis, H., & Morton, J. (1991). Newborns Preferential Tracking of Face-Like Stimuli and its Subsequent Decline. *Cognition, 40*, 1-19.

Jones, B.T., Jones, B.C., Smith, H., Copley, N. (2003). A flicker paradigm for inducing change blindness reveals alcohol and cannabis information processing biases in social users. *Addiction, 98*, 235-244.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye. In J. Long & A. Baddeley (Eds.), *Attention & Performance IX* (pp. 187-203). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Kanwisher, N. (2000). Domain specificity in face perception. *Nature Neuroscience, 17*, 4302-4311.

Kanwisher, N., McDermott, J., & Chun, M.M. (1997). The fusiform face area: a module in the human extrastriate cortex specialized for face perception. *The Journal of Neuroscience, 17*, 4302-4311.

Khurana, B. (2000). Not to be and then to be: visual representation of ignored unfamiliar faces. *Journal of Experimental Psychology: Human Perception & Performance, 26*, 246-263.

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

Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 451-468.

Lavie, N. (1997). Visual feature integration and focused attention: Response competition from multiple distractor features. *Perception & Psychophysics, 59,* 543-556.

Lavie, N. (2000). Selective attention and cognitive control: Dissociating attentional functions through different types of load. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention & Performance XVIII* (pp. 175-194). Cambridge, MA: MIT Press.

Lavie, N. (2001). Capacity limits in selective attention: Behavioural evidence and implications for neural activity. In J. Braun, C. Koch, & J.L. Davis (Eds.), *Visual attention and cortical circuits* (pp. 49-68). Cambridge, MA: MIT Press.

Lavie, N., & Cox, S. (1997). On the efficiency of visual selective attention: Efficient visual search leads to inefficient distractor rejection. *Psychological Science, 8,* 395-398.

Lavie, N., & Fox, E. (2000). The role of perceptual load in negative priming. *Journal of Experimental Psychology: Human Perception & Performance, 26,* 1038-1052.

Lavie, N., Ro, T., & Russell, C. (2003). The role of perceptual load in processing distractor faces. *Psychological Science, 14,* 510-515.

Lavie, N., & Tsal, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. *Perception & Psychophysics, 56,* 183-197.

Le Gal, P.M., & Bruce, V. (2002). Evaluating the independence of sex and expression in judgments of faces. *Perception & Psychophysics, 64,* 230-243.

Leder, H. (1996). Line drawings of faces reduce configural processing. *Perception, 25,* 355-366.

Lewis, M.B., & Edmonds, A.J. (2003). Face detection: Mapping human performance. *Perception, 32*, 903-920.

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

Logan, G. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *CognitivePsychology, 12*, 523-553

Mack, A., Pappas, Z., Silverman, M., & Gay, R. (2002). What we see: Inattention and the capture of attention by meaning. *Consciousness & Cognition, 11*, 488-506.

Mack, A., & Rock, I. (1998). *Inattentional blindness.* Cambridge, MA: MIT Press.

MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology, 95*, 15-20.

McNeill, A., & Burton, A.M. (2002). The locus of semantic priming effects in person recognition. *The Quarterly Journal of Experimental Psychology, 55A*, 1141-1156.

McNeill, A., Burton, A.M., & Ellis, A.W. (2003). When sex isn't special: Priming onto a sex decision. *Visual Cognition, 10*, 641-650.

McNeill, J.E., & Warrington, E.K. (1993). Prosopagnosia - a face-specific disorder. *The Quarterly Journal of Experimental Psychology, 46*, 1-10.

Merikle, P.M., & Gorewich, N.J. (1979). Spatial selectivity in vision: Field size depends on noise size. *Bulletin of the Psychonomic Society, 14*, 343-346.

Miller, J.O. (1987). Priming is not necessary for selective-attention failures: Semantic effects of unattended, unprimed letters. *Perception & Psychophysics, 41*, 419-434.

Mogg, K., & Bradley, B. (1999). Selective attention and anxiety: A cognitive-motivational perspective. In T. Dalgleish & M. Power (Eds.), *Handbook of cognition and emotion* (pp. 145-170). Chichester, England: Wiley.

