Attentional biases in sub-clinical anxiety

Martin Joseph Batty

The University of Southampton

Thesis presented to the Faculty of Medicine, Health and Life Sciences for the degree of Doctor of Philosophy.

January 2004.

UNIVERSITY OF SOUTHAMPTON
ABSTRACT
FACULTY OF MEDICINE, HEALTH AND LIFE SCIENCES
SCHOOL OF PSYCHOLOGY
Doctor of Philosophy
ATTENTIONAL BIASES IN SUB-CLINICAL ANXIETY
By Martin Joseph Batty

There is agreement that the fast detection of threat might be evolutionarily adaptive. Cognitive models of visual attention suggest that items can only be detected preattentively if they contain basic features that can be used to guide visual search. It is assumed that search efficiency is determined only by features shared between targets and distractors, and that the visual system cannot make use of a category as broad as threat in order to guide search. On the other hand, clinical theories of anxiety purport that individuals preferentially attend towards threat-related stimuli. Some clinical theories even assume that threat can be captured at a preattentive level.

Previous studies have tried to examine this issue using visual search paradigms in which threat-associated targets were embedded amongst neutral distractors. However, since realistic pictures were used, close experimental control of low-level physical properties was not possible. The experiments presented here were designed to address experimental problems due to uncontrolled low level features by using highly controlled schematic stimuli. The emotional valance of these stimuli was changed by a conditioning procedure in which they became associated with either neutral or negative pictures. In order to test the effectiveness of this conditioning, subjects undertook a behavioural task (Implicit Association Test; IAT), which measured implicit valence to exclude the possibility of a null result simply reflecting ineffective conditioning.

In Experiment 1, 32 participants searched for either threat-associated or neutral targets. There was no advantage for the threat-paired relative to the neutral-paired targets, although a non-significant result in the IAT meant that findings were inconclusive. To increase the effectiveness of the conditioning, Experiment 2 used a within-subjects design. 46 students searched for both threat-associated and neutral conditioned targets. Again, there were no differences in search for the neutral and threat-paired targets. Furthermore, search was unaffected by anxiety level. However, conditioning was significant as evidenced by a main effect of response congruity (subjects were faster to associate the negatively paired target with bad words than good words). As phylogenetic based theories of fear conditioning purport that snakes and spiders are particularly potent sources of threat, in Experiment 3, 63 undergraduates screened for high and low snake and spider fear searched for snake or spider and neutral-associated targets. Again, there were no differences in search efficiency across target type and fear level, although importantly, once more there was a significant effect of conditioning.

As studies 1-3 showed no evidence of preattentive biases towards threat-associated stimuli, Experiment 4 used a cueing paradigm to investigate whether biases in attention reflect difficulty in disengaging attention from threatening stimuli once identified. 48 undergraduates pressed buttons corresponding to the direction of targets that were either validly or invalidly cued. Cues were either negative or neutral pictures. Subjects were significantly slower to respond to invalid negative relative to invalid neutral cues, reflecting a difficulty in disengaging attention from threat pictures. However, in a fifth experiment that used conditioned abstract cues instead of pictures, (using 48 undergraduates), subjects showed no significant differences in their RTs to both negative and neutral conditioned cues.

In sum, three experiments found no evidence that threat-associated stimuli controlled for low-level features could be detected preattentively. However, there was evidence that subjects do find it harder to disengage attention from a negative stimulus relative to a neutral stimulus, providing the stimulus is complex enough to warrant further processing.

Contents

Chapter 4: Experiment 1

Chapter 5: Experiment 2

Chapter 6: Experiment 3

Chapter 7: Experiment 4

Chapter 8: Experiment 5

Chapter 9: General Discussion

List of Tables

List of Figures

Acknowledgements

I owe a huge intellect debt to Dr Kyle Cave and Professor Paul Pauli for their outstanding input and supervision over the course of this thesis. Thanks also to Professors Karin Mogg, Brendan Bradley and Jim Stevenson for invaluable discussions and Drs Jan de Houwer, Aiden Gregg and Steven Glautier for advice and assistance in the development of conditioning paradigms. Any errors or misrepresentation of any of the above peoples' ideas is of course entirely my fault.

Finally, on a more personal note, I would like to express my sincere gratitude to my family and friends (especially Beth), for all their love and support.

Chapter 1:An Overview of Visual Attention

"Everybody knows what attention is" -William James (1890)

1.1 Introduction

The present chapter is a brief history of the most important discoveries in attention research over the past hundred years. Although findings from auditory attention are discussed, particularly with reference to the early/late selection debate, the emphasis of this chapter (and indeed the whole thesis) centres firmly upon visual attention. Attention is one of the most thoroughly studied aspects of cognition, with experimental paradigms dating back to the nineteenth century and the work of James, Helmholtz and Titchener. As the twentieth century dawned, attention remained a popular topic within cognitive psychology. Titchener (1908) for example, defined attention in terms of exogenous (or automatic), stimulus driven control, where sudden changes or movement, including pitch, distract the individual from the task that they are performing. The distinction between exogenous and endogenous (or voluntary) orienting of attention continues today and will be discussed in more depth later in the chapter.

1.2 What is Visual Attention and How is it Measured?

Research over the past century has seen a decline and resurgence of interest in the field of attention. Despite the quote by William James at the head of this page, attention, like emotion is arguably *not* very well understood by everyone, psychologists included. Like emotion, attention should not be regarded as a unitary concept. Rather the term attention can be used to describe a whole plethora of psychological processes. Over the course of this thesis, unless otherwise stated, the term attention is used to refer primarily to visual attention, although reference will obviously be made to other facets of attention, particularly the early experimental work in the field of auditory attention. Attention in its most general sense is often equated with the ability to focus on the relevant and to ignore the irrelevant (Medin & Ross, 1996). Put another way, attention can be regarded as a selection process, whereby certain stimuli are given precedence over others. More specifically, visual attention can be thought,

> "to be those processes that enable an observer to recruit resources for processing selected aspects of the retinal image more fully than non-selected aspects." Palmer (1999), p.532.

The difficulty we experience when we try to give our undivided attention to any one thing for a protracted period is something that most people are intuitively familiar with. Indeed, our *attentional capacity*, (the amount of perceptual resources available for a given task) can vary due to a number of factors, including alertness and motivation (Kahneman, 1973). Given this, it is likely that the Fear Relevant (FR) stimuli used in subsequent experiments may alter the motivational aspects of participants' capacity. In particular, the ability to concentrate and pay close attention to stimuli when other demands compete for resources due to heightened affective states will be addressed in detail in subsequent chapters. For the present, the issue of what attention is will be addressed with respect to normal populations.

Somewhat paradoxically however, the importance of attention can be seen at its most extreme in the disorders of attention experienced by patients with brain damage. One example is unilateral visual neglect, in which one side of visual space is ignored. While neuropsychology and lesion experiments have shown us *what* brain systems are implicated in attention demanding tasks, cognitive models of attention have attempted to elucidate *how* these systems operate in healthy subjects. Indeed, the invaluable contribution of cognitive neuroscience to the study of attention will be considered throughout this thesis. However, the primary methodology reviewed over the current chapter is of cognitive models of attention, beginning with a brief history of their development.

1.3 The Cognitive Revolution and Early Vs Late Selection Models

During the late 19th Century, perceptual processes were by and large studied using analytical introspection (Bruce, Green & Georgeson, 1996). However, by the turn of the 20th Century, behaviourism was the dominant movement, led largely by John Watson, who advocated an end to mentalist notions in favour of more 'concrete' discriminatory responses (Watson, 1913; 1924). Cognitive processes such as memory and attention were considered unsatisfactory topics for study, as their mechanisms, although indirectly measurable (e.g., via response times), could not be directly observed. Watson and the behaviourists purported that, as one cannot know how animals (or indeed other people) experience the world, inferences based upon observation were the only way of obtaining valid scientific data. Modern day psychology continues to be heavily influenced by behaviourism, although participants' verbal responses in perceptual tasks are usually considered legitimate data in the field of perception (Bruce, Green & Georgeson, 1996). By the post war era of the 1950's, interest in the field of attention and cognitive psychology had snowballed, partly as a result of Broadbent's (1958) Filter Theory. During this time, research was primarily concerned with auditory attention and the so-called 'Cocktail Party' phenomenon, described by Cherry (1953). Of fundamental concern to these researchers was the question of how auditory attention enabled individuals to attend to the relevant and ignore the irrelevant in a crowded and noisy environment such as a cocktail party. These researchers also worked to determine how much processing the unattended information received and at what stage filtering or gating occurred. This latter point was the forerunner to the early versus late selection debate, a dichotomy that continues to the present day and is discussed in more detail shortly.

Central to the study of auditory attention was the dichotic listening paradigm in which participants (through a set of headphones) 'shadowed' a message presented to one ear whilst ignoring a simultaneous message presented to the other ear. The task of the participants was to repeat aloud the relevant passage, with data showing little, if any, recall of the unattended channel when questioned later. Broadbent's contribution (initially designed to test the 'real life' scenario of attentional load on pilots) was seminal in that he attempted to integrate the data into an encompassing theory of attention and selection, which adopted a computational approach. In particular, Broadbent differentiated between parallel "preattentive" processing, in which basic physical properties were encoded, and serial "attentive" processing in which more elaborate abstract properties were encoded (Driver, 2001). This distinction, also made by Neisser (1967) was to be heavily influential in subsequent models such as Feature Integration Theory (Treisman & Gelade, 1980), reviewed later in this chapter. Indeed, Feature Integration Theory can be regarded as a sophisticated version of early selection theory (Pashler, 1998).

Whilst Broadbent's theory argued for early selection in perceptual processing (i.e. it was *an early selection theory),* whereby stimuli were selected prior to stimulus identification, *late* *selection* theorists (e.g. Deutsch & Deutsch, 1963; Duncan, 1980; Norman, 1968) argued that selection occurred much later on, following stimulus identification, and justified their stance by providing data to support deep processing for unattended information. Early selection is based upon the supposition that the processing system is of limited capacity and that stimulus identification is performed serially. In addition, low-level perceptual features such as colour and location, which are presumably detected early in visual processing, are posited to be the basis for selection. Late selection on the other hand, assumes that high-level features such as stimulus identity are the basis for selection and that all the stimuli in a scene are processed in parallel and without capacity limitations (Mozer, 1991). Somewhere in the middle ground between these two extremes is Treisman's attenuation theory (1960, 1969), which proposed that whilst deep processing of unattended stimuli could and did sometimes occur, such processing was the exception rather than the rule. Rather than being filtered, such stimuli were attenuated, usually to the point where abstract properties were rarely processed at a deep level. However, if particularly salient stimuli (such as a person's name), were presented to the unattended ear in a dichotic listening paradigm, then the deep level of processing could be accounted for by the fact that we are primed to detect personally significant information which requires less attention than other non significant words (Driver, 2001). This principle when applied to the paradigm of visual attention is particularly relevant for phylogenetic (evolutionarily relevant) stimuli and will be discussed in more depth in subsequent chapters.

Leaving aside the early experiments of selective attention conducted across the auditory modality, by the 1960's interest had shifted more towards the field of visual attention. Sperling's classic iconic memory experiments (Sperling, 1960), in which participants were found to have brief visual memories for cued rows of letters, suggested a form of limited capacity. Once again, the early vs. late selection debate continued with a roughly 50:50 split of evidence supporting each standpoint. For a detailed discussion of these findings see Driver (2001), or Mozer (1991) for a detailed analysis of computational approaches from a more hybrid perspective.

1.4 Cueing Experiments and the Spotlight Metaphor

Cueing experiments came to prominence in the early eighties and provided a simple yet elegant technique with which to measure visual attention. The basic tenant of the paradigm was that a cue (generally a brief illumination of one location point) could be used to direct attention to the forthcoming target. When a cue was presented at the same location as a target (a valid cue), participants were typically faster to detect the target than they were if no cue was present. In other words, the cue produced a *benefit*, or facilitation of target detection. However, if the cue appeared at a different location to the target (an invalid cue), then target detection was slowed (e.g. Posner, 1980) a phenomenon referred to as the *cost*. Quite how this facilitation works is disputed as it raises two issues. Firstly, is the facilitation a result of improved information coming from the cued location, or is it due to inhibition from other locations? Secondly, does facilitation occur only at the cued location, or does it spread across a gradient that includes regions surrounding the cued location?

Stemming from this work, one of the most popular metaphors for visual attention was that attention could be likened to a spotlight, so that stimuli that fell within its beam were given heightened processing, whilst stimuli that fell outside of the attentional spotlight were less likely to receive attention. As cueing experiments illustrated that attention could be guided to a specific location, this generated the question of how the movement of attention was guided. An experiment by Downing & Pinker (1985) manipulated the presentation of cues to different parts of the retina and found that RTs for cues close to fixation in valid trials were rapid, whereas invalid trials demonstrated substantial costs the further away the cue became from the target. Less substantial costs and benefits were found for peripheral locations, suggesting that the spotlight of attention can be more sharply focused at foveal locations than peripheral locations. Downing & Pinker suggest that this makes intuitive sense, due to the high concentration of neurones devoted to this area.

Further development of the spotlight theory revolved around a number of questions. One was whether it could be considered as a variable spotlight or a zoom lens. A summary of selectivity in visual displays was provided by Broadbent (1982), in which selectivity is portrayed as a searchlight with the option of altered focusing. While many authors assume that attention can be focused on only one spatial location (e.g. Eriksen & Yeh, 1985; Posner, Snyder & Davidson, 1980) others believe that under certain conditions it may be divided to 2 or more points in the visual field (e.g. Egly & Homa, 1984). The spotlight of attention may move across the field in one of two ways; either at a constant speed (e.g. Tsal, 1983) or in a constant time similar to ocular saccades (e.g. Remington & Pierce, 1984). According to Rizzolatti et al (1987) participants are faster to respond to valid cues as a result of ocular motor programming. If a cue is valid, then a response is immediately emitted, irrespective of ocular movements. However, in invalid cues, the distance feature must be modified and this process takes time. In unclear circumstances the beam is kept wide, whilst in clearer situations the beam is sharpened. For a more comprehensive review of the spotlight metaphor see Cave & Bichot (1999).

1.5 Is Attention Based on Objects or Locations?

Equally contentious in the attentional literature is whether visual attention is space or object based. According to Duncan (1984), theories of visual attention fall into three broad classes, object based, space based and discrimination based. It is with the former two categories that this brief review will concentrate. Whilst the spotlight metaphor implies that attention is directed to spatial locations (Egly, Driver & Rafal, 1994), Duncan (1984) argued that attention was directed to objects rather than their locations, primarily as a result of grouping principles. Evidence for this will be discussed shortly.

Much of the evidence for space-based attention has come from studies in which a test stimulus is preceded by a cue or some other stimulus that brings attention to the test stimulus location. (See Cave & Bichot (1999) for a review). The spotlight (Posner, 1980) and zoom lens (Eriksen & Yeh, 1985) metaphors used to account for their facilitated performance for valid cues. For example, the costs and benefits observed by central cues (arrows that point to the appropriate side of space or the word “left” or “right” indicating the likely position of the target) in Posner's (1980) experiment suggest that participants attended the location rather than an object, as the field was entirely empty. Stronger evidence for space-based attention has come from studies of neurological patients suffering from unilateral neglect, in which information on the side of space contralateral to their lesions is ignored. The impairment here is attentional rather than visual, as afferent pathways for the ignored information have been found to be intact (Egly et al, 1994). Specifically, participants were exceptionally slow in their RTs for invalid contralesional targets, interpreted as an inability to disengage attention.

An important experiment to measure whether attention takes the form of a spotlight was devised by Duncan (1984). Subjects were presented with small (<1°), brief foveal displays of two objects superimposed over one another, and required to judge different properties of these two targets. One was an outline box (rectangular in shape, but with a gap so that the shape was not fully completed), and the other was a line. Each object had two relevant properties. In the case of the box, these were size (small or large) and the position of the gap (facing left or right). The line comprised either dots or dashes and was tilted clockwise or anticlockwise. Displays were presented briefly, followed by a mask, and participants were required to judge either one or two properties. Performance was compared for discriminations on the same object and on different objects. For example, size and position were properties belonging to the box, whilst line composition and orientation were properties of the line. It was found that when two properties belonged to the same object (e.g. small and facing left), they were judged as effectively as one. However, when the two properties belonged to different objects (e.g. small and clockwise), performance was significantly worse. Neither the similarity, or the difficulty of the required discriminations, or the spatial distribution of information could account for the results, nor could it be explained by an attentional spotlight (assuming that the spotlight was constrained to have a roughly circular shape and a size of at least 1°. Thus it was proposed that parallel, preattentive processes segmented the field into separate objects. This was then followed by focused attention, which dealt with one object at a time.

In order to resolve the space based versus object-based dichotomy, Egly et al (1994) devised a paradigm in which they tested both space and object based attention simultaneously. Their task examined how cueing one part of an object affects the processing of other parts of the same object and equidistant other objects. Two experiments were conducted, using healthy controls (Experiment 1) and participants with parietal lesions (Experiment 2). In the interest of brevity, only data from the normal subjects are considered here.

Using a modified cueing paradigm, participants were presented with 2 rectangular boxes to the left and right of or above and below fixation. Cues were presented in one of the four corners of the rectangles prior to the display of the target stimuli, which occurred in either valid or invalid locations. The critical manipulation occurred in invalid trials (i.e. the target was not at the cued location). In these trials the target could appear on the same object as the cue but at the opposite end, or at one end of the uncued other object. Both these locations were the same distance from the cued location. The results showed that both space-based and object-based attention could occur within the same task. RTs to invalid cues within an object were slowed compared to valid cues, thus supporting space-based theories of attention. However, when the invalid cue concerned an equidistant location in the other rectangle, detection was significantly slower than for invalid cues within an object, thus demonstrating object-based allocation of attention. As these data provide evidence to support both forms of attention, the dichotomy between space-based and object based attention remains unresolved. Again, brevity dictates that a fuller analysis of the many experiments in this area is it beyond the scope of this review. However, Egly, Driver & Rafal (1994) have suggested that the form of attention taken (i.e. space or object-based) may be dependent upon the level of visual representation that the task demands. For example, if the task requires discriminating between different locations, then location-based attention is favourable. In contrast, if the task requires the comparison of object shapes, then object-based attention is preferable.

1.6 Traditional Visual Attention Paradigms

In the real world, humans frequently search for targets in a crowded visual array. Laboratory paradigms for testing visual attention often utilise similar visual search tasks, whereby participants are required to find a target hidden amongst a field of distractor items. Sometimes the experimenter will require that this task be accomplished in the absence of saccades (rapid, ballistic eye movements), which allow the object of interest to be foveated for greater visual acuity. However, eye movements are uncontrolled in the majority of visual search tasks (Wolfe, 1998a). Although we usually move our eyes to an object or spatial location in order to attend to it, attention and fixation are not necessarily coincidental (Palmer, 1999; Wolfe, 1998a). For example, if one attends to a fixation point it is perfectly possible to detect stimuli in the periphery without shifting fixation. In a heavily cited example used in textbooks (e.g. Palmer, 1999; Wolfe, 1998a), an asterisk is printed and the reader is invited to read nearby words without averting their eyes from the marker. Other examples in everyday life can be seen in the instance of words that carry personal meaning; for example one's name or hometown seems able to somehow "grab" attention away from what is being read and fixated on at the time (similar to the notion proposed by Moray, 1959). Shifting attention without eye movements is referred to as *covert orienting of attention.* More often however, people will generally move their eyes to the object to be fixated on, in order that it may be foveated for greater visual acuity. The deployment of visual attention in this manner is known as *overt orienting of attention*.

When one considers that there is only so much attention to go around, focusing our attention on any one stimulus carries a cost. This cost is often implicitly stated; for example, the teacher asking the class for their full attention. Continuing this train of thought, over the course of a visual search task, certain regions of the visual field receive more thorough processing than other regions. For example, when a number of objects appear simultaneously competing for our attention, information from each object can interfere with the processing of other objects. Our ability to engage *spatial attention*, i.e. selecting visual information by location (Kim & Cave, 1995) sets a limit on the amount of this interference by favouring the location of certain objects over others.

As previously discussed, in a standard visual search task, participants' search for a target hidden in a field of distractor items. This paradigm has been the bedrock of visual attention studies for the past 20 years (Wolfe, 1998b) as it provides a vast amount of useable data. The number of items in a display, the *set size*, typically varies from trial to trial, and the target may be either present (positive trials) or absent (negative trials). Generally, the participant will be asked to press one button if he or she has detected a target and another if they cannot find a target. Data are analysed according to accuracy and response times (RT). Hence, set size, salience and duration of stimuli are the independent variables and error rates and response times the dependent variables. Changes in the dependent variable as a function of set size constitute the preferred measure of search performance (Wolfe, 1998b). In studies in which

RT is the primary interest of the experimenter, a graph of RT against set size may be plotted and the underlying search mechanism inferred from the resulting slope. Accuracy methods are particularly important when display duration is very rapid in order to prevent eye movements (Wolfe, 1998a).

With respect to the interpretation of RT search slopes, the general inference is that shallow, fairly flat slopes indicate a parallel type of search, whilst steep slopes arise from serial search. The logic behind serial and parallel processing is as follows. If a target is presented in a variety of set sizes and the slope x set size function is almost flat, then the time taken to find the target is relatively independent of the number of distractors. It is therefore assumed that the underlying search mechanism is a *parallel* search, by which all the targets can be processed simultaneously. However, if there is a substantial slope (in other words if the RTs increase with corresponding increases in set size), then this is assumed to reflect a *serial* search mechanism (i.e. items are processed one at a time).

In "serial" type searches on target present trials, it is assumed that participants start their search at the first item (wherever that may be) and work their way around the display until the target is selected. The target could be the first item they fixate. Equally it could be the last. However, on average it will be the middle item in the display if a number of trials are averaged together. In target absent trials however, the participant must search each item in the display before ruling out the presence of a target, thus yielding a 2:1 slope for negative to positive trials. Searches of this type are generally referred to as *serial self- terminating searches.* The distinction between serial and parallel search is not always useful or indeed accurate (Wolfe, 1998a) even though it continues to be used by many modern researchers. This potential inaccuracy will be addressed later. Many of these concepts are fundamental to Feature Integration Theory (Treisman & Gelade, 1980) reviewed below.

1.7 Feature Integration Theory. Visual Search may be Either Serial or Parallel

In many respects Feature Integration Theory (FIT) addressed issues raised at the beginning of the last century. Titchener (1908) for example noted that attention is captured by rare or novel stimuli for basic features. By the 1960's, basic features were defined as properties that could support parallel or preattentive processing (i.e. without focused attention) (Wolfe, 1998a). This distinction was fundamental to Feature Integration Theory. The basic tenant proposed by FIT is that attention is directed serially to each stimulus in a display whenever conjunctions of more than one feature are needed for its identification. But what exactly is the relationship between features and search?

A common assumption of visual search is that any basic feature that differs from other nearby stimuli will capture attention automatically (e.g. Joseph, Chun & Nakayama, 1997). If a participant's RTs do not increase with corresponding increases in set sizes, then the search is assumed to be *parallel* (i.e. all items are processed simultaneously). There is some dispute as to what constitutes a basic feature in visual search. However, certain stimuli are almost universally recognised as basic features. Properties in this category include those that fall along the dimensions of colour, orientation, curvature, size, motion and depth (Wolfe, 1998a). Features seem to be basic properties that can be identified by simple computational mechanisms, and generally correspond to the properties that trigger responses in the early part of the visual system, e.g. orientation in area V1 and motion in V5 (Luck, 1998). Although certain simple stimuli are almost undisputedly regarded as being basic features, more complex properties have also been proposed as basic features. For example, Enns & Rensink, (1990) have suggested that three-dimensionality is identified preattentively, whilst more controversially, Kristjánsson & Tse (2001) propose that curvature discontinuities constitute basic features.

Returning to FIT, a feature that differs from the features surrounding it (e.g., a green bar in a field of red distractor bars; see Figure 1a) is said to "pop out" from its surroundings (Treisman & Gelade, 1980). As the number of distractors in the field increases, there is no additional cost. However, evidence suggests that if an additional attentionally demanding task is imposed concurrently, detection of simple features is severely impaired (Joseph, Chun & Nakayama, 1997). Conjunctions are composed of a combination of features. In this respect they can be considered to be an 'and'. For example, in a search for a red vertical bar in a field of green (vertical and horizontal) and red (horizontal) distractors (see Figure 1b), search is inefficient as the target is defined by a conjunction of properties (the colour red *and* the orientation vertical).

Since the target shares one of the relevant properties (orientation and colour) with each of the distractor items, there is a substantial cost in RTs to find the target with subsequent increases in set size (Treisman & Gelade, 1980; Wolfe, 1998a; Wolfe, Cave & Franzel, 1989). Searches of this nature are assumed to reflect a serial search, in which each item in the array is searched until the target is found, in contrast to the parallel search mechanism assumed to underlie feature searches. The distinction between what constitutes a serial and a parallel search, and the problems with describing searches in this manner will be discussed in more detail later.

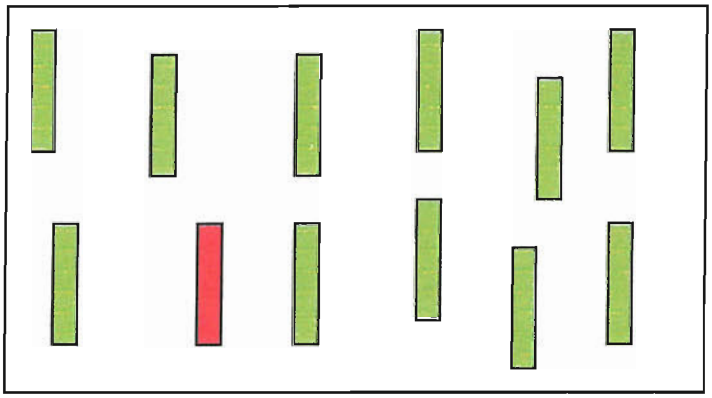

Figure 1.1 (a) An example of two basic search displays. The red target in the left display is easily detected, irrespective of the number of distractors present as it consists of a single basic feature (colour), which distinguishes it from distractors. *Efficient* searches of this type are often described as supporting "parallel" search as response times to detect the target are unaffected by the set size.

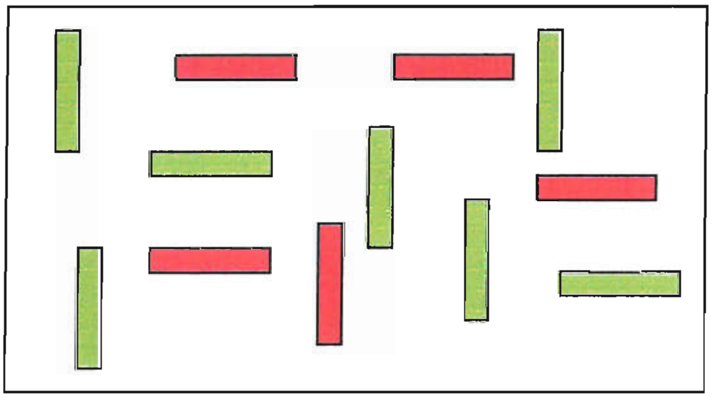

(b) Searches of this type (for the same red vertical bar) are *inefficient* as the target object is composed of a conjunction of features shared between the target and distractor objects (colour and orientation). RT's to detect the target increase linearly with corresponding increases in set size. Subsequently, they are said to support *serial* search.

One of the central tenants of FIT is that features are registered early, automatically, and in parallel across the visual field (within limits of acuity and interference etc). At this stage, there are separate mechanisms for each feature dimension, e.g., colour, orientation, spatial frequency, brightness and direction of movement. This early parallel process of feature registration mediates texture segregation and figure ground grouping. However, objects are identified separately and at a later stage. Focused attention is required for this later processing, and thus it is limited to one region of the visual field. In order to recombine these features, processing must be serial and this requires focused attention. Features within the fixation of attention are combined to form a unitary object, with focal attention acting as the "glue". Once registered, they are perceived and stored as unitary objects, although interference and memory decay may result in stimuli disintegrating and "floating free", or recombining in false guises. For example, a 1 and a 3 may combine to form a B, a phenomena termed an *illusory conjunction* (Treisman & Schmidt, 1982). Attention is necessary for the correct perception of conjunctions. However, participants can use their knowledge of typical co-occurrences of features to direct feature conjunction top down processing when attention cannot be used to ensure the correct combinations

Whilst FIT provided an excellent framework and explanation for serial and parallel processing, it was incapable of providing answers to certain questions. For example, why do some participants demonstrate almost flat functions for conjunction searches indicative of what we would traditionally term parallel processing? Over the past 20 years newer models of visual search have been proposed to account for some of the failings of FIT, including revisions by Treisman herself (e.g. Treisman, 1993), Guided Search (Wolfe, Cave & Franzel, 1989), Guided Search 2.0 (Wolfe, 1994) and FeatureGate (Cave, 1999). These newer models offer plausible accounts, some of which roughly correspond with physiological accounts of vision. These will be briefly reviewed later. For the present however, our focus shall turn to Top Down control and Bottom Up processing.

1.8 Top Down Versus Bottom Up Processing

Influential in this approach is the distinction between exogenous (or reflexive) and endogenous (voluntary) attentional control, proposed by Posner (1980) and Jonides (1981). Although both authors use different terms with slight differences in connotation, the basic premise is the same, i.e. automatic versus voluntary control of attention. In both examples, exogenous, or as it is interchangeably called, automatic attentional control, implies that attention is outside of the individual's control, and is governed by the stimuli impinging upon an individual's senses (Pashler, Johnston & Ruthruff, 2001). As a result, exogenous attention is often thought of as stimulus driven, or *bottom up* control. As discussed previously, certain features are thought to elicit this kind of automatic attention (e.g. colour, motion, etc.). In contrast, endogenous (or voluntary) control is thought to reflect higher level processing. Put another way, in the example of a visual search for a red O hidden amongst green O's and red X's, we *choose* to attend to the stimulus. Whilst a cognitive decision is hypothesised, no particular stimulus event is required (Pashler, Johnston & Ruthruff, 2001). Both approaches are reminiscent of Titchener's (1908) approach to attention discussed briefly at the beginning of the chapter. Universally agreed stimulus properties that automatically attract attention have already been covered. Other properties, such as transients (stimuli which change over time) have also been shown to "grab" attention, particularly transients with an abrupt onset (Yantis & Jonides, 1996). A suggestion of Yantis & Hillstrom (1994) is that from an evolutionary perspective, it is ecologically useful to detect new stimuli in a field in order to elude predators or to detect potential prey. This idea of preferential bias towards threat is crucial to the work of Öhman (e.g. 1992, 2001) and indeed this thesis, and will be discussed in greater depth over the following chapters. Presently, there is little consensus over precisely which stimuli are capable of exogenously attracting attention (with the exception of the stimuli already mentioned). Some researchers believe that all stimuli have at least some tendency to attract attention and that it is only their salience (or strength) that differs (e.g. Pashler et al, 2001). For the present it is perhaps sufficient to say that reflexive (exogenous) stimuli are free from the constraints of top down processing and instead, rely on bottom up processing.

In many instances the results of preattentive processing (processing carried out before the selection of a particular location) may determine where attention should be deployed. This idea is central to several contemporary theories of visual attention e.g. Guided Search (Wolfe, Cave & Franzel, 1989), Guided Search 2.0 (Wolfe, 1994) and FeatureGate (Cave, 1999), the latter of which will be subsequently reviewed. As previously discussed, there are two ways in which attention may be guided; top down (user driven) and bottom up, (stimulus driven). If bottom up processing directs our attention, then it is logical that we need also to be able to direct preattentive processes to find locations or objects that have properties that make them important for the current task, but that do not have a unique feature. Put another way, we need to be able to control preattentive processes in a top down manner. Evidence for this control comes from colour search tasks with heterogeneous distractors, in which it is possible to conduct efficient search for specified colour targets, thus resulting in parallel type search slopes (Wolfe, 1998a). This type of search cannot rely on bottom-up guidance, as bottom-up processing only locates unique features. As a heterogeneous display does not contain unique features, it must utilise top-down processing.

The division between top down and bottom up processing is addressed in FeatureGate (Cave, 1999). One of the main advantages of this particular model is that it addresses issues such as how efficient conjunction search may be possible. At the heart of FIT was the premise that conjunction search must be serial, as preattentive processing could not detect conjunctions. However, in models such as Guided Search and FeatureGate attention is *guided* to likely targets in the manner described below.

1.9 The FeatureGate Model of Visual Attention

FeatureGate is an attempt to explain visual search and is essentially an upgrade of Guided Search (Wolfe, Cave & Franzel, 1989), which has itself been superseded by Guided Search 2.0 (Wolfe, 1994). As such, it is not the intention of this paper to review numerous models. Each model has a number of strengths, with few outstanding problems. However, brevity necessitates a cursory review of just one model. For the purpose of this paper the focus centres on FeatureGate, as this is arguably the most up to date and plausible cognitive model of visual attention available. Furthermore, the model roughly corresponds with physiological accounts of visual attention and like certain other models (e.g. Humphreys & Müller, 1993) can be successfully implemented as a neural network. FeatureGate consists of a hierarchy of spatial maps with attentional "gates" controlling the flow of information through this hierarchy driven by both top down and bottom up processes.

FeatureGate assumes that if a person is presented with a number of objects simultaneously, information from each object can interfere with processing and identification of the other objects. In order to prevent interference, spatial attention is required. FeatureGate begins with an input being encoded as a pattern of activation across a 2-D arrangement of neurones. Each location has a set of units, each specific to a simple visual feature (e.g. colour, orientation, size etc; see Figure 1.2). Connections from all these different units converge at a single set of output units, which represent an object's visual features but not its location. The regulatory mechanism controlling the gates relies on two sub-systems, the top-down and bottom-up systems, to select the location most likely to contain important information.

The top-down system operates when a person knows some or all of the visual features present in a target (e.g. a red vertical bar). It operates by closing gates at locations with non-target features in order to inhibit processing of information at those locations. The more a feature differs from the target, the more the gate will be closed and the stronger the inhibition to that location. Inhibition can be turned either up or down depending upon the needs of the task.

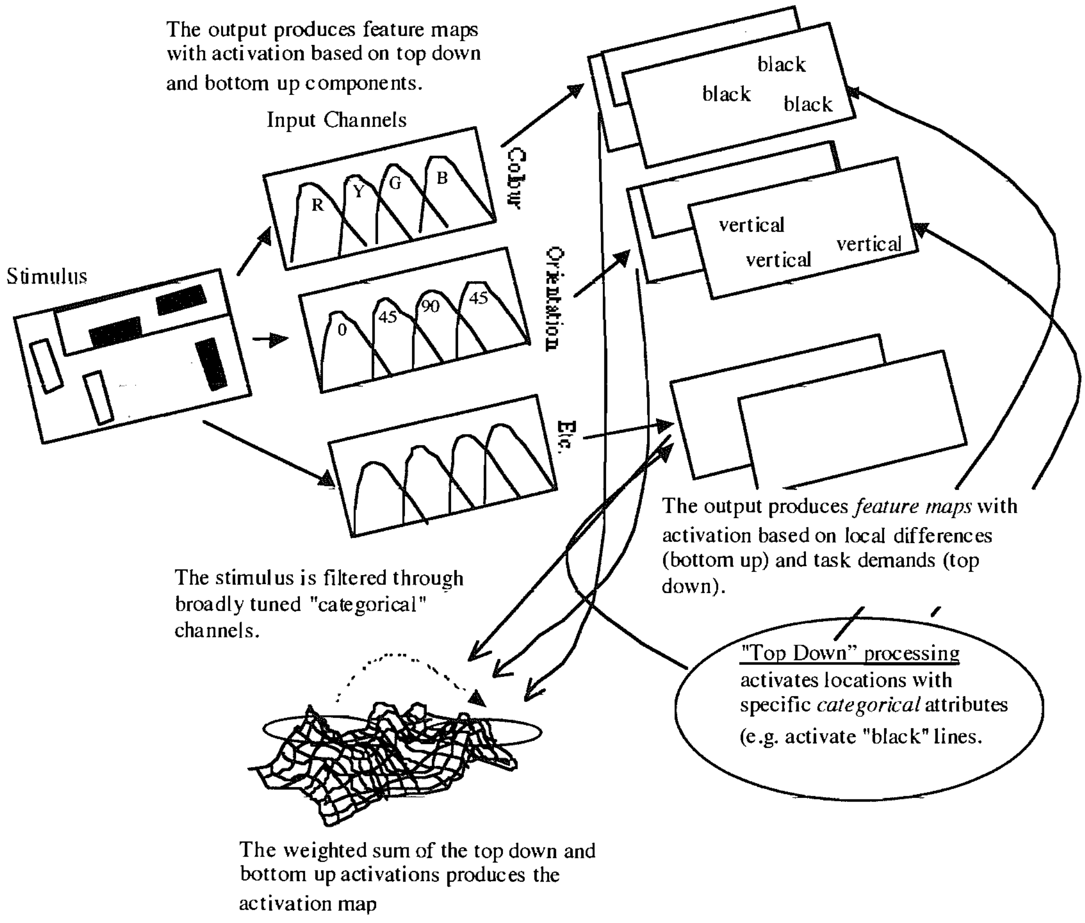
The output produces feature maps with activation based on top down and bottom up components.