Morrison, D.J., Bruce, V., & Burton, A.M. (2000). Covert face recognition in neurologically intact participants. *Psychological Research, 63*, 83-94.

Morton, J., & Johnson, M.H. (1991). CONSPEC and CONLERN: A Two-Process Theory of Infant Face Recognition. *Psychological Review, 98*, 164-181.

Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience, 9*, 555-604.

Näsänen, R., & Ojanpää, H. (2004). How many faces can be processed during a single eye fixation? *Perception, 33*, 67-77.

Norman, D.A. (1968). Toward a theory of memory and attention. *Psychological Review, 75*, 522-536.

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

Palermo, R., & Rhodes, G. (2002). The influence of divided attention on holistic face perception. *Cognition, 82*, 225-257.

Palermo, R., & Rhodes, G. (2003). Change detection in the flicker paradigm: Do faces have an advantage? *Visual Cognition, 10*, 683-713.

Parry, F.M., Young, A.W., Saul, J.S., & Moss, A. (1991). Dissociable face processing impairments after brain injury. *Journal of Clinical and Experimental Neuropsychology, 13*, 545-558.

Pashler, H.E. (1998). *The Psychology of Attention.* Cambridge, MA: MIT Press.

Pinker, S. (1997). *How the mind works.* New York: Norton.

Posner, M.I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology, 32*, 3-25.

Posner, M.I., & Petersen, S.E. (1990). The attention systems of the human brain. *Annual Review of Neuroscience, 13*, 25-42.

Posner, M.I., Snyder, C.R.R., & Davidson, B.J. (1990). Attention and the detection of signals. *Journal of Experimental Psychology: General, 109*, 160-174.

Rees, G., Frith, C.D., & Lavie, N. (1997). Modulating irrelevant motion perception by varying attentional load in an unrelated task. *Science, 278*, 1616-1619.

Reinitz, M.T., Bartlett, J.C., & Searcy, J. (1997, November). *Role of attention in encoding facial features and their spatial relations.* Paper presented at the Psychonomic Society Conference, Philadelphia, PA.

Reinitz, M.T., Morrissey, J., & Demb, J. (1994). Role of attention in face encoding. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 20*, 161-168.

Remington, R.W., Johnston, J.C., & Yantis, S. (1992). Involuntary attentional capture by abrupt onsets. *Perception & Psychophysics, 51*, 279-290.

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

Ro, T., Russell, C., & Lavie, N. (2001). Changing faces: A detection advantage in the flicker paradigm. *Psychological Science, 12*, 94-99.

Roediger, H.L. (1990). Implicit Memory: Retention without remembering. *American Psychologist, 45*, 1043-1056.

Rossion, B. (2002). Is sex categorization from faces really parallel to face recognition? *Visual Cognition, 9*, 1003-1020.

Rumiati, R.I., & Humphreys, G.W. (1997). Visual agnosia without alexia or prosopagnosia: Arguments for separate knowledge stores. *Visual Cognition, 4*, 207-217.

Scapinello, K.F., & Yarmey, A.D. (1970). The role of familiarity and orientation in immediate and delayed recognition of pictorial stimuli. *Psychonomic Science, 21*, 329-330.

Schmitt, M., Postma, A., & De Haan, E.H.F. (2000). Interactions between exogenous auditory and visual spatial attention. *The Quarterly Journal of Experimental Psychology, 53A*, 105-130.

Schmitt, M., Postma, A., & De Haan, E.H.F. (2001). Cross-modal exogenous attention and distance effects in vision and hearing. *European Journal of Cognitive Psychology, 13*, 343-368.

Schweich, M., & Bruyer, R. (1993). Heterogeneity in the cognitive manifestations of prosopagnosia – The study of a group of single cases. *Cognitive Neuropsychology, 10*, 529-547.

Schweinberger, S.R., Burton, A.M., & Kelly, S.W. (1999). Asymmetric dependencies in perceiving identity and emotion: Experiments with morphed faces. *Perception & Psychophysics, 61*, 1102-1115.