The stimulus is filtered through broadly tuned "categorical" channels.

The output produces *feature maps* with activation based on local differences (bottom up) and task demands (top down).

"Top Down" processing activates locations with specific *categorical* attributes (e.g. activate "black" lines.

The weighted sum of the top down and bottom up activations produces the activation map

Figure 1.2. Basic model of FeatureGate (shown without the hierarchical structure) taken from Guided Search 2.0 (Wolfe, 1994, p205).

The Bottom-Up System, on the other hand, selects targets that differ significantly from distractors in any of a number of ways (e.g. orientation, colour etc). It compares the features at a location with neighbouring locations and opens gates for targets with features that differ from their surrounds (using a 3x3 matrix in the implementation). It ensures that feature singletons (e.g. a green object in a field of red objects) will have an open gate. The more the features differ from surrounding features, the stronger the influence will be to keep the gate open wide.

The Top-Down and Bottom-Up systems initially work independently, each calculating an activation for each location. Each dimension is summed to form an overall top-down and bottom-up activation and then the two systems are summed to produce an overall activation (as in Guided Search). The activations are compared across different locations and their gates partially closed unless that location has the highest activation. Hence only the features at the location with the highest activation are made available for higher level processing. (Note the hierarchical organisation in FeatureGate roughly corresponds to the organisation of visual cortex). In order to work efficiently, local operations are used in order to avoid long connections (as in biological and computational models). The visual field is split up into "neighbourhoods", with the 'winning' location from each region combined with the winning locations from other regions to form a new spatially organised representation at the next level. These winning locations then compete with each other and so on, with fewer and fewer units at each level. Thus, unlike Guided Search, FeatureGate consists of a hierarchy of spatial maps. At the upper levels of the hierarchy each location represents a single location selected from a broad region of the visual field, with only a single location at the top level. This is the 'winner'.

In many respects FeatureGate is very similar to Guided Search. Like Guided Search, it is imperfect, where very occasionally a target will be inhibited and a distractor selected instead. This rarely happens when the target differs from the distractors by a single salient feature (e.g. vertical amongst horizontal or red among green). In feature searches the bottom-up component activates a large dose of attention to the target but not the distractors. However, in conjunction searches as in Guided Search, a certain amount of random noise may be added. Under some conditions FeatureGate makes errors without the noise. The bottom-up system activates locations with differing features. However, distractors of one type can be highly activated if near distractors of another type (e.g. a red vertical distractor near to green horizontal distractors can be selected instead of a red horizontal target). Once the distractor has been selected and found to be a non-target the location can be inhibited by Inhibition Of Return (Posner & Cohen, 1984), reducing activation at the winning location. A new cycle will then be activated and a serial search performed. In general, the more distractors present in a conjunction search, the more distractors will be selected before the target, and hence the longer search times as in previous studies. FeatureGate explains the wide discrepancy in conjunction search times between subjects by assuming that subjects with steep slopes have stronger bottom-up than top-down systems, whilst subjects with fast RTs have stronger top-down systems. Unlike searches for feature targets, if a target is composed of a conjunction of features, then a strong bottom-up system is not beneficial, as the target does not differ from the distractors by a single basic feature. In this scenario, strong top-down systems are much more useful, as the search for the target is based upon the target location having *all* the relevant features.

1.10 The Problems of Adopting a Serial/Parallel Distinction

Returning to the point made previously regarding the problem of a serial or parallel distinction, it should now be obvious that visual attention cannot be simply placed into either category and left at that. For example, many conjunction searches yield slopes that would be considered indicative of parallel processing as their search slopes are more efficient than traditional serial searches (Wolfe, 1998b). While the serial/parallel dichotomy is indeed useful, it is too simplistic. Furthermore, inferring mechanisms from search slopes this way can be wrong (Wolfe, 1998a). For a more detailed discussion of this topic see Wolfe (1998a). For the present however, the terms efficient and inefficient are used interchangeably with parallel and serial respectively. Where the underlying search processes in later experiments are proposed, these will be stated directly.

1.11 Summary

Over the course of this chapter, some of the key studies and issues pertinent to modern attention research have been discussed. Although the focus of this paper centres upon visual attention, a brief history of auditory attention has been included in order that some of the fundamental issues directly relevant to visual attention can be detailed in a logical and coherent manner. Whilst some of these points are fairly contentious, (e.g. early versus late selection, and space versus object based theories of attention), certain issues are more definite.

A consistent finding to have emerged from the literature is that certain properties are easier to detect than others when hidden in a crowded visual array. This notion has been incorporated into a number of theories (e.g. Cave, 1999; Treisman & Gelade, 1980; Wolfe, Cave & Franzel, 1989; Wolfe, 1994) with a vast body of data in support. While each theory differs from the others in a number of ways, a central tenant common to all is that search will be fast and efficient when the target differs from distractors by a simple feature that is identified early in visual processing. However, when the target is defined by a more complex configuration, search will be slow and less efficient, as detection occurs at a later stage in the visual system. These findings are discussed in relation to clinical models of attentional bias in the Chapter 3, where the rationale for the studies is outlined. Prior to this, Chapter 2 provides an overview of anxiety and some of the key theories and findings pertaining to it.

Chapter 2: Anxiety

"Emotion is that which leads to one's condition to become so transformed that his judgement is affected, and which is accompanied by pleasure and pain. Examples of emotions include anger, fear, pity and the like, as well as the opposites of these." – Aristotle (384-322 B.C.) "Rhetoric".

2.1 Introduction

Just as the previous chapter described some of the key landmarks in visual cognition, the present chapter outlines key research pertaining to anxiety and anxiety disorders. The chapter begins with a brief discussion of the mind as an information-processing device, before considering some key historical and philosophical issues relevant to emotion. After discussing the relationship between fear and anxiety, the aetiology and nosology of anxiety disorders are described, with particular reference to conditioning and anxiety. Finally, neuropsychological and cognitive theories of anxiety are summarised, (although a more comprehensive treatment of cognitive theories of anxiety can be found in Chapter 3) before the chapter summary in section 2.11.

2.2 The Mind as an Information-Processing Device

The initial approach in this paper is to treat cognition (primarily, visual attention) and emotion (specifically fear and anxiety) as separate but interacting entities. Indeed, the dichotomy between cognition and emotion is still very much evident in today's judiciary. For example, 'crimes of passion' generally receive more lenient sentences than premeditated acts, as the heightened affective states experienced by the perpetrator are considered to interfere with their ability to behave in a rational manner (Maguire, Morgan & Reiner, 1996). Up until the mid 1980's, emotion was considered an unsuitable topic for study within cognitive psychology, possibly because of cognitive psychology's heavy reliance of the computer metaphor. Put simply, the mind (in its broadest sense) was treated as an information-processing device, analogous to that of a computing system. Cognitive processes such as perception, attention and memory were studied analytically and systematically, with logical reasoning emphasised over the traditionally so-called "illogical" emotions (LeDoux, 1996).

Computers process information but do not "experience" the world or have what we would consider consciousness (or at least, not as yet). However, this is not to say that it is not possible to programme a computer to experience emotion, whether in a primitive form or in a more advanced "human like" experience. While the affective responses of the cyborgs featured in Stephen Spielberg's latest blockbuster AI may be far fetched at the present time, this does not preclude the notion of "feeling" computers. Part of the problem here may lie in the fact that both philosophers and psychologists have yet to reach a consensus of opinion on what constitutes consciousness. As such, it is hardly surprising that the holy grail of Artificial Intelligence (AI) research (namely, conscious emotionally capable machines) is still some way off. Part of the problem in viewing emotion as a purely human trait lies in the mystification of what are arguably purely cognitive (and biological) processes that are treated as irrational states rather than evaluative mechanisms. For example, fear serves a valuable function in protecting us from danger and aversive consequences. Therefore, although AI is some way off producing computers with levels of consciousness or affect on a par with humans, regarding the mind as an information processing device makes a great deal of sense, and a fully developed information processing account of the mind can be expected to include emotion as well as cognition.

An important point to note at this juncture is in the logic of brain science and its reliance upon rules. From our current understandings in the field of neuroscience, brain processes are very logical and rule driven. For example, the neural pathway in visual recognition involves stimulus input from the retina passing through the lateral geniculate nucleus of the thalamus to area V1 (Luck, 1998). From here, the output is transferred to areas V2, V4 and so on for higher processing. A series of feedback loops also exist to allow top down processing, to enable processes like attention to exert control. While it is highly improbable that exactly the same neurons are used to process information in all individuals, the mode by which information is governed by the neurons (e.g., the regulation of key neurotransmitters) seems to be fundamentally the same for all people for any given process. Thus, all brain systems follow logical rules and laws of physics and biochemistry across all known processes. While individuals differ, these differences are arguably due to differences in the connections between the neurons in the brain. Experience plays a part in shaping personality and one's outlook on the world, but it is experience that strengthens or weakens synaptic connections and hence behaviour. Thus, whilst each individual experiences emotion in a different way, to state categorically that emotion is not rule governed and open to scientific manipulation and measurement is not only simplistic, but it is also incorrect. If one accepts that emotion is governed by rules, then adopting a cognitive approach to the study of emotion is a logical approach.

As humans, our feelings and affective states often appear to impinge upon our reasoning and behaviour. We rightly or wrongly tend to view ourselves as more than information processing devices, and contrary to the views expressed in the former paragraph, often recognise when we are acting irrationally; for example, when experiencing feelings of jealousy. However, whilst jealousy is frequently viewed as a negative and irrational state, it is important to note that this 'irrational' feeling arguably serves an important adaptive function of alerting us of potential rivals that threaten our parental investment (Rolls, 1999).

If the more 'advanced' affective states experienced by humans are a fundamental tenant, which distinguish the human mind from AI attempts to model it, then our elevated emotionality comes at the expense of poorer processing efficiency in comparison to computer systems. This trade off in terms of lower task efficiency and increased affective responses serves a valuable function in alerting us to potential hedonic or threatening situations, and can be regarded as the 'cost' of consciousness. Normally there is very little evidence of this cost. However, in individuals with affective disorders (particularly anxiety disorders), this cost becomes more apparent, as there is a breakdown in the logic system, whereby the threshold for appraising threat becomes pathologically lowered (e.g. Mogg & Bradley, 1998). This is not to say that in individuals with anxiety disorders, emotion is not governed by rules. Rather, in populations with pathological anxiety, their rules are just less effective. Thus, where emotions can be considered to be logical, rule driven processes, emotional disorders can be viewed as the corollary of a breakdown in this system. This topic will be addressed more extensively later in the thesis. Specifically, Chapter 3 will examine the link between the two once disparate fields of cognition and emotion, and there the rationale behind the studies will be expanded. For the present however, the initial chapters aim to define and explain some of the key concepts, issues and methodological problems inherent in much of the current research in this area.

2.3 Emotion and Attention to Threat

One of the fundamental questions of this thesis is whether threat stimuli can be detected preattentively if there are no low-level perceptual features signalling threat. In other words, if target and distractor stimuli share different arrangements of the same basic features, based upon data from contemporary models of visual cognition (e.g. Cave, 1999), there should be no differences between participants' search times for neutral versus threat conditioned stimuli. However, theories developed within Clinical Psychology (e.g. Mogg & Bradley, 1998; Williams et al, 1988, 1997) suggest that there are attentional biases towards threat, raising the question of whether these biases are produced at a preattentive level. Furthermore, it has been suggested that the direction of these biases is mediated by individual differences in anxiety (Williams et al, 1988, 1997). This point will be expanded upon in greater detail later in the chapter. The initial studies (to be presented in later chapters) centre upon the effect of emotional stimuli on the cognitive system in student populations separated by differing levels of trait anxiety, while later studies examine the effect of specific phobia on the cognitive system. For practical and ethical reasons, over the course of the thesis all experiments focus upon student populations with sub-clinical levels of anxiety.

The study of anxiety is important for two reasons. Firstly, the prevalence of anxiety in the general population is reasonably high, with anxiety disorders being a common psychiatric illness causing considerable morbidity and social cost (Gratacòs et al, 2001). Secondly, and perhaps more importantly, a substantial number of recently published studies suggest differential cognitive processing in individuals with affective disorders. Schema (e.g., Beck, 1976) and network models (Bower, 1981), propose that anxiety and depression are associated with information processing biases, including attention and memory. Furthermore, these biases are congruent with the relevant emotion, so that in the case of anxiety, the associative link (i.e. hypervigilance towards threat) should be stronger for danger, whilst for depression, the bias should be associated with loss or failure. During the 1980's a large number of experiments were carried out to test these hypotheses. Subsequently, information-processing biases have reliably been found for memory in individuals with depressive disorder, whilst attentional biases are more prevalent in anxiety disorder (e.g. Eysenck, 1992; Mathews, 1990; Mogg & Bradley, 1998; Williams, Watts, MacLeod & Mathews, 1988,1997). The ensuing surge of interest in the field of cognition and emotion was unsurprising given that the rigorous methodology of cognitive science is ideally suited to the study of affective processing. If one accepts the logic of a cognitive approach to the study of emotion (or at least a narrow spectrum of emotion) in normal individuals, then this approach may provide valuable information about what happens when the emotional system malfunctions and breaks down. This approach will be discussed further in Chapter 3.

2.4 Emotion, Fear and Anxiety: Historical and Philosophical Issues

The focus of this chapter is to provide a backdrop to some of the key issues pertaining to both current and previous research in the field of emotion. The perspective taken is predominantly cognitive, although neurobiological evidence will also be considered to a lesser extent. The primary emotion under scrutiny in this investigation is fear and its relationship with anxiety. It is hoped in this chapter to address some of the key issues raised over the course of this research, although inevitably there still remain more questions than there are definitive answers.

The seemingly simple task of defining emotion with all its subtleties and nuances is surprisingly difficult. However, to a large extent we are all intuitively aware of what constitutes an emotion. Indeed, in the course of our everyday lives we frequently run through a whole gauntlet of emotions. Theories of emotion are often tied in with theories of mind (Power & Dalgleish, 1997) and as such, are a concept held dear by philosophers, both past and present. However, fundamental questions concerning the very nature of emotion often seem difficult to answer in a scientific manner. Basic questions, such as what distinguishes emotion from non-emotion? How many emotions are there? And what is the relationship between normal emotion and emotional disorders? (See Power & Dalgleish, 1997 for a review) are often problematic and difficult to put into a scientific framework. However, as this chapter has and will continue to argue, emotion can be studied empirically and cognitive psychology offers the best tool with which to dissect its intricacies.

Emotions are sometimes defined as states elicited by reward and punishment, whereby rewards are anything that an animal will work for, whilst punishments are anything that an animal (or human) will work to avoid (Rolls, 1999). It is worth noting however, that certain stimuli may be excluded from this list. For example, when positively rewarding stimuli are relevant to the drive state, e.g., food to a hungry animal, even though they may produce pleasure, they do not elicit affect (Rolls, 1999). That is to say, pleasure and emotion are not the same. A hungry animal receiving food will experience a number of positive states, predominately satiety, and the taste of certain foods may be more pleasurable than others. These states are not emotional states, however. It is not the intention of this paper to expound upon these points. Rather, the focus will be upon a more limited subset of emotion, i.e., fear, and more particularly, anxiety. For the present, the focus will be on behaviourist and recent cognitive theories of emotion, with a foray into recent findings in the field of neuroscience.

Just as Freud is regarded as the father of Psychoanalysis, John Watson can be considered the father of Behaviourism. In 1924 Watson put forward a theory of fear in which only 3 kinds of innate stimuli (loud noise, pain and sudden loss of support) were capable of causing fear without prior learning (Cohen, 1979). Anything else was said to come about as a result of learning, primarily through the process of Classical Conditioning (or as it is synonymously referred, Pavlovian Conditioning). However, while Watson's theory was elegantly simple, there were a number of holes in his argument, not the least of which being the fact that numerous other stimuli are innately capable of eliciting fear, (Gray, 1982).

Ethologists (zoologists who specialise in the study of behaviour) have provided many valuable insights into the nature-nurture debate in the acquisition of fear in animals. One important tenant to have sprung from this work is that of the maturation of fear. Evidence from human studies suggests that ophiciophobia (fear of snakes) and arachnophobia (fear of spiders) are innate. (Öhman, Esteves & Soares, 1995). Interestingly, these are among the most common phobias (Pinker, 1997) and form the basis for phylogenetic arguments of attentional bias (covered in subsequent chapters and the rationale for Experiment 3). Lifetime prevalence rates of phobia in general range from 10% to 11.3% (APA, DSM-IV). In an experiment to investigate the maturation of fear in humans, Jones & Jones (1928, cited in Gray, 1982) released a large harmless snake into a children's play enclosure. For the youngest children (0-2 years), no evidence of fear was shown. However, by the age of 3-4 years old, the children had started to show signs of apprehension, which increased with age into definite signs of fear. Whether this fear was an artefact of social learning by the children who may have heard or read 'horror' stories about snakes is impossible to determine. However, as a study by Hebb (1946, cited in Gray & McNaughton, 2000) revealed a similar process of fear acquisition in the chimpanzee, it appears likely that there is some form of "hard wiring" already in place, suggesting an evolutionary bias in the acquisition of certain fear relevant (FR) stimuli. Furthermore, the similarity between chimpanzees and humans extends to other fears, including the fear of strangers and the fear of dead or mutilated bodies.

Before continuing further it may be useful to define a few terms. It is not the intention in this thesis to provide a definitive definition of emotion per se. As LeDoux (1996) has pointed out,

> "Emotion is only a label, a convenient way of talking about aspects of the brain and mind...There is no such thing as the "emotion" faculty and there is no single system dedicated to this phantom function. If we are interested in understanding the various phenomena that we use the term "emotion" to refer to, we have to focus on specific classes of emotions. We shouldn't mix findings about different emotions all together independent of the emotion that they are findings about. Unfortunately, most work in psychology and brain science has done this." (p.16)

2.5 The Relationship Between Fear and Anxiety

In keeping with this premise, the specific emotion of fear and its relationship to anxiety (and visual attention) occupies centre stage in this investigation. Anxiety is a normal feeling or response to stress, and is something that most, if not all of us, have experienced at some stage. Although fear and anxiety are closely related, it is important that one realises that they are not simply referents of the same thing. The distinction between fear and anxiety is often described in terms of internal and external stimuli (LeDoux, 1996). Where anxiety is elicited by internal states, such as ruminating over current worries and concerns, fear is elicited from external sources; for example, coming across a dangerous animal or a potentially bruising encounter with an assailant. For others, (e.g. Öhman, 1992), the acts of escape and avoidance best differentiate between the two states. For example, a spider phobic will become fearful if the object of their fear (a spider) is put in a room in front them. If the door to that room were then shut and locked, making escape impossible, their fear would then be transformed into anxiety. Where anxiety can be regarded as the inability to escape a threatening situation, fear is the anticipation of pain (Gray, 1982). Again, fear and phobias are closely related, and it has been suggested that phobias are simply innate fears, which have not been unlearned (Marks & Nesse, 1994). Regardless of the subtle nuances between these definitions, a common theme running through each explanation is that fear is the precursor of anxiety.

Although the approach taken throughout this thesis is guided by methods from Cognitive Psychology, findings from other related disciplines will be integrated whenever appropriate. As the neurobiology of fear circuits in the brain are particularly important to the understanding of anxiety disorders, the following section briefly describes some important findings from neurophysiology experiments. LeDoux (1996) proposed that anxiety disorders arise when the fear system detaches itself from cortical controls that otherwise inhibit our behaviour. Numerous studies implicate the importance of the amygdala in fear and fear conditioning (e.g. Lang, Bradley & Cuthbert, 1998; LeDoux, 1996, 1998, 2000: Morris et al, 1996, 1999; Rosen & Schulkin, 1998) and the septo-hippocampul system (Gray, 1982; Gray & McNaughton, 2000, discussed later in the chapter). In addition, the neural pathways in specific anxiety disorders have been mapped with varying levels of success (e.g. Coplan & Lydiard, 1998; Heller et al, 1997; LeDoux, 1998; for a summary, see LeDoux, 1996). From these studies, most authors concur that pathological anxiety is an exaggerated fear state involving the amygdala and extended amygdala, and that normal fear and anxiety are emotional responses to danger or threat (Rosen & Schulkin, 1998).

Fear is a precursor of pathological anxiety, while pathological fear is simply an exaggerated emotional state (Rosen & Schulkin, 1998.). Interestingly, recent findings suggest a genetic component to certain forms of anxiety (panic and phobic disorders) due to an interstitial duplication of Chromosome 15 (Gratacòs et al, 2001). However, the main gene implicated in anxiety (the serotonin transporter gene, 5-HTT) is encoded on chromosone 17q12 (Lesch et al, 1996). The fundamental point highlighted by these findings is the heterogeneous nature of anxiety. Thus, whilst the fear system may be engaged in a similar manner and use essentially similar circuitry for different subgroups of anxiety, the genetic basis may be very different for each subtype. This point may be particularly salient in explaining the inconsistencies in hemispheric asymmetry in anxiety disorders, as different types of anxiety may account for the somewhat erratic data (Heller et al, 1997). An important differentiation between anxious apprehension (i.e. worry), typically defined in terms of cognitive processes, and anxious arousal (i.e. panic) characterised by somatic tension and arousal, has often been disregarded (Heller et al, 1997). Hence, unequivocal findings in the physiological measurement of anxiety, whether measured by rCBF, PET or EEG are unlikely to be found unless one carefully controls and compares studies by anxiety type.

Two important points to arise from these findings are the general consensus that the amygdala plays a pivotal role in the control of fear, and the heterogeneous nature of anxiety disorders. Although we may not have experienced anxiety at its extreme, or at least in the same prolonged manner as an individual with any of the pathological anxiety disorders described in table 2.1, anxiety is a feeling that all of us have at least rudimentary experience of (for example the anxiety we may experience in the run up to a presentation or job interview). As previously discussed, anxiety is closely related to fear. However, despite the fact that fear in its evolutionary context was generally related to physical threat (see experiment 3), this does not imply that anxiety is elicited by physical threat alone, as excessive cognitive rumination and worry is fundamental to many anxiety disorders (e.g. Generalised Anxiety Disorder).

2.6 From Normal to Clinical Anxiety

Many people are afraid of something, whether it is a 'real' physical threat (e.g. arachnophobia) or a more 'cognitive' concern (e.g. fear of failure). As such, fear to some extent is a fact of life. Fear responses, such as startle, heart rate and blood pressure changes, are evolved defence responses, elicited to facilitate escape and avoidance from such fear (Charney et al, 1996; Rosen & Schulkin, 1998). Furthermore, seemingly 'maladaptive' responses to fear, such as fainting at the sight of blood, highlight adaptive strategies. As fainting is usually attributable to a drop in blood pressure, such actions would lessen the likelihood of one's own blood loss (Pinker, 1997).

Whilst fear and anxiety motivate an individual to escape from negative states, fear is a major precursor for the development of pathological anxiety, a state defined as an exaggerated fear that manifests itself into the various symptoms of the anxiety disorders covered later in this chapter. The aetiology of pathological anxiety is hypothesised to stem from normal fear responses, which, in the presence of danger, lead to increased activation of the brain's fear circuits. In normal individuals, this activity subsides with the dissipation of the danger. However, in individuals with a sensitivity or predisposition to anxiety, fear circuits may become over-activated following a stressor, and through a combination of behavioural and biological processes, they may develop into hyper-excitable fear circuits, culminating in a switch from normal fear into pathological anxiety (Rosen & Schulkin, 1998).

Anxiety can appear at any age. For example, in their first year of life, babies exhibit strong separation anxiety and fear of strangers (Pinker, 1997). However, anxiety disorders most frequently manifest themselves in early adult life (LeDoux, 1996; Gray & McNaughton, 2000). Anxiety and depression are the most common affective disorders, with anxiety disorders occurring more than twice as frequently as depression amongst older adults, the highest rates being for phobias and generalised anxiety disorder (Stanley & Beck, 2000). But how do they begin, and is the transition from normal to pathological anxiety a gradual process or of sudden onset? Certainly, trait anxiety has most frequently been cited as one of the strongest personality dimensions predicting vulnerability to anxiety disorders (Eysenck, 1992). However, as previously stated, anxiety is not a unitary concept. Therefore the question posed as to the time course involved in the aetiology of anxiety is perhaps a little inappropriate and should perhaps be made to address each subtype individually. Regardless of this, genetic factors seem likely to play a causal role in individuals' susceptibility to anxiety (Gratacòs et al, 2001; Lesch et al, 1996; Rosen & Schulkin, 1998). More prominent in the literature, however, is the idea of fear conditioning being responsible for the genesis of pathological anxiety. Most famously cited is the case of "Little Albert", an 11-month-old boy who was conditioned to develop a fear of rats due to classical conditioning. Conditioning theories of anxiety will be covered in detail later in this chapter.

2.7 Nosology of Anxiety and DSM (IV) Criteria

Under the general umbrella of Anxiety Disorders the following disorders are classified and defined as shown in table 2.1 (APA, DSM-IV):

Disorder	Description of symptoms
Panic Attacks.	Discreet periods of sudden onset characterised by feelings of intense apprehension, terror and impending doom. Associated symptoms include shortness of breath, palpitations and fear of losing control.
Agoraphobia.	Refers to anxiety of, or avoidance of, places and situations where escape may be impeded or embarrassing, and help unavailable in the event of a panic attack.
Panic Disorder with and without Agoraphobia.	
Agoraphobia Without History of Panic Disorder.	
Specific Phobia.	Specific phobia is characterised by clinically significant anxiety in the event of exposure to the feared stimulus. The individual will often engage in avoidance behaviour of situations where an encounter with the feared stimulus may occur.
Social Phobia.	Is characterised in a similar manner, with the exception that the feared situations are social encounters or performance situations.
Obsessive-Compulsive Disorder.	Is characterised by obsessions that cause significant anxiety or stress, and/or compulsions (often ritualistic in nature), which serve to neutralise the anxiety.
Posttraumatic Stress Disorder.	Is characterised by the re-experiencing of an extremely traumatic event. It is often accompanied by somatic arousal and avoidance of associative stimuli.
Acute Stress Disorder.	Is similarly characterised, occurring immediately in the aftermath of a highly traumatic event.
Generalised Anxiety Disorder.	GAD is characterised by persistent and excessive anxiety and worry of at least 6 months duration.
Anxiety Disorder due to a General Medical Condition.	Is characterised by symptoms of anxiety, which are judged to be directly attributable to a general medical condition.
Substance-Induced Anxiety Disorder	Is judged in a similar manner and attributed directly to drug abuse, medication, or exposure to toxins.
Anxiety Disorder Not Otherwise Specified.	Is used to code disorders with prominent anxiety or phobic avoidance that do not fit or meet the criteria of the above categories.

Table 2.1. Nosology of Anxiety Disorders (taken from DSM-IV p 403)

2.8 Conditioning and Anxiety

One of the most prominent theories of the aetiology of anxiety is that of fear conditioning (e.g. Watson, 1924). A particularly interesting observation by Öhman (1986) is that as social phobias typically develop during puberty and young adulthood, there may be a critical period for their development (see also LeDoux, 1996; 1998). However, research into the onset of panic disorder and its relationship with prior traumatic events is equivocal, as conditioning theories are unable to adequately explain PD in the absence of prior traumatic events (Rosen & Schulkin, 1998). The idea of genetically inherent or *phylogenetic* fear is especially prevalent in the work of Arne Öhman and will be reviewed in more depth in Chapter 3. Equally important is the function of learning and classical conditioning in the acquisition of fear. Classical conditioning, particularly Pavlov's conditioning experiments with dogs, is possibly one of the most famous topics in popular psychology and paved the way for an explosion of interest in learning behaviour studies. Briefly, the process involved in classical conditioning involves the pairing of a neutral stimulus (the conditioned stimulus, CS) with a biologically relevant stimulus, the unconditioned stimulus (UCS). In the case of Pavlov's dogs, the CS was a tone and the UCS, food. Initially, the UCS elicits unconditioned responses (UCR). However, after repeated pairings, the CS eventually elicits a conditioned response (CR). In Pavlov's experiment, the CR was salivation at the sound of the buzzer. Put another way, the UCS is a natural trigger, whilst the CS is a learned trigger (LeDoux, 1996). Classical conditioning is dependent upon the relationship between two stimuli, so that a neutral stimulus eventually elicits an affective response, due to the strengthening of the memory representation between the CS and UCS (Hermans et al, 2002).

Numerous behaviours and indexes of fear can be classically conditioned in animals (and human beings, where ethically acceptable) in the laboratory. Of particular interest is the tendency for stimuli similar to the CS to elicit the CR, a phenomenon termed by Watson as *stimulus generalisation*. In the case of "Little Albert", Albert's fears extended to white rabbits and fluffy toys, in addition to the conditioned animal. If the similarity between a stimulus and the CS is very strong, the magnitude of the response is strong. Similarly, as the similarity between the stimulus and the CS decreases, the magnitude of the response decreases too (Gray, 1982). It is much easier to transfer fear to biological fear relevant (FR) stimuli than to other stimuli as, according to Öhman & Mineka (2001), we have an in-built predisposition to fear these stimuli, or as Valentine (1936, cited in Eysenck, 1992) memorably describes it, a "lurking fear in the background". Intuitively, this idea makes sense, particularly given what we know about the maturation of fear. From an evolutionary perspective, certain animals or stimuli possess an inherent danger to us. Given this, it seems logical that it should be easier to condition a fear of snakes or spiders than it should be to condition non-FR stimuli. In an experiment by Öhman (1992) using conditioned snake and spider pictures, extinction was slower for FR snake and spider pictures than for ontogenetic stimuli (e.g. pictures of guns and other weapons). Whilst the latter category clearly serve a dangerous function, their evolutionary significance is low; hence the faster rate of extinction for ontogenetic versus phylogenetic stimuli.

The processes of conditioning and extinction reveal some subtle differences between humans and animals. However, fear conditioning can be very rapid in both humans and animals, sometimes involving a single trial. Again this makes intuitive sense, as evolutionarily speaking, both humans and animals need to be able to quickly learn what are and are not inherently threatening in order to survive. A mouse for instance, cannot afford the luxury of numerous encounters with a snake to find out that it poses a significant danger. Similarly, conditioned fears are long lasting, as testified by people's phobic responses, often based upon one traumatic experience. The passage of time does nothing to diminish this fear. However, in classical conditioning, repeated exposure to the CS in the absence of the UCS can over time, lead to *extinction* of the fear. In other words, if the CS is presented a number of times without the UCS, the elicited fear will lessen and eventually, may even disappear altogether. It is important to note however, that extinction does not imply a total cessation of the relationship between the CS and the UCS. For example, rats receiving electric shocks paired with tones in one chamber, when placed in a second chamber and subjected to the same process followed by extinction, will exhibit fear when placed back in the first chamber (Bouton, 1994). Similarly, an extinguished CR can for no apparent reason become renewed almost out of the blue, a process Pavlov referred to as *spontaneous recovery*. Similar phenomena occur in humans, where previously controlled phobic or treated panic disorder patient's relapse. One of the problems with animal work on fear, however, is that fear learning may not be the same as in humans, particularly with respect to extinction (LeDoux, 1998). In addition, neuropsychological accounts of fear and anxiety, (such as Gray's (1982, 2000) model reviewed below), assume that fear and anxiety processes are mediated in a similar fashion across both human and animal subjects. Whilst the biological structures mediating this experience may indeed be similar, it is very doubtful that the cognitive experience is similarly perceived (Eysenck, 1992).

Finally, a particularly interesting finding, although equivocal, is the relationship between temperament/personality and conditioning. Introverts have generally been found to acquire conditioned fear responses more readily than extraverts, and with higher levels of acquired fear (Gray, 1982). As such, these findings could go some way to explaining the relationship between personality and anxiety, particularly with respect to the elevated neuroticism scores consistently exhibited by individuals high in anxiety.

2.9 Neuropsychological Theories of Anxiety

As the fields of psychology and neuroscience became more amalgamated, investigators began to study the fear system and formulate neuropsychological theories of anxiety. Possibly one of the most ambitious theories was that of Gray (1982; Gray & McNaughton, 2000), who studied the behavioural effects of anxiolytic (anti-anxiety) drugs. Gray examined drugs that worked on the septo-hippocampal system and showed behavioural effects related to anxiety, primarily by reducing the release of noradrenalin and serotonin into the septo-hippocampal system, thus altering the encoding of theta activity. Gray assumed that the effects of these drugs indicated the location of the impairment in functioning. He termed this system the Behavioural Inhibition System (BIS) and hypothesised that the septo-hippocampal system played a key role in mediating anxiety.

In this system, the hippocampus acts as a comparator, receiving information about the world and how it should be, and checks the ‘fit’ between predicted and actual events. If there is a discrepancy between predicted and actual events, the BIS is activated. Individuals high in trait anxiety have a much more attuned and active BIS than individuals low in trait anxiety, and are more sensitive to novelty (itself a prime for fear), and any other indicators of punishment and non-reward. The BIS governs the approach and avoidance of stimuli that indicate potential dangers, whether these be innate (IS, Innate Stimuli) e.g. cat odour to a rat, or as a result of conditioning. The BIS mediates responses to secondary punishing (CS-Pun+) and secondary rewarding (CS-Rew+) stimuli, with the comparator acting as the cognitive 'heart' of the BIS. Central to the theory is that the anxiolytics act upon the BIS, thus attenuating its output. However, psychophysiological evidence to support this claim is limited, particularly in studies of normal subjects (Eysenck, 1992). Furthermore, it is unlikely that the comparator always activates the BIS when a mismatch is detected, particularly when the mismatch is positive e.g. expecting a punishment (CS-Pun+) but receiving a reward (CS-Rew+; Eysenck, 1992).