Schweinberger, S.R., & Soukup, G.R.. (1998). Asymmetric relationships among perceptions of facial identity, emotion, and facial speech. *Journal of Experimental Psychology: Human Perception & Performance, 24*, 1748-1765.

Schweinberger, S.R., Klos, T., & Sommer, W. (1995). Covert face recognition in prosopagnosia: A dissociable function? *Cortex, 31*, 521-536.

Sergent, J., & Signoret, J.L. (1992a). Varieties of functional deficits in prosopagnosia. *Cerebral Cortex, 2*, 375-388.

Sergent, J., & Signoret, J.L. (1992b). Implicit access to knowledge derived from unrecognized faces in prosopagnosia. *Cerebral Cortex, 2*, 389-400.

Sergent, J., Ohta, S., MacDonald, B., & Zuck, E. (1994). Segregated processing of facial identity and emotion in the human brain: A PET study. *Visual Cognition, 1*, 349-369.

Shapiro, K.L., Caldwell, J., & Sorensen, R.E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology: Human Perception & Performance, 23*, 504-514.

Simons, D.J., & Levin, D.T. (1998). Failure to detect changes to people during real-world interaction. *Psychonomic Bulletin & Review, 5*, 644-649.

Smith, M.C., & Magee, L.E. (1980). Tracing the time course of picture-word processing. *Journal of Experimental Psychology: General, 109*, 373-392.

Smith, S.L. (1962). Color coding and visual search. *Journal of Experimental Psychology, 64*, 434-440.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L.A. Hirschfeld & S. Gelman (Eds.), *Mapping the mind:*

Domain specificity in cognition and culture (pp. 37-67). Cambridge, UK: Cambridge University Press.

Styles, E.A. (1997). *The Psychology of Attention.* UK: Psychology Press Ltd.

Suzuki, S., & Cavanagh, P. (1995). Facial organization blocks access to low-level features: an object inferiority effect. *Journal of Experimental Psychology: Human Perception & Performance, 21*, 901-913.

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

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

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics, 51*, 599-606.

Theeuwes, J. (1991). Cross-dimensional perceptual selectivity. *Perception & Psychophysics, 50*, 184-193.

Theeuwes, J. (1994). Stimulus-driven capture and attentional set: Selective search for color and visual abrupt onsets. *Journal of Experimental Psychology: Human Perception & Psychophysics, 20*, 700-806.

Theeuwes, J., De Vries, G.-J., & Godijn, R. (2003). Attentional and oculomotor capture with static singletons. *Perception & Psychophysics, 65*, 735-746.

Tipper, S.P. (1985). The negative priming effect: Inhibitory effects of ignored primes. *Quarterly Journal of Experimental Psychology, 37A*, 571-590.

Tipper, S.P., & Cranston, M. (1985). Selective attention and priming: Inhibitory and facilitatory effects of ignored primes. *Quarterly Journal of Experimental Psychology, 37A*, 591-611.

Tipper, S.P., & Driver, J. (1988). Negative Priming between pictures and words in a selective attention task: Evidence for semantic processing of ignored stimuli. *Memory & Cognition, 16*, 64-70.

Treisman, A. (1988). Features and objects: The fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology, 40A*, 201-237.

Treisman, A. (1993). Representing visual objects. *Attention & Performance, 14*, 163-175.

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

Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14*, 107-141.

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

Van Honk, J., Tuiten, A., De Haan, E.H.F., van den Hout, M., & Stam, H. (2001). Attentional biases for angry faces: Relationships to trait anger and anxiety. *Cognition & Emotion, 15*, 279-297.

Vuilleumier, P. (2000). Faces call for attention: evidence from patients with visual extinction. *Neuropsychologia, 38*, 693-700.

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.

Waters, A.J., Shiffman, S., Bradley, B.P., & Mogg, K. (2003). Attentional shifts to smoking cues in smokers. *Addiction, 98*, 1409-1417.

Williams, J.M.G., Mathews, A., & MacLeod, C. (1996). The emotional Stroop task and psycho-pathology. *Psychological Review, 120*, 3-24.

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

Wolford, G., & Morrison, F. (1980). Processing of unattended visual information. *Memory & Cognition, 8*, 521-527.