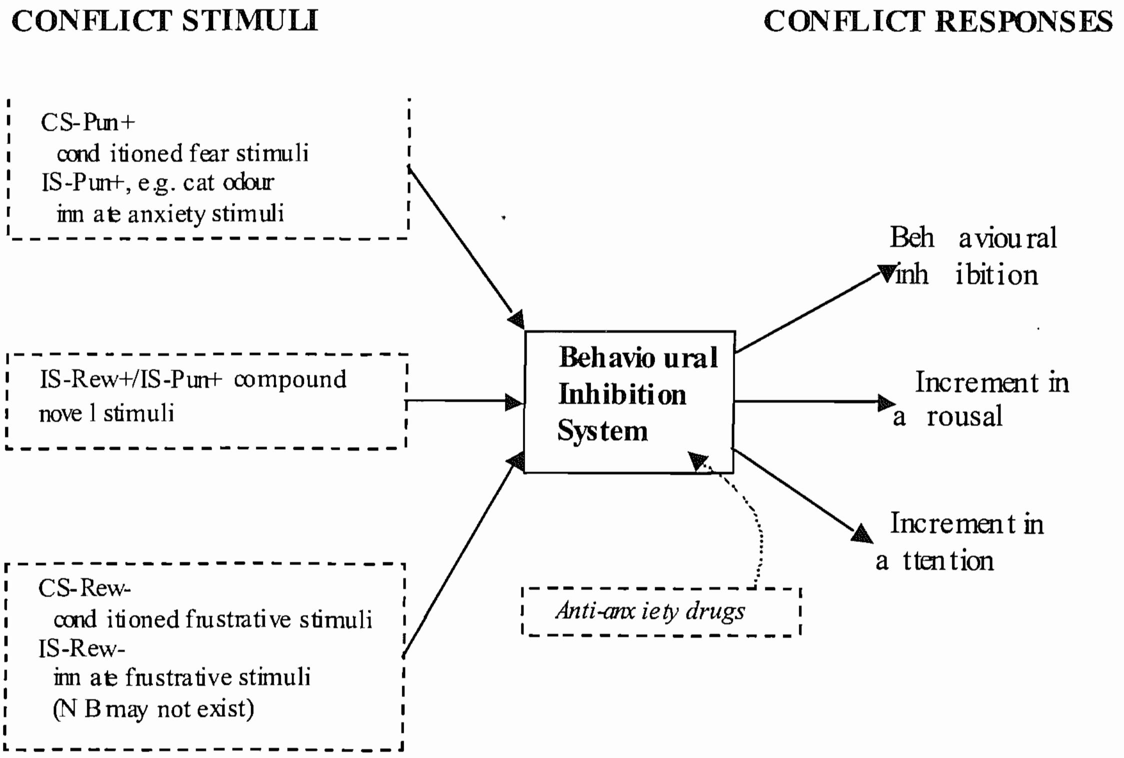

Figure 2.1. The Behaviour Inhibition System (BIS) revised version. Taken from Gray & McNaughton (2000), p.54. Inputs in the boxes on the left hand side are described using language derived from learning theory in place of the original terms (signals of punishment/non-reward, novel stimuli and innate fear stimuli). For a more comprehensive account see Gray & McNaughton, Chapter 3.

Other theories proposing that emotion is centred on just two basic systems have roots that can be traced back to Darwinian evolutionary theory, in which the foundation of emotion has a simple two-factor organisation. This two-factor organisation of emotion has been proposed by a number of theorists over the years and essentially centres upon emotional actions being either appetitive (or preservative) or defensive (commonly referred to as aversive or protective). Whilst emotional states may be composed of a number of highly varied and often universally recognised responses (e.g. Ekman, 1984) their fundamental function may be described by the two factors above. Examples of preservative emotions include sexual passion, joy and nurturance, whilst protective affect can be seen in states such as fear and anger (Lang, Bradley & Cuthbert, 1998). The link between these theories and cognitive models of attention will be returned to in later chapters.

2.10 Cognitive Theories of Anxiety

If the surge of interest in the neurobiology of fear and anxiety has grown from a trickle into a stream, then using the same analogy, the interest in adopting a cognitive approach to the study of emotion has grown into a tidal wave. One of the most influential cognitive models of anxiety was proposed by Beck (e.g. Beck, 1976; Beck, Emery & Greenberg, 1985). Beck speculated that the aetiology of anxiety was not attributable to cognition alone, but to a complex interaction of genetic predisposition, psychological and sociological factors. Although distorted cognitions could contribute to the acquisition of anxiety disorders, the primary dysfunction in anxiety disorders was in the cognitive system. Beck further extended this claim to include pathological cognition in depressive disorders, reasoning that as there is a substantial co-morbidity between depression and anxiety, similar mechanisms may exist in the two disorders (DSM-IV).

One of the central tenants of Beck's model was a construct called schema (basically a body of knowledge, or way of thinking about properties stored in long term memory). An additional organising structure, the *mode*, consisted of rules, organised into themes. In the instance of anxiety disorders, the principle mode is the danger mode. Schemas serve to direct congruent processing resources. For example, anxious patients will attend more readily to threatening stimuli, whether physical or psychological (Eysenck, 1992). Whilst there is merit in Beck's theory, particularly the notion that pathological schemata influence threat related processing in a top down fashion (Eysenck, 1992), Beck's theory owes little to experimental psychology and more to clinical observation. (In some respects it can be regarded as an amalgam of behaviourism and introspection). In order to address the lack of scientific rigour in this theory, newer models based on experimental approaches have been proposed in recent years (e.g. Mogg & Bradley, 1998), which propose that vulnerability to anxiety stems from a lower threshold in the appraisal of threat rather than an attentional bias. These models will be reviewed in greater depth in Chapter 3.

2.11 Summary

The purpose of this chapter has been to provide a brief overview and definition of anxiety and its relationship with fear. Along the way, some of the key movements and theories have been discussed and evaluated, albeit in very minor detail. The impact and importance of neuroscience findings in relation to cognitive models of anxiety is outlined as a framework for the current set of experiments. In addition, the rationale behind adopting a primarily cognitive approach to the study of anxiety and biases in attention has been discussed, with the assumption that emotions occur as a result of cognitive interpretations of situations. The eclectic approach taken will become more apparent in the following chapter, which links the fields of cognition and emotion and details the rationale for the forthcoming experiments.

Chapter3: Attentional biases to Emotional Stimuli

"We do not see things as they are. We see them as we are." The Talmud.

3.1 Introduction

During the early part of the 20th Century, researchers were very much interested in the brain mechanisms underpinning emotion. The limbic system theory of emotion (MacLean, 1952), had been forwarded as an evolutionary theory of mind and behaviour, with the claim that the limbic system acted at a middle level between primitive urges and sophisticated reasoning (Greenfield, 2000). The exact brain structures within the limbic system were poorly defined, but included structures between the brain stem and sensory relay regions (i.e., the thalamus) and the overlying cortex (Greenfield, 2000). Indeed, such was the dominance of this theory that many researchers felt that the mysteries of emotion had been all but solved (LeDoux, 2000). Around this time, a group of British researchers (most notably Donald Broadbent and Colin Cherry) were making big waves in what was to become known as the 'cognitive revolution'. Coinciding with the advent of computer systems, many researchers embraced the cognitive approach and turned away from the study of emotion to focus on more 'traditional' cognitive processes, such as perception, attention and memory.

Today, almost half a century later, the limbic system remains a popular notion, although it has been subsequently revised, as experiments implicating the hippocampus in long term memory (a distinctly cognitive function) discredit the simplistic notion of the limbic system's architecture mediating only 'older' aspects of behaviour (LeDoux, 2000). Affective research in the 21st Century now takes a much more integrated approach, with the merger of cognitive psychology and neuroscience paving the way for a wealth of findings from the new sub-discipline of cognitive neuroscience. Equally exciting is the bridge between the once separate fields of cognition and emotion. Until relatively recently, cognitive psychology paid little attention to the topic of emotion, whilst clinicians and neuroscientists studied emotion using psychodynamic and cellular approaches respectively. Where Chapters 1 and 2 have treated attention and affect separately, the present chapter develops the link between these fields, as well as providing a review of the literature and current findings. Additionally, the chapter provides a critique of the methodology used in current research and the rationale for the paradigm developed for set of experiments outlined in chapters 4-8.

3.2 Evidence for Attentional Biases in Anxiety

Historically speaking, the bridge between cognition and emotion is still very much in its infancy. Thus, much of the work reviewed in this current section is at the forefront of research and as such, is often fairly tentative in its conclusions. However, one relatively consistent finding to have emerged is of a bias in the directing of attention to favour stimuli representing a potential threat, particularly for individuals high in trait anxiety (e.g. Eysenck, 1992; Mathews 1990; Mogg & Bradley 1998; Williams, Watts, MacLeod & Mathews, 1988, 1997, Öhman, 1992, 1996, 2001; Yiend & Mathews, 2001). This attentional bias has been shown in tasks using words (generally the emotional Stroop task; Williams, Mathews & MacLeod, 1996) as well as in more 'traditional' visual attention paradigms, such as dot probe tasks (e.g. MacLeod, Mathews & Tata, 1986) and visual search paradigms (e.g. Öhman, Flykt & Esteves, 2001). The present chapter will provide a critique of the methodology and results obtained from these experiments before outlining the rationale for experiments 1, 2 and 3.

3.3 Cognitive Models of Attentional Bias in Anxiety

In the same way that cognitive scientists have tried to extrapolate the processes underlying visual attention, a number of clinical models have been proposed to account for attentional biases in anxiety. Two of the more influential theories include the cognitive-motivational theory of Mogg & Bradley (1998) and the integrative model of Williams, Watts, MacLeod & Mathews (1988, 1997). Both theories derive, in part, from the shortcomings of schema (e.g. Beck, 1976; see section 2.10) and network models (e.g. Bower, 1981), both of which are based on the assumption that anxiety and depression are associated with biases towards emotionally congruent information. Whilst the bias in depression concerns loss or failure, anxiety is associated with danger; i.e. the 'danger node' is more active in anxious individuals (Bower, 1981). However, while Beck and Bower's models allude to biases in information processing, they are not explicit in the content of these biases. Although both anxiety and depression are associated with information processing biases, the nature of these biases are more specific than either model states, with attentional biases prevalent in anxiety, while biases in memory are commonly evidenced in individuals with depressive disorders (Mogg & Bradley, 1998).

To circumvent these shortcomings, Williams et al (1988, 1997) proposed a model in which attentional biases in anxiety were posited to result from a preattentive bias towards threat. However, the nature of this bias is influenced by the individual's level of trait anxiety. Whereas individuals with low levels of trait anxiety tend to move their attention away from threat, individual with high levels of trait anxiety evidence an opposite pattern of results, whereby their attention is automatically captured by threat. The model was revised in 1997, to incorporate an interaction hypothesis, whereby as state anxiety and/or stimulus threat is increased, differences between individuals low and high in trait anxiety become more apparent. In other words, individuals with high levels of trait anxiety become hyper-vigilant towards threat, whilst people with low levels of trait anxiety become more avoidant. Consequently, individuals evidencing attentional biases towards threat are speculated to be at greater risk of developing anxiety disorders, in keeping with findings that trait anxiety is a strong predictive factor for the development of pathological anxiety (Eysenck, 1992).

Criticisms of the Williams et al (1988, 1997) model have been forwarded by Mogg & Bradley (1998), who propose that the interactive hypothesis is counter-intuitive for attentional biases to severe threat. They argue that from an evolutionary perspective, avoiding stimuli with a high threat value (e.g. snakes) could be potentially catastrophic. Accordingly, Mogg & Bradley suggest that if the threat value is severe, then *all* subjects should orient towards it, irrespective of their level of anxiety. However, if the threat value is mild, participants with low trait anxiety may focus attention in the opposite direction as a mood regulation strategy (i.e. to maintain a positive mood). As participants with high trait anxiety characteristically attend towards threat, this may be detrimental to positive mood states and be a contributory factor to the high co-morbidity between anxiety and depression.

Although both the Mogg & Bradley and Williams et al model suggest that individuals with high levels of trait anxiety preferentially attend towards threat, a crucial difference between the models is in their interpretation of the aetiology of anxiety. Whereas Williams at al propose that vulnerability to anxiety is attributable to biases in the direction of attention, Mogg & Bradley suggest that it is the appraisal of threat, rather than the direction of attention, that is accountable for an individual's susceptibility to developing anxiety. Although the aetiology of anxiety is an important theme, as the issue of the direction and level at which attentional biases occur is fundamental to this paper, the topic of exactly where these biases occur will be given more extensive treatment. Data from paradigms measuring attentional biases and how they fit within the frameworks discussed here will be considered in more detail later in the chapter.

3.4 Preattentive Biases and Attention

Of particular interest to the set of studies outlined over the course of this paper are the experimental data suggesting that threat stimuli are detected at a very early stage of visual processing. As the detection of fear relevant (FR) stimuli (biologically prepared, genetically 'hardwired' aversive stimuli; Mogg & Bradley, 1998) is seemingly independent of the number of distractors, this has been taken as evidence to support the hypothesis that threat can be captured preattentively. This bias has been shown to occur for both normal subjects, and more particularly, in individuals with high levels of anxiety or specific phobia (e.g. Öhman, Flykt & Esteves, 2001). Much of the evidence for these findings can be attributed either directly, or indirectly, to the work undertaken in Arne Öhman's laboratory over the past two decades. Öhman has adopted a psychological and neurobiological approach to the study of emotion, whereby both components are speculated to play a key role in the aetiology and mediation of anxiety (e.g. Mineka & Öhman, 2002; Öhman & Mineka, 2001, 2003). A central tenant to this work is the importance of evolutionary "hardwiring" of the brain, as an adaptive response to fear. Reiterating a point made in Chapter 2, the fear of evolutionarily threatening stimuli is intuitively more salient than the fear of more contemporary *learned* fear stimuli, (e.g. guns and weapons). Learning Theory (e.g. Seligman, 1970) supports this notion, in which evolutionary (FR) stimuli are hypothesised to be more easily and rapidly learned, as such events presented a recurrent threat to our ancestral survival (Öhman, Esteves & Soares, 1995; Öhman & Mineka, 2003). Furthermore, this learning may occur independently of cognition (c.f. Zajonc, 1984). Thus, if we as a species are inherently predisposed to associate fear and anxiety with stimuli and events that have provided persistent threats in our evolutionary history, this may help to explain the rapid and often 'irrational' nature of many common phobias (snakes, spiders, heights and closed spaces) in addition to providing a possible explanation for preattentive bias to FR stimuli (Öhman, Esteves & Soares, 1995).

Researchers have used the labels *ontogenetic* and *phylogenetic* respectively to refer to learned fear stimuli (e.g. guns) and 'evolutionary inherent fears' such as snakes, spiders and angry faces. The role of ontogenetic versus phylogenetic stimuli in the acquisition of fear will be discussed in more depth later in this chapter (as will the fundamental question of whether threat can be truly captured preattentively). Initially however, a brief overview of the fear system is provided to link these concepts together before the evidence for preattentive bias in visual attention is discussed. Following on from this will be the rationale for the initial studies, where the clinical evidence for preattentive bias will be reviewed against contemporary models of visual attention.

3.5 The Fear System; Evidence for an Enhanced Pathway for Rapid Visual Attention to Threat

The central role of the amygdala and its related areas in fear conditioning is relatively undisputed (e.g. LeDoux, 1996, 1998, 2000), and extends to both seen and 'unseen' fear stimuli (Morris, Öhman & Dolan, 1999; Whalen et al, 1998). By presenting stimuli very briefly (usually <30 ms) followed by a mask, conscious representation is prevented. Hence, whilst objects are 'seen' in the literal sense, their conscious awareness is blocked. However, psychophysiological evidence (e.g. galvanic skin responses) suggests that although participants are unable to report the perception of objects, activity in the autonomic nervous system is automatically activated (e.g. Öhman, Esteves & Soares, 1995). Öhman & Mineka (2001) have used this as a basis for a proposed Fear Module, a mechanism used to activate defensive behaviours selective to the input to which it (the fear module) responds. Consequentially, the fear module enables the ready-made neural mechanisms to respond rapidly, with minimal neural processing to threat stimuli. As ontogenetic stimuli require more processing (Öhman & Mineka, 2001), the fear module is less sensitive (and less rapid) in responding to them. Classical conditioning can greatly expand the range of items that the fear module is responsive to. However, more salient FR stimuli are more easily and durably conditioned (and have a slower rate of extinction, e.g. Öhman, Esteves & Soares, 1995). The fear module operates at a preconscious level, particularly for phylogenetic stimuli (Öhman & Mineka, 2001). Although ontogenetic stimuli can be processed automatically (through extensive training), they are usually considered to reflect 'postconscious' automaticity (Öhman & Mineka, 2001).

As the fear system is of ancient evolutionary origin, it is organised at a sub-cortical level and similar across humans and other mammals. Its ancient origin[1] makes it automatic and relatively independent of cognition, with fast acting pathways such as the thalamo-amygdala pathway providing "quick and dirty" transmission of fear signals (LeDoux, 1996). This is done at the expense of an increased number of false positives (reacting to neutral threats as if they were dangerous). While this is costly in terms of unnecessary activation of the nervous system, it is less costly than the potentially fatal situation of not reacting to real dangers (false negatives). Öhman has found evidence for a preattentive bias in a number of studies (e.g. Öhman, Flykt & Esteves, 2001), which are discussed in greater depth later in the chapter (section 3.7). For the present, the following section integrates both methodology and data from the most popular experimental paradigms used to investigate attentional biases in anxiety.

3.6 Paradigms used in the Study of Attentional Biases: The Emotional Stroop

Until the mid 1980's, the main cognitive paradigm used in affective research was the emotional Stroop (ES). The ES has fairly reliably demonstrated interference effects (and hence attentional bias) in both clinical populations and individuals with high levels of state and/or trait anxiety (e.g. Williams, Mathews & MacLeod, 1996). The premise of the emotional Stroop is similar to that of the original Stroop task (Stroop, 1935), in which participants are required to name the colour ink in which a word is printed. If the word name and colour are congruent (e.g. the word blue is printed in blue ink), responses are consistently faster than for incongruent pairings (e.g. the word yellow printed in red ink). The ES uses a similar paradigm, the key difference being in the semantics of the words used (affective versus non-affective words, e.g. idiotic vs. antique). Results from the ES are generally consistent with the hypothesis that anxious individuals show attentional biases for threat stimuli, with high anxiety subjects demonstrating increased latencies to threat compared to neutral words (e.g. Mathews

---

[1] This notion is in some respects similar to that of the limbic system. However, the limbic system is a poorly defined concept (LeDoux, 2000) with as yet, no firm criteria as to what constitute or precludes brain areas from this system. For a fuller review see LeDoux (2000).

& MacLeod, 1985). Interestingly in this study, the increase in latency to name threat related words (using Generalised Anxiety Disorder patients) interacted with the type of words used and the participant's particular concern. Thus, although all GAD participants were slower to name words related to social threat, only participants primarily concerned with physical threat were slower to name words related to this domain. However, this interaction has not always been reliably demonstrated (e.g. MacLeod, Mathews & Tata, 1986).

As in the original Stroop task, it appears that words are appraised for meaning at an early stage at the expense of task performance, even when the participant knows it is beneficial to ignore the semantics of the word. Nevertheless, the emotional Stroop is not without its problems, particularly as a measure of attention, and has yielded somewhat mixed results (Mathews & MacLeod, 1994). Indeed, it has been suggested that the emotional Stroop may actually suppress rather than enhance interference effects. However, this interpretation may be somewhat obscured, as these effects may arise in response selection and not just selective attention (Mogg et al, 2000). Furthermore, it has been suggested that these effects may be additionally clouded by increased effort on the part of the participant as a compensatory mechanism to overcome the detrimental effect elicited as a result of processing task irrelevant information (Ibid.). A further problem inherent in the Stroop (and also in the shortly to be discussed dot-probe paradigm) is that the material to be ignored is generally presented within the foveal region (Fox, Russo, Bowles & Dutton, 2001). As noted in the first chapter, there is a strong relationship between foveal vision and attention. Therefore, across both the ES and the dot-probe, one cannot determine whether attention is drawn towards threat, or whether, once discovered, attention is held by threat (Fox, Russo, Bowles & Dutton, 2001). This argument is discussed in more detail in subsequent chapters. Finally, a further methodological problem (also applicable to the dot-probe paradigm) is a potential confound between stimulus threat value and subjective familiarity. Put simply, threat related words are likely to have a larger subjective frequency for high trait (and particularly clinically anxious) subjects than for low trait individuals. The reasons for this are twofold. Firstly, as subjects typically orient towards threat related words (MacLeod, Mathews & Tata, 1986), threatening words are more likely to reach conscious awareness in high anxious individuals. Secondly, anxious individuals typically have higher levels of depression than low anxious subjects (DSM-IV). As individuals with depression commonly evidence stronger recall of negative words and events than individuals low in depression (Williams et al, 1997), these affective biases are likely to result in increased subjective familiarity with threat and/or negative words. As such, any attentional bias evidenced toward threat may be mediated by the effect of familiarity rather than threat (Mogg & Bradley, 1998,1999).

3.7 Paradigms used in the Study of Attentional Biases II: The Dot-Probe Paradigm

Although the emotional Stroop has yielded some very interesting (and generally consistent) findings, as a direct measure of *spatial* attention it is of little value. To overcome some of these shortcomings, the dot-probe paradigm has become a more widely used measure of visual attention. Stemming from paradigms used in cognitive psychology, the methodology behind the dot-probe is elegantly simple, and its popularity as a measure of spatial attention to reveal attentional biases in anxiety has increased steadily since its introduction in the mid eighties by MacLeod, Mathews & Tata (1986). The basic premise of the dot-probe is as follows. A probe is presented in the position previously occupied by a target or distractor, and the subsequent RT to the probe taken as a measure of direct spatial attention to that location. In order to measure attentional biases to affective stimuli, threat and neutral stimuli are presented simultaneously on a monitor and a probe is displayed at either location, immediately following the stimulus offset. Stimuli used are generally negative and neutral words, although in recent years pictures (by and large, angry and neutral faces) have been used to overcome some of the shortcomings of word stimuli (discussed shortly). The common finding is that participants with high levels of state and/or trait anxiety demonstrate more rapid RTs to probes at locations that contained threat words rather than neutral words. For example, in an experiment by Asmundson, Sandler, Wilson & Walker (1992), patients with Panic Disorder (PD) evidenced significantly faster RTs to probes at locations that had words pertaining to physical threat compared to controls, consistent with the notion that patients with PD have a selective bias towards physical threat. These findings are consistent with clinical models of anxiety, which propose hypervigilance to threat as a significant component in the aetiology and maintenance of anxiety (e.g. Williams, Watts, MacLeod & Mathews, 1988, 1997).

In the original dot-probe study of MacLeod, Mathews & Tata (1986), clinically anxious (GAD) patients showed a significant bias to threat words compared to neutral words, while controls evidenced a strong trend in the opposite direction (i.e. they tended to shift attention away from the threat stimuli), thus supporting the theoretical model of Williams et al (1988, 1997). While this evidence is credible in offering support for the Williams' model, it is important to note that data were obtained from only a small number of trials (probes occurred on 96 out of 288 trials). More importantly, a possible confound existed, whereby probes were more likely to occur on trials containing threat words. Thus, as the presence of threat words could serve as a cue for the forthcoming probe, the interpretation of the results is subsequently complicated (Mogg & Bradley, 1998,1999). In contrast to the findings of Mathews & MacLeod (1985) with the emotional Stroop task, there was no interaction between vigilance to the type of threat words and the subject's dominant concern in this dot probe study. Put another way, GAD participants whose dominant concern was related to physical threat did not demonstrate a bias only to physically threatening words, while similarly, GAD participants whose dominant concern was social threat did not show a bias only towards socially threatening words. Rather GAD participants as a whole showed an overall bias towards threat related words in general.

However, results were somewhat different in a more recent study by Keogh, Dillon, Georgiou & Hunt (2001), with individuals high and low in physical anxiety sensitivity (a personality trait in which individuals show an exaggerated fear of anxiety related sensations). Participants high in physical anxiety sensitivity oriented towards physically threatening words, while individuals low in physical anxiety sensitivity showed avoidance of threat related words. For participants high in physical anxiety sensitivity, this attentional bias was evidenced only towards physically threatening words and not socially threatening material, suggesting attention is not directed to threat per se, but rather towards the subject's primary concern. Again, while these data suggest that the nature of attentional biases is very much related to the individual's dominant concern, the evidence is somewhat mixed.

Although attentional biases in anxiety have most frequently been studied using participants with high levels of trait anxiety (Broadbent & Broadbent, 1988), there have been a number of studies in which state anxiety has been used as the critical manipulation (e.g. Fox, Russo, Bowles & Dutton, 2001). Furthermore, evidence to suggest that these biases begin at an early age can be seen in studies using children (e.g. Vasey, ElHag & Daleiden, 1996). However, in the Vasey et al (1996) study, in contrast to the theory of Williams et al (1988, 1997) an attentional bias away from threat was found only for low anxious boys. For girls low in anxiety, no significant differences were evidenced between threat and neutral words.

While the dot-probe paradigm has proven to be fairly robust in demonstrating an attentional bias towards threat in both individuals with elevated state and/or trait anxiety and clinically anxious populations (see Mogg & Bradley, 1999), a number of studies have failed to replicate these findings (e.g. Asmundson & Stein, 1994; Pett, 2001, Yiend & Mathews, 2001). Thus, the evidence although persuasive, remains equivocal. Over time, the dot-probe paradigm has evolved subtly in order to address some of the weaknesses inherent in its design. (For a more extensive review of this topic, see Mogg & Bradley, 1999).

As previously discussed, the use of single word stimuli introduces a potential confound of subjective familiarity, whereby negative words may be more frequently used/encountered by anxious individuals. In addition, as words may lack ecological validity and are of limited threat value, their suitability is questionable. To address these potential problems, picture stimuli have been used to good effect in a number of studies (e.g. Bradley, Mogg & Millar, 2000; Mogg & Bradley, 1999; Mogg et al, 2000). While the use of picture stimuli undoubtedly addresses some of the above weaknesses, the dot-probe is not without a number of other problems. Stimulus display in dot-probe experiments is usually presented for a period of 500 ms (Bradley, Mogg & Millar, 2000). A potential problem lies in the fact that a temporal interval of this duration allows participants the chance to shift attention between the two stimuli. In a neutral/threat pair of words, the subject may initially attend to the neutral item before switching attention to the threat word. Hence, the attentional bias may not reflect the initial orienting response, but rather, a shift of attention to more salient material. Equally, attention may be initially captured by the threat item, but rather than switching to the neutral word, it may stay at the more salient item. If the former assumption is true, then the attentional bias should be considered more as increased post attentive awareness (Broadbent & Broadbent, 1988) rather than a preattentive bias. However, empirical evidence showing biases in eye movements towards threatening facial expressions for individuals with GAD compared to controls and individuals with depressive disorder (Mogg, Millar & Bradley, 2000) indicates that the initial orienting response is very much dependent upon the symptomatic state of the individual. Furthermore, an experiment by Bradley, Mogg & Millar (2000) demonstrated that the RT measures for attentional bias towards negative faces were strongly related to the direction of initial orienting to affective stimuli, suggesting initial capture by threat. Again however, this interpretation is not necessarily valid. Whilst eye movements and covert attention are strongly associated (Palmer, 1999) it is important to note that covert attention and eye movements can operate independently. Therefore, whilst eye movement is indicative of overt attention, covert attention is not so easily measured.

Although combining physiological measures within the dot-probe paradigm allows one to interpret data to a much greater extent, the restricted number of items in the display limits the scope of the findings. While the dot-probe enables tight experimental control, it is somewhat offset by a lack of external validity. The real world is composed of complex visual scenes in which the observer is rarely (if ever) afforded the luxury of such a limited number of distractors. As such, interpretations of findings using this paradigm are limited in scope, not only because of the restrictive number of distractors present in any display, but also because of the difficulty in interpreting the level at which attentional biases might operate. As the dot-probe does not use any variation in the number of distractors present, one cannot calculate the efficiency with which participants search for target stimuli (or make inferences about the underlying type of search mechanism). Therefore, any attentional biases could equally reflect either the rapid engagement of attention to threat or a difficulty in disengaging attention once a threat stimulus has been located.

According to Posner & Peterson (1990), the attentional system comprised 3 sub-systems (disengage, move and engage), each controlled by separate brain areas. Posner & Peterson posited that in order to attend to an object, an individual must first disengage their attention from whatever is being attended to, before moving attention and finally engaging upon the object of interest. Although this account of attention may be too simplistic (see Chapter 1), it highlights the fact that bias in attention might be attributable to factors other than biases in initially orienting to stimuli. In keeping with this, data from a recent experiment by Fox, Russo, Bowles & Dutton, (2001) suggest that the attentional biases observed in the dot-probe toward threat stimuli may reflect increased dwell time and difficulty in disengaging from threat.

3.8 Paradigms used in the Study of Attentional Biases III: Cueing Experiments

In keeping with the proposition of Fox, Russo, Bowles & Dutton (2001) that biases towards threatening stimuli might reflect differences in the disengage component of attention, Yiend & Mathews (2001) ran a series of experiments to investigate the mechanisms of attention using students with low and high levels of trait anxiety. In the initial experiment (dot-probe), no significant differences between groups were observed in subjects' initial allocation of attention (although there was a trend for high trait anxious participants to preferentially attend to negative images). However, further experiments using a cuing paradigm, whereby threatening or neutral pictures were briefly presented as location cues (see also Chapter 1), revealed that participants in the high trait anxiety group evidenced significantly slower RTs on invalid threat trials. Furthermore, these differences were more pronounced for high than mild threat, in keeping with both theoretical models (e.g. Mogg & Bradley, 1998) and empirical data (Mogg et al, 2000). Yiend & Mathews interpreted this as evidence that individuals high in trait anxiety do not necessarily preferentially attend to threat stimuli, but rather, their difficulty lies in disengagement from threat stimuli. Rather unexpectedly, however, there were no apparent costs or benefits for neutral cues in either group. In other words, when a neutral cue appeared in a different location to the target, there was no significant *cost*, i.e. no slowing in subjects RTs. Similarly, when the neutral cues appeared at the same location as the target (valid cues), there were no discernable *benefits*, (i.e., the RTs were no faster). These findings are contrary to the data generally obtaining using cueing studies (e.g. Posner, 1980), in which valid cues generally demonstrate significant benefits, whilst invalid cues show significant costs. As the paradigm used was similar to that of normal cueing experiments (with the exception that the cue stimuli were affective pictures instead of briefly brightened boxes), these results are rather puzzling and merit further investigation. However, the SOA between the cue and target was 500 ms rather than the 150-300 ms generally used in experiments showing a clear RT advantage for cued targets. The longer SOA's in this experiment may leave open the possibility for effects from inhibition of return (Posner & Cohen, 1984).

In a final experiment, using an identical paradigm, but with a longer SOA between cue and target (changed from 500ms to 2 s), the phenomenon of Inhibition Of Return (IOR) was investigated. In a normal cueing experiment, participants are reliably faster to detect targets that appear in valid locations compared to targets, which appear in invalid locations. However, if the temporal interval between cue and target onset is greater than 300 ms (Posner & Cohen, 1984) target detection in the valid location is reliably slower than detection of targets in invalid locations. This phenomenon, known as IOR, is the withdrawal of attention from a location in order to inhibit attention visiting the same location repeatedly (Hoffman, 1998). Such a mechanism affords an obvious advantage with respect to the efficiency of visual search. Following this increase in the SOA between cue and target, both low and high trait anxiety subjects demonstrated significant benefits in detecting targets following invalid neutral pictures, but not for invalid threat pictures. While these results show an effect of threat on disengagement, there were no main effects of anxiety; all participants were faster to the invalid neutral cues. Such findings (i.e. benefits for invalid cues) are in keeping with results from studies on IOR (e.g. Posner & Cohen, 1984). However, as the effects of IOR usually arise following relatively brief temporal intervals between cue and target onset (generally around 300 ms), one would expect to see evidence of IOR in the small SOA condition also.

Another experiment, again based on the Posner cueing task, used threat and neutral words as cues in an attempt to extrapolate the neuropsychological basis of threat related attentional biases in anxiety (Avila & Parcet, 2002). Subjects were either given information about the cue validity (informed condition) or no information concerning the cues' predictive validity (non-informed condition), and SOA between cue and target was either short (100 ms) or long (500 ms). Contrary to the results of Yiend & Mathews (2001), as expected, significant costs were observed for both affective and neutral invalid cues. With respect to groups, anxiety differences were observed for high anxiety subjects in the non-informed short SOA condition, whereby there were greater benefits for threat related valid cues than neutral valid cues. For low trait anxious subjects, these results were reversed, i.e., benefits were greater for neutral valid cues than threat valid cues. As these findings were not found in the longer SOA or the informed condition, the authors interpreted this as evidence supportive of an automatic effect, temporally related to the appearance of the cue. However, as this study used word stimuli rather than pictures, it is difficult to compare the results with those of Yiend & Mathews (2001), as picture cues may elicit a more potent source of threat. In addition, as the long SOA used in the Avila & Parcet (2001) study was the same duration as the short SOA used in the Yiend & Mathews (2001) study, comparisons are problematic because of the possible complications of IOR. Nonetheless, if methodological problems are carefully addressed, the cueing paradigm shows promise as a means to investigate the level at which attentional biases may occur.

3.9 Paradigms used in the Study of Attentional Biases IV: Visual Search

Like the dot-probe, visual search is a valuable paradigm with which to measure attentional biases in anxiety. Visual search has been widely used to differentiate between preattentive processing and focused attention, by making inferences about the search mechanisms as a function of RT by set size (although see Wolfe (1998a) for a comprehensive account of the dangers of labelling searches 'serial' and 'parallel'). The basic premise behind the visual search paradigm is relatively simple and is covered in detail in chapter one. However, while visual search has been a popular research tool within visual cognition for almost a quarter of a century, its use as a paradigm to investigate attentional biases in anxiety is a relatively recent development. One of the earliest studies to utilise this paradigm was a classic study by Hansen & Hansen (1988) in which participants were required to find a discrepant face in a crowd. The data showed that participants were significantly faster to detect a threatening face amongst friendly faces than vice-versa. More importantly, as the time taken to detect an angry face did not increase with corresponding increases in set size (Experiment 3), the findings were taken as evidence to support the notion that threat stimuli can be detected preattentively.

While Hansen & Hansen's findings were important, their work soon came under criticism, including from the authors themselves (Hampton, Purcell, Bersine, Hansen & Hansen, 1989). A number of explanations for their findings have since been posited, such as the possibility that friendly faces may be processed more efficiently than threatening faces because of their familiarity (Öhman, Lundqvist & Esteves, 2001; Treisman & Souther, 1985; Wolfe, 2001). In addition, the faster processing of friendly targets, demonstrated by faster RTs to detect the absence of discrepant targets in 'friendly' arrays compared to 'threatening' arrays of faces (Hansen & Hansen, 1988) may be due to the subject's ability to disregard non threatening faces more rapidly. Evidence in favour of this interpretation can be found in an experiment by Treisman & Souther (1985), in which subjects were faster to reject familiar, normal letters than unfamiliar, mirror reversed letters. Thus, as it is easier to reject familiar than unfamiliar items, subjects should be faster to reject the ubiquitous smiling face than the less familiar angry face. More important was the discovery of a low level perceptual confound, whereby conspicuous dark areas were apparent on the necks of the angry and happy female faces and on the chin of the angry male faces (Purcell, Stewart & Skov, 1996). This confound meant that the more rapid response to identify the threatening faces could be attributable to unique features rather than affective expression per se. In order to extrapolate whether the preattentive bias apparent in the Hansen & Hansen (1988) study was indeed a real phenomenon, Öhman, Lundqvist & Esteves (2001) carried out a series of 5 experiments using schematic faces (rather than photographic images) in what was essentially a replication of Hansen & Hansen's (1988) original study.

Although the use of schematic faces compromises ecological validity, the obvious advantage afforded by this methodology lies in allowing the experimenter tight control over the physical features of the faces presented. While most individuals find it difficult to produce a convincing threatening face on demand, most people are inherently capable of producing a reasonably convincing happy face (Öhman, Lundqvist & Esteves, 2001). Consequentially, the use of photographic images is likely to result in a more heterogeneous set of angry faces than it is for happy faces. As heterogeneity of distractors is an important variable in RTs within visual search paradigms (Duncan & Humphreys, 1989; Wolfe, 1998a), the lack of ecological validity is somewhat offset by the tight control afforded by the schematic faces. Furthermore, as most people are able to identify the affective expression of schematic faces accurately (Fox et al, 2000) their use may be less problematic than is initially apparent.

The results of the Öhman, Lundqvist & Esteves (2001) study showed a generally consistent threat advantage, whereby schematic threatening faces were detected more rapidly than happy faces, particularly when the distractors were neutral. The latter finding is unsurprising given that the neutral faces were defined by a feature, i.e. a horizontal line for the mouth, whereas the 'emotional' (happy or angry) schematic faces were defined by a conjunction of features (e.g. an angry face and a happy face shared all of the same features, with the orientation of the 'mouth' and 'eyebrows' delineating the affective expression). As such, the homogenous nature of the neutral distractors led to faster searches. Importantly however, although there was a modest threat superiority effect as evidenced by the RT intercepts, there was no threat advantage in the search slopes. Interestingly, when emotional distractors were used, the threat advantage was only significant in experiments that used a longer exposure time (2 seconds rather than 1 second). Such findings seem rather at odds with the idea that the detection of threat is very rapid (sometimes automatic) as argued by proponents of the phylogenetic hypothesis. With emotional distractors, search was more effortful, indicating the need for post attentive processing. Hence, whilst the search asymmetries indicate that threatening faces have some unique properties that lead to more efficient search, searches of this nature still require focused attention.