Yantis, S., & Johnston, D.N. (1990). On the locus of visual selection: Evidence from focused attention tasks. *Journal of Experimental Psychology: Human Perception & Performance, 16*, 135-149.

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

Yiend, J., & Mathews, A. (2001). Anxiety and attention to threatening pictures. *The Quarterly Journal of Experimental Psychology, 54A*, 665-681.

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

Yin, R.K. (1970). *Face recognition: A special process?*, Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge.

Young, A.W. (1998). *Face and mind.* UK: Oxford University Press.

Young, A.W., & Burton, A.M. (1999). Simulating face recognition: implications for modelling cognition. *Cognitive Neuropsychology, 16*, 1-48.

Young, A.W., Ellis, A.W., Flude, B.M., McWeeney, K.H., & Hay, D.C. (1986). Face-name interference. *Journal of Experimental Psychology: Human Perception and Performance, 12*, 466-475.

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

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

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

Young, A.W., McWeeney, K.H., Hay, D.C., & Ellis, A.W. (1986b). Access to identity-specific semantic codes from familiar faces. *The Quarterly Journal of Experimental Psychology, 38A*, 271-295.

Young, A.W., Newcombe, F., De Haan, E.H.F., Small, M., & Hay, D. (1993). Face perception after brain injury: Selective impairments affecting identity and expression. *Brain, 116*, 941-960.

Appendix A

FACE STIMULI USED IN EXPERIMENT 7

American:	British:
Woody Allen	Billy Connolly
Groucho Marx	John Cleese
Tom Cruise	Hugh Grant
Robert de Neiro	Sean Connery
Elvis Presley	Liam Gallagher
Eminem	Mick Jagger
George W. Bush	John Major
Bill Clinton	Tony Blair
Pete Sampras	David Beckham
Andre Agassi	Tim Henman

Appendix B

FACE STIMULI USED IN EXPERIMENTS 9 & 10

American:	British:
Al Pacino	Ali G
Angelina Jolie	Catherine Zeta-Jones
Ben Affleck	Chris Evans
Brad Pitt	Chris Tarrant
Britney Spears	Cilla Black
Bruce Willis	Craig David
Cameron Diaz	David Beckham
Courtney Cox	Davina McCall
Leonardo Di Caprio	Lady Di
Drew Barrymore	Duncan James
Eminem	Emma Bunton
George Clooney	Ewan McGregor
Halle Berry	George Michael
Harrison Ford	Geri Halliwell
Jack Nicholson	Graham Norton
Janet Jackson	Hugh Grant
Jennifer Aniston	Jarvis Cocker
Jennifer Lopez	Jonathan Ross
Jim Carrey	Judy Dench
Josh Hartnett	Kate Winslet
Julia Roberts	Liam Gallagher
Keanu Reeves	Liz Hurley
Kevin Costner	Michael Owen
Kevin Spacey	Michael Caine
Mariah Carey	Michael Palin
Marilyn Monroe	Paul McCartney
Matt Damon	Pierce Brosnan
Matt leBlanc	Robbie Coltrane
Mel Gibson	Robbie Williams
Mike Myers	Roger Moore
Nicolas Cage	Ronan Keating
Penelope Cruz	Ross Kemp
Richard Gere	Rowan Atkinson
Robert Redford	Sean Connery

Sandra Bullock	Sporty Spice
Sarah Jessica-Parker	Stephen Gately
Sarah Michelle Gellar	Victoria Beckham
Tom Cruise	Vinnie Jones
Val Kilmer	Zoe Ball
Will Smith	Will Young

Appendix C

FACE STIMULI USED IN EXPERIMENT 11 WERE THE SAME AS FOR EXPERIMENTS 9 & 10 IN ADDITION TO THE FOLLOWING:

American:	British:
David Schwimmer	Anne Robinson
JohnTravolta	Elton John
Justin Timberlake	Gareth Gates
Liv Tyler	Jamie Oliver
Madonna	John Major
Mark Wahlberg	Mick Jagger
Pamela Anderson	Prince Charles
Paul Newman	Richard Branson
Tom Hanks	Sting
Woody Allen	Tony Blair