There are two reasons to question whether efficient search performance in Öhman, Lundqvist & Esteves' (2001) search experiment might be due to relatively simple combinations of features rather than to the preattentive identification of threat. Firstly, if faces are inverted, then their perception is disrupted (Yin, 1969). Thus, if affective expression is the critical factor, then RTs will be slower for inverted faces as holistic processing is disrupted. On the other hand, if subjects are finding the angry faces by searching for relatively simple configurations of lines, then their performance should be similar whether the faces are inverted or upright. Öhman et al's data showed an equivalent anger superiority effect for both inverted and upright faces.

The second reason to question the role of threat comes from a series of face search experiments by Purcell & Stewart (2002). They asked participants to press a key whenever a discrepant face was present and to withhold their response when only distractors were present. Displays consisted of happy and angry faces, and also control faces that consisted of little or no emotional expression, but included certain configural properties that were also found in the happy and angry faces. There were "line-in" faces with inward pointing eyebrows and mouth to match configurations in the happy faces, and "line-out" faces with outward pointing eyebrows and mouth to match configurations in the angry faces. Results showed superior detection for the "line out" configuration, with the pattern of results being almost identical to the findings of Öhman et al (2001). Hence, the threat superiority effect evidenced in Öhman's study is likely once again to be attributable to relatively simple configurations of lines rather than affective expression per se.

Eastwood, Smilek & Merikle (2001) employed a methodology similar to that of Öhman, Lundqvist & Esteves' (2001), in which participants were required to search for positive or negative schematic faces embedded amongst neutral distractors. Again, participants were faster to detect negative faces compared to positive faces. As in the Öhman Lundqvist & Esteves experiment, while the differences in the search slopes were significant, no 'pop out' was evidenced. Furthermore, when the faces were inverted in a second experiment to limit the recognition of affective expression, the slope difference between searches for angry and happy faces diminished and was no longer significant. Eastwood et al (2001) interpreted this as evidence to suggest that the differences in attention were due to differences in affective valence between the stimuli rather than featural differences. A close look at the results reveals a puzzling pattern, however. There was very little difference in for RTs to inverted negative faces compared to upright negative faces, whereas inverted positive faces were detected more quickly than upright positive faces. Hence, the more rapid detection for negative faces might somehow be connected with the fact that they are more novel than the ubiquitous smiley face rather than enhanced detection due to their affective expression.

In a series of experiments by Fox et al (2000), again using schematic faces, subjects were again faster to detect an angry face amongst happy faces than vice-versa. Furthermore, when set size increased, the increase in RT in detecting angry faces was less than the increase for happy faces. However, whilst the search rate of approximately 17 ms per item for the angry faces was rapid, it cannot be considered to reflect ”pop-out”, which usually yields figures of 10 ms/item or less (Fox et al, 2000. Treisman & Souther, 1985). In keeping with the findings of Hansen & Hansen (1988), participants were slower to detect the absence of a discrepant face when the display contained angry faces compared to happy face displays. The authors interpreted this as reflecting stroop-like interference, possibly due to increased dwell time or problems in disengaging attention away from the threat stimuli. Again, this difference could be due to the fact that participants are slower to discard angry distractors compared to happy distractors, in line with the proposed explanation for the Hansen & Hansen (1988) results. Again, when displays were inverted these differences disappeared, supporting the notion that emotionality rather than feature detection plays a pivotal role in visual searches of this nature. In summary, although face search experiments have evidenced search asymmetries in which threatening faces are located more rapidly than friendly faces, "pop-out" (Treisman & Gelade, 1980) has not been demonstrated. Furthermore, results remain equivocal due to simple featural differences between threatening and happy faces (e.g. Hansen & Hansen, 1988; Öhman, Lundqvist & Esteves, 2001) and ambiguity over whether threatening faces are detected faster or happy faces are detected more slowly (e.g., Eastwood et al, 2001). Although the results from the Fox et al (2000) experiment are less ambiguous, their data do not clearly indicate whether threatening faces are detected faster or reflect increased dwell time/difficulty in disengaging attention from threatening faces.

While the use of schematic faces is advantageous for a number of reasons, their lack of ecological validity remains problematic. In order to redress some of the weaknesses inherent in this approach, and to enable comparisons across studies, Öhman, Flykt & Esteves (2001) carried out a series of experiments in which participants searched through matrices comprised of colour photographs. For each stimulus display, participants responded as quickly as possible whether all items came from a single category (e.g., all flowers) or whether one item belonged to a different category (e.g., a spider among mushrooms). Displays were comprised of either FR targets amongst FI distractors (i.e. a snake or spider amongst flowers or mushrooms), FI targets amongst FR distractors (i.e. a flower or mushroom amongst snakes or spiders) or all items from one category (all snake or spider photos or all flower or mushroom photos). It was found that participants were faster to detect FR stimuli amongst FI distractors than vice-versa. Furthermore, when the size of the matrix was increased from a 2x2 to a 3x3 design (increasing the number of distractors from 3 to 8), there was no significant effect on the participants' RTs in locating FR targets. In a final experiment, participants were selected on the basis of high snake or spider fear or low snake/spider fear (control group) and carried out the same experiment. Again, all participants were faster to detect FR targets than FI targets, and once more, their detection was unaffected by location or set size. More importantly, fearful participants were faster to detect their feared than their non-feared FR target (i.e. high spider fear participants were faster to locate spider targets than snake targets, while high snake fear participants were faster to detect snakes than spiders). Importantly, the faster response times were not at the expense of a speed/accuracy trade off.

Whilst the use of photographic stimuli circumvents the problem of ecological validity, it raises the possibility that the results may be attributable to low-level perceptual features rather than threat value. However, the fact that the high snake fear participants were faster to detect snakes than spiders, whilst the high spider fear participants showed a reversed pattern of results, supports the notion that the detection of FR targets was enhanced because of threat value rather than low level perceptual features. In addition, as a mixed category of pictures was used, the heterogeneous nature of the displays limits the chances of a common perceptual artefact, although it is still possible that snakes and spiders have simple defining features that mushrooms do not. As phobic subjects are more motivated to locate the object of their fear and may have access to 'templates' of the features of their feared object (e.g. a 'coiled' spherical representation for snakes), this may be responsible for their more rapid detection. These findings, in conjunction with data showing an anger superiority effect for schematic faces (Eastwood, Smilek & Merikle, 2001; Fox et al, 2000; Öhman Lundqvist & Esteves, 2001), offer persuasive evidence in support of phylogenetic explanations of attentional bias (whereby threat is posited as an important factor in the automatic capture of attention). Furthermore, it is suggested that the search asymmetries revealed in these studies might be due to the fact that FR stimuli possess characteristic elements that automatically capture attention, as they contain a signature feature, which because of evolution, we have become highly attuned to (Öhman Flykt & Esteves, 2001).

Although these findings are persuasive, the evidence that threat stimuli are a unique source for the automatic allocation of attention is equivocal. In a recent series of experiments using visual search methodology and picture stimuli (Tipples, Young, Quinlan, Broks & Ellis, 2002), subjects evidenced faster search rates for threatening animals (snakes, bears and dogs poised to attack) than plants. In line with findings of Öhman Flykt & Esteves (2001), when the display size of the search array was increased from 4 to 9 items, RTs for the threatening animals increased at a rate of only 11ms/item, whereas for plants, the increase was 28 ms/item. While the former category fall close to the <10 ms/item search rate required for 'quite efficient' searches (Wolfe, 1998a), and 'pop-out' (Treisman & Gelade, 1980), the search rate for plants is best described as 'inefficient' (Wolfe, 1998a) Although these data at first glance appear to provide firm evidence for enhanced attention to threat stimuli, data from a second near identical experiment, in which participants searched for pleasant animals (e.g. horses, kittens, 'relaxed' dogs etc.) hidden amongst plants and vice-versa revealed a different interpretation of the results. As in the first experiment, the smaller set size effect for animals than for plants was repeated, with a significantly faster rate of search for pleasant animals than plants (8 ms/item vs. 24 ms/item respectively). As participants were able to search for both threatening and pleasant animals at a roughly equal rate, it appears that threat per se confers no advantage. To eliminate the possibility that one target category may be easier to either locate or reject than another (as previously discussed), and to enhance ecological validity (as real world searches for threat are rarely afforded the luxury of homogenous distractor categories), a further experiment incorporated neutral distractors (e.g. umbrellas, mugs, computers etc.) in addition to the picture categories used in experiments 1 and 2. Again, participants were faster to detect both threatening and pleasant animals than plants.

A further experiment in which targets were presented either close to or far from fixation (to test whether peripheral threat stimuli confer any advantage) showed no threat advantage for targets near or far from fixation. Rather, in keeping with cognitive models of visual attention, RTs to targets close to fixation were more rapid than those presented in peripheral locations. Finally, fruit pictures were added to the search matrices, and participants asked to search for both threatening and pleasant animals as a unified category in order to examine the possibility that *all* animals are seen as a potential threat and to ensure that any RT advantage reflects enhanced target detection rather than target rejection. Once again, RTs to threatening and pleasant animals *and* fruit were roughly the same, whilst detection of flowers was significantly slower. Thus, while search for all animals was more rapid than for flowers, it was no more efficient than search for fruit. These findings discredit the notion that all animals (whether pleasant or threatening) are detected more efficiently than other stimuli as an evolutionary advantage because they are all a potential threat, as there is no clear explanation why fruit should be detected with equal efficiency. In summary, these data offer a convincing argument against the findings of Öhman, Flykt & Esteves (2001). Although the detection of threatening animals was more rapid than the detection of plants, in keeping with data from Öhman, Flykt & Esteves (2001), search for pleasant animals and even fruit was equally efficient.

Although the results from the above study present a strong case against a threat superiority effect, it is important to note that in all the experiments, neither state or trait anxiety was measured or manipulated. Thus these results are not inconsistent with claims that attentional biases occur primarily in individuals with high levels of state and/or trait anxiety (e.g. Williams et al, 1988, 1997). However, these results do raise questions about claims that certain stimuli (snakes, spiders and angry faces) capture attention automatically (Öhman & Mineka, 2001) and elicit different neural responses (Morris et al, 1996; Morris, Öhman & Dolan, 1998). It should be noted that this study, like the study by Öhman, Flykt & Esteves (2001), used colour pictures, which undoubtedly differ in their complexity, colour and numerous other attributes. While the use of a large number of different images reduces the possibility of perceptual confounds, the study by Tipples et al has the same potential for perceptual confounds as the original study by Öhman, Flykt & Esteves.

3.10 Paradigms used in the Study of Attentional Biases V: Imaging Studies

A newer methodology with which to measure the effects of emotion on spatial attention is the use of brain imaging. In order to extrapolate whether fearful faces can be processed when attention is engaged elsewhere, Vuilleumier, Armony, Driver & Dolan (2001) employed fMRI measures to investigate the amount of processing that task-irrelevant emotional stimuli receive. Participants were presented with briefly displayed (250 ms) pairs of houses or faces and instructed to make same/different discriminations. The critical measure was the extent of processing allocated to fearful versus neutral unattended faces. The RTs showed that fearful expression in the task irrelevant condition slowed down performance compared to neutral expression, thus indicating some degree of processing of the task irrelevant fearful stimuli. Fusiform gyrus activity was modulated by the attentional condition, with greater right fusiform activity to fearful than neutral faces. Importantly, amygdala responses (particularly the left amygdala) to fearful faces were unaffected by the attentional manipulation, indicating automatic processing of fearful faces, independent of attention.

A similar study in which subjects were required to indicate the gender of a briefly (200ms) presented face (attended trials) or make same/different discriminations of the orientation of bars presented to the top left and right of the monitor (whilst ignoring the centrally presented face; unattended condition) was conducted by Pessoa, Kastner & Ungerleider (2002). As expected, fearful faces produced greater bilateral (particularly left) activation of the amygdala in the attended condition. However, in the unattended condition, there were no significant differences in the amygdala by stimulus type. As such, it would appear that, contrary to the notion advocated by LeDoux (1996), amygdala responses to affective stimuli are not automatic and are in fact, mediated by attention. These data are contrary the findings of Vuilleumier et al (2001) and as such, any conclusions are tentative. However, Pessoa et al suggest that as their study evidenced error rates of 36% compared to 14% in the Vuilleumier et al study, the attentional load was greater. Thus, although emotional stimuli receive prioritised processing when the task is relatively undemanding, as the perceptual load increases, attention becomes necessary (cf. Lavie, 1995).

3.11 Summary

The aim of this chapter has been to review the literature and experimental data to date, in order to examine the extent to which emotional stimuli affect visual attention. The methodologies used (and the limitations inherent in each of the different paradigms) have been discussed with reference to current findings in visual cognition, and are summarised in Table 3.1. From this discussion of the data and the paradigms used, it is hoped the rationale behind the paradigm used in Experiments 1-5 will become more apparent. More specifically, the set of experiments detailed over the following chapters have been designed to capitalise on the strengths of experiments using picture stimuli (ecological validity) and schematic stimuli (tight control of perceptual features) whilst circumventing problems inherent in each (i.e. a lack of ecological validity or control over perceptual features). Should the experiments be successful, the findings will either require a revision of current models of visual attention or clinical theories of attentional biases in anxiety.

Paradigm	Measurement/DV	Advantages	Disadvantages	References
Emotional Stroop	Interference effects/ RTs, errors	Simple to administer; can be used as a 'paper and pen' test, therefore highly portable. Generally reliable findings.	Cannot be used as a direct measure of spatial attention. Distractors presented in fovea. Cannot determine stage at which biases occur. Words may lack a strong threat value.	Mathews & MacLeod, 1985, 1994; Williams *et* 1996
Dot-Probe	Spatial attention/ RTs, errors	Generally reliable. Provides a more direct measure of spatial attention. Greater ecological validity (when photographs are used).	Often long temporal interval (~500ms) allows for switches in attention. One cannot accurately determine the level at which attentional biases occur.	Asmundson *et al*, 1992; Bradley *et a* 2000; MacLeod *et* 1986; Mogg *et al*, 2000.
Cueing	Spatial attention/ RTs, error rates.	Allows inferences to be made regarding the stage at which bias occurs (although cannot differentiate between serial and parallel mechanisms). Generally reliable findings across both word and picture stimuli.	Cannot be used to measure preattentive processing. Limited number of target positions compromises ecological validity. Not generally used with distractor stimuli.	Avila & Parcet, 2002; Fox *et al*, 2001; Yiend & Mathews, 2001.
Visual Search	Attention/ RTs, error rates.	Enables inferences to be made re underlying search mechanisms (i.e., preattentive/efficient vs. serial/inefficient) search. Photographs increase ecological validity.	Low-level perceptual properties cannot be controlled if photographs are used. Schematic stimuli are less ecologically valid and might still combine to form features that can be used to guide search (e.g. Purcell & Stewart, 2002).	Eastwood *et al*, 2001; Fox *et al*, 2000; Hansen & Hansen, 1988; Öhman, Flykt & Esteves, 2001; Öhman, Lundqvist Esteves, 2001, Tipples *et al*, 2002
Imaging Studies	Neural /BOLD activity, (RTs, error rates).	Enables neural activity/ structures subserving psychological functioning to be measured. Good spatial resolution. Can be combined with behavioural measures.	Poor temporal resolution. Expensive. Needs combining with other measures.	Pessoa *et al*, 2002; Vuilleumier *et al*, 2001

Table 3.1. Summary of the main paradigms used in the study of attentional biases to threat.

Chapter 4: Experiment 1

4.1. Introduction

Thus far, the literature review of the preceding chapters has provided the backdrop to the set of studies outlined over the following chapters. The methodological limitations inherent in many of these studies (e.g. the Emotional Stroop and Dot-Probe paradigm) have already been discussed. More significant to the current set of studies is the visual search paradigm adopted by Öhman, Flykt & Esteves (2001), used to measure the allocation of visual attention to FR stimuli. While the results from Öhman's search experiments are very persuasive, particularly when combined with other recent research carried out in the same laboratory (e.g. Öhman, Lundqvist & Esteves, 2001), they are not without methodological weaknesses. For example, the former study used pictures of snakes, spiders, flowers and mushroom in matrices of different sizes. As a consequence, both target and distractor stimuli were composed of complex visual scenes allowing the experimenter little or no control over the perceptual features present. While there is undeniable merit in the use of ecologically valid stimuli, potential confounds exist with the possibility of differences in discrimination being attributable to low-level physical properties. Although this problem was to some extent redressed in the latter study by the use of schematic faces, there is still some opportunity for basic featural differences in the schematic faces, especially if the mouth and/or eyebrows can combine with other components to make different shapes. These points are covered in greater detail in the preceding chapter.

In an effort to redress some of these issues, Experiment 1 utilised a novel paradigm, in which target letters were paired with affective images to become FR targets. Essentially, the letters were the CS and the affective pictures the UCS. As such, any differences between RTs to neutral and negative targets could be attributable to conditioned responses to the FR stimuli. Thus, should response time, or more specifically, search slope, vary according to association with threat, then this would indicate preferential (or avoidant) search for threat. In particular, should response times be independent of set size for FR targets, this would indicate a 'parallel', or more accurately, an efficient search (Wolfe, 1998a). Whilst evidence of such search mechanisms has been demonstrated in previous studies (e.g. Hansen & Hansen, 1988; Öhman, Flykt & Esteves, 2001), such evidence is difficult to reconcile with some current cognitive models of visual attention, which purport that search rate is determined only by features shared amongst targets and distractors (e.g. Cave, 1999; Cave & Wolfe, 1990; Treisman & Gelade, 1980; Wolfe, 1994).

If attentional biases to negative valenced conditioned target letters are evidenced (whether by faster RTs or more specifically, flatter search slopes) then this would indicate that emotion affects visual search. If no changes are evidenced, however, this may simply reflect the fact that the conditioning did not work rather than the fact that threat does not preferentially capture visual attention. In order to rule out this possibility, an Implicit Association Test (IAT; Greenwald, McGee & Schwartz, 1998) was adapted to determine whether the conditioning procedure had worked effectively. The IAT measures differential associations of two target concepts with an attribute, where concepts appear in a two-choice task (e.g. flower vs. insect names or in this case T vs. L) along with attributes (e.g. pleasant vs. unpleasant words as evaluative attributes) to measure individual's implicit beliefs. For example, as the word flower is implicitly associated with pleasant attributes, the response to flower should be faster when it is made with the same hand as the response to pleasant words than when it is made with the same hand that is used to respond to unpleasant words. Similarly, if the conditioning procedure has been effective in the present experiment, then participants in the group in which the letter T was paired with negative pictures should be faster to associate the letter T with the same response as negative words, rather than the response used for positive words. The structure of the IAT is described in more detail in the method section.

To measure subjective anxiety, a number of anxiety inventories are in general use, including The Taylor Manifest Anxiety Scale (Taylor, 1953) and The Hospital Anxiety and Depression Inventory (Zigmond & Snaith, 1983). However, possibly the most widely used anxiety inventory is the Spielberger State-Trait Anxiety Inventory (Spielberger, Gorsuch& Lushene, 1970). The State-Trait Anxiety Inventory (STAI) was designed,

"to measure a stable propensity to experience anxiety, and tendencies to perceive stressful situations as threatening." (Bieling, Anthony & Swinson, 1998, p. 780) and has a high test-retest reliability of between 0.73-0.86 for the trait scale (Ibid.). A more detailed description of this inventory can be found in the main body of the experiment (see stimuli section).

Many current models of visual attention (e.g. Cave, 1999; Treisman & Gelade, 1980; Wolfe, 1995; Wolfe, Cave & Franzel, 1989) predict that that there will be no differences in the RTs of participants searching for threat-associated targets compared to participants searching for neutral associated targets. As the number of distractors increases, RTs to find both negative and neutral targets will increase linearly. Assuming that conditioning has worked, there will be a significant main effect of response pairing in the IAT. In other words, participants given the pairing of T and negative pictures in the visual search task will be faster to associate T with negative words than positive words. For participants given the pairing of T and neutral pictures in the visual search task, there should be no difference between RTs to T and negative words and T and positive words in the IAT.

Experiment 1

4.2. Method

Participants

32 Students (12 male and 20 female) from The University of Southampton participated in this experiment. Participants were recruited from the psychology undergraduate subject pool and gained course credits, or were postgraduate volunteers. Participant age ranged from19 to 43 years. Mean age was 25.22 years (SD =5.87). All participants gave written informed consent and reported normal or corrected to normal vision, including normal colour perception. Participants were screened for trait anxiety using the Spielberger State-Trait Anxiety Inventory (STAI; Spielberger, Gorsuch & Lushene, 1970) although their allocation to neutral and negative conditions was randomised.

Design

A mixed design was used. Picture Type was a between subjects factor. 16 participants (6 male, 10 female) took part in the visual search task for targets conditioned with negative valence. 16

participants (5 male, 11 female) undertook an identical visual search task, with the exception that target letters were paired with neutral stimuli.

Stimuli

Two types of picture sets were relevant for this study (see Appendix I): 24 negative valence pictures (e.g., a mutilated body, an aimed gun) and 24 neutral valence pictures (e.g., a light bulb, an office block). Images were taken from the International Affective Picture System (IAPS; Lang, Bradley & Cuthbert, 1997), a standardised set of normative emotional stimuli. Selection was done on the basis of mean ratings of valence and arousal (see appendix and also Table 1 of the IAPS technical manual). As the IAPS has been shown to have two primary dimensions, affective valence (ranging from unpleasant to pleasant) and arousal (calm to excited), images for the negatively evaluated targets were selected on the basis of low scores for valence and high scores for arousal. In contrast, images for the neutral valence condition were selected on the basis of midline scores for valence and low scores for arousal. A third, less strongly related dimension, variously called 'dominance' or 'control' was omitted when selecting appropriate images for each condition. Each scale ranges from 1 to 9, where 9 represents a high rating for the particular dimension (e.g., high pleasure or high arousal) and 1 represents a low rating (e.g., low pleasure, low arousal). Mean valence and arousal values for the negative images were 1.77 and 6.63 respectively. For the neutral images, mean valence was 4.92 and mean arousal was 2.87. In order to simplify proceedings, images were selected on the basis of having appropriate scores for both male and female subjects, thus eliminating the need for a separate set of pictures for each group. As the IAPS images varied in size, they were aligned onto a black background (as this was judged not to detract attention from the displayed image) and centred for fixation. The display size of each image, along with the black background, was 1152 X 870 pixels, and occupied the entire screen.

Stimuli used in the visual search task were target letters (T) and distractor letters (L). Both targets and distractors were presented at one of four orientations; upright (0°), inverted (180°), facing left (270°) and facing right (90°) and their presentation was randomised. Target and distractor letters were black. Each letter was 10 mm tall and 10 mm wide. Throughout all trials the background was white. Target and distractor letters were spaced around an imaginary circle with a diameter of 100 mm around fixation (see Figure 4.1). Throughout the course of the experiment, target and distractors could appear at any of the stimulus locations, which were evenly spaced around this imaginary circle. A black fixation cross was displayed at the centre of the screen throughout presentation.

IAT stimuli were either positive (e.g., lucky, ecstasy, etc.) or negative words (e.g., vomit, scar etc.) and the letters T and L (presented at the same 4 orientations used in the visual search task). Stimuli were presented just below the centre of the monitor and appeared at the same location on each trial. All stimuli were black on a white background. Letters making up words were 10mm tall and 10mm wide, whilst the letters T and L were 25mm tall and 25mm wide. A full list of the words used appears in Appendix III.

Finally, the STAI (Spielberger, Gorsuch & Lushene, 1970) was used as a subjective measure of anxiety. The STAI consists of two short questionnaires, measuring components of both state and trait anxiety. Each questionnaire consists of 20 items rated on a 4 point scale, where 1 = almost never, 2 = sometimes, 3 = often and 4 = almost always. Hence the minimum possible score for each scale is 20 and the maximum possible score is 80. State items address the current mood of the participant, using short statements such as. "I feel calm", or "I am tense" etc, whilst trait questions are used to measure more enduring personality characteristics rather than current mood, e.g. "I lack self-confidence", "I am a steady person" etc. A full list of these items can be found in Appendix II.

Apparatus

The experiment was conducted on Apple Macintosh Power PC 400 MHz G4 computers with 19" colour monitors. Participants responded with their dominant hand on a Superlab button box (Cedrus Corporation, 1998), interfaced with the computer via a serial port adapter connected to a USB port.

Procedure

Visual Search Task: Upon entering the laboratory, participants were asked to complete The Spielberger State-Trait Anxiety Inventory and then presented with a card on which the letter T was printed at 4 orientations (0°, 90°, 180°, 270°). Participants were asked to rate the letter for valence. Ratings were scored on a 7-point scale, where 1 represents low valence and 7 represents high valence. A second card on which the letter L was printed at the same 4 orientations was subsequently presented, and ratings recorded using the same 7-point scale.

Participants were randomly divided into two groups, and both groups performed exactly the same visual search task. Participants in Group 1 were presented with targets conditioned with negative IAPS pictures, and for Group 2, targets were associated with neutral IAPS pictures. Participants in both groups were shown a sample IAPS image appropriate for their group and informed that target letters would be paired with similar images for the duration of the experiment. This procedure served two functions; firstly conditioning is facilitated if participants are explicitly told the association between the CS and UCS and secondly, it provided participants an easy opportunity to quit the experiment without viewing any more stimuli that might have been distressing.

Participants were asked whether they felt happy to continue with the investigation, and then seated in front of the apparatus and allowed to move the response box to a comfortable position to their left or right depending on their dominant hand. The distance away from the monitor for each participant was approximately 60 cm. Following this, participants began the first block of trials, the results of which were excluded from the data analysis. This served a dual purpose, as both a conditioning phase and a practice run. The initial practice block consisted of 24 trials, and the set size was randomised. In target present trials, a picture (negative or neutral depending upon the participant's group) always followed their response. As the number of target-present trials equalled the number of target-absent trials, a total of 12 IAPS pictures were viewed in the practice run.

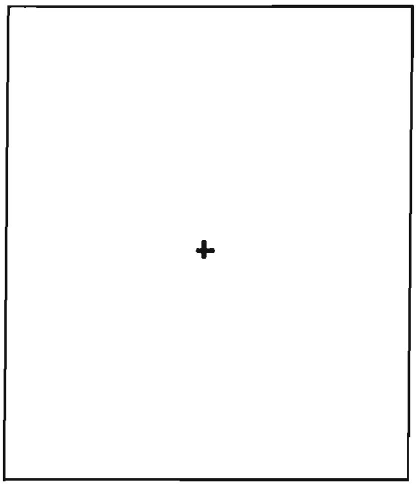

Figure 4.1 Sequence of events per trial. (A) The first event in each trial consisted of the presentation of a fixation cross in the centre of the monitor for 1200ms.

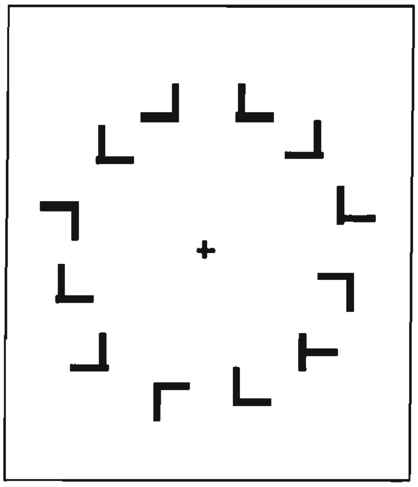

(B) The fixation cross remained visible while the visual search display was presented. There was no time limit on the duration of the search display, which remained visible until the participant responded. Set Size was randomised across 4 sizes (1, 3, 6 and 12 items). In this example (target present) the set size is 12. The target could appear at any of the 12 locations around an imaginary circle. Trials were randomised across target position and trial type (target present or absent). In target absent trials, a fixation cross was displayed as above after the response indicating the start of the next trial.

In target present trials an IAPS picture was presented immediately following the participants' response for 5s. The content of the picture was dependent upon the participants' group. For participants in the neutral group (as in this example), pictures consisted of low arousal/medium valence images, e.g. mushrooms, hairdryers etc (see Appendix I). Images in the negative group were chosen on the basis of high arousal/low valence scores (see also appendix I). If participants missed the target, a 'negative' tone was sounded and a congruent image was displayed for 5s (using the same procedure used for correctly identified target present trials).

Once the practice trial was finished, participants completed 4 blocks of 48 trials, with breaks after every 24 trials. The procedure used was identical to that in the practice run and is illustrated in Figure 4.1. At the beginning of each trial, a fixation cross was displayed in the centre of the screen. Following a stimulus onset asynchrony (SOA) of 1200 ms, a randomised search array appeared around an imaginary circle (see Figure 4.1B). The set size varied between 1, 3, 6 and 12 items, and set size was randomised and balanced across the trials in the 4 blocks. Participants responded (via a key press on the button box) as quickly and accurately as possible whether a target was present or absent. For approximately half the participants, the left hand button designated target present, and the right hand button designated target absent. For the other half of the participants, the mapping was reversed. Stimulus display remained until participants responded. If an incorrect response was given, the display disappeared and feedback was given via a ‘negative’ auditory tone. On target present trials (including trials in which the participant missed the target), an IAPS image congruent to the group condition was displayed. The SOA between the initial fixation cross and IAPS image varied depending upon the speed of the participant's response. However, the display duration of the IAPS image was always 5 seconds. In target absent trials, no image was presented. Instead, the fixation cross was immediately displayed following the participant's response, indicating the start of the next trial.

At the end of the visual search task, participants were again presented with cards showing the letters T and L at each of the 4 orientations and asked to rate them for valence using the same 7-point scale as before.

IAT Procedure: The second half of the experiment consisted of an Implicit Association Test (Greenwald, McGee & Schwartz, 1998) to measure whether targets had been successfully conditioned to carry affective valence. The premise behind the IAT is outlined briefly in the introduction. Stimuli in the IAT are presented one at a time, and are of two types, words (positive and negative) and letters (T and L). Each time a word or letter is presented, it must be classified into one of two categories by pressing the left or right button. The IAT consists of five different phases, or blocks, of trials. In critical phases, stimuli of both types are randomly intermixed. Thus, participants must classify a word or letter as belonging to one of two categories (e.g. T and Good *or* F and Bad and vice-versa). If conditioning has been successful, participants in the negative group (in which the letter T was paired with negative pictures in the earlier part of the experiment) should evidence faster RTs to respond to the letter T when it requires the same response as negative words than when it requires the same response as positive words. However, for participants in the group in which T was paired with neutral pictures, RTs to T and positive words and T and negative words should be roughly equal. Thus the IAT should effectively discriminate valence between the conditioned letters from the visual search task (i.e. T and L's). The structure and layout of the IAT is shown in Table 4.1 and the 5 stages are described over the page.

Sequence	P1	P2	P3	P4	P5
Task Description	Initial attribute discrimination	Associated target-concept discrimination	Initial combined task	Reversed attribute discrimination	Reversed combined task
Task instructions	● Bad Good ●	● T L ●	● T ● Bad L ● Good ●	Bad ● ● Good	● T ● Good L ● Bad ●
Sample stimuli	✓ Abuse Freedom✓ ✓ Filth Laughter✓ ✓ Vomit ✓ Rotten Paradise✓ Peace ✓	L ✓ Γ ✓ ✓ T ✓ ⊥ L ✓ ✓ T ✓ ⊥ L ✓	L ✓ Gentle ✓ ✓ T ✓ Cancer L ✓ ✓ ⊥ ✓ Pollute Happy ✓	✓ Friend ✓ Happy Tragedy✓ Stink ✓ Assault ✓ Hatred ✓ ✓ Rainbow Love ✓	✓ T ✓ Friend L ✓ Tragedy✓ ✓ ⊥ Γ ✓ ✓ Peace Assault ✓

Table 4.1. Schematic description of the IAT. The 5 phases of the IAT are detailed in the above table (columns 1-5). A pair of target concepts and an attribute dimension are introduced in the initial two phases. Categories for each of these discriminations are assigned a left or right key on the button box (indicated by the black circles). These are combined in the third phase and then recombined in the fifth phase, following a reverse response assignment (in the fourth

phase) for the attribute discrimination. Ticks indicate correct responses. Adapted from Greenwald et al (1998) p 1465.

Phase 1: The words Bad and Good were displayed in the top left and right of the monitor respectively. Written instructions detailing the format and correct key responses were displayed on screen at the beginning of this phase until the participant depressed a key on the button box to start the trials. Only words appeared in this phase (T and L were never present). As there was no time limit to the display duration of targets, targets remained visible until a response was given. In the first phase, (attribute discrimination task), participants were instructed to press the left key whenever an unpleasant word (e.g. vomit) was presented. Words appeared in the same location (slightly below centre) throughout each experimental stage and as in all stages, their order of presentation was randomised. Likewise, whenever a target word that could be classified as 'good' (e.g. laughter) was presented, participants were instructed to respond by pressing the right key. In each trial, a single word was presented alone. At the beginning of this stage (and all other stages), verbal and written instructions emphasised the need for participants to respond as quickly and accurately as possible. If an incorrect response was made, an "X" was displayed at the location of the prior target word, and a written instruction was displayed stating that an error had been made and to press any key to resume the experiment. Participants completed 16 trials in this initial stage.

Phase 2: Associated target discrimination. The letters T and L were displayed in the top left and right of the monitor respectively. Participants followed exactly the same procedure as in phase 1, except that target stimuli were the letters T and L rather than positive and negative words. Presentation was randomised so that they appeared at any of the 4 orientations seen in the previous search task. Participants were instructed to press the left key for T at any orientation and the right key for L at any orientation. Again, participants completed 16 trials.

Phase 3: Initial combined task. The word Bad and the letter T and the word Good and the letter L were displayed in the top left and right of the monitor respectively. In this experimental phase, participants responded to categories displayed at the top of the screen as before by pressing the appropriate key. For example, depressing the left key for the letter T *and* unpleasant words and the right key for the letter L *and* pleasant words. Throughout the trials, as in all phases, incorrect responses were followed with an X and a message stating that an incorrect response had been given. As before, the experiment resumed once any key was pressed or 1500ms had passed. Participants completed 1 block of 96 trials in this phase.

Phase 4: Reversed attribute discrimination. The procedure was exactly the same as in phase 1, except that target categories (and their displays) were swapped over, i.e. participants pressed the left key for words that could be classified as "good" and the right key for words that could be considered "bad". The letters T and L did not appear in this phase. As in Phase 1, participants completed 16 trials.

Phase 5: Reversed combined task. In this final phase, letters and categories were recombined, but this time with the reversed order for “good” and “bad” as in Phase 4. The target categories L and bad appeared to the upper right hand side of the monitor, and T and good to the upper left of the monitor. As in phase 3, participants were required to classify stimuli accordingly. Again, as in phase 3, participants completed 1 block of 96 trials.

As the response keys used in the IAT (i.e. press the left button for T and the right button for L) are the same as those used in the original visual search task, a strong implicit association may have already been built between the left key and the letter T and the right key and the letter L. Therefore, the concepts T and L were always presented on the left and right of the screen respectively. As the IAT shows strong order effects, the order of presentation was counterbalanced across participants. Therefore, for 50% of participants, T and Bad (P3) were associated with the same response in the third phase, whilst T and Good (P5) were associated with the same response in the final phase. For the remaining 50% this order was reversed, i.e. P5 (T and Good) was presented before P3 (T and Bad). Therefore, the order of presentation for the IAT was P1, P2, P3, P4 and P5 for half the participants and P4, P2, P5, P1 and P3 for the remainder. Data were compared from phases 3 and 5 so that each participant provided data from a total of 192 trials. These data were analysed by a series of ANOVAs.

4.3. Results

Visual Search Task: Data were analysed using a 4 x 2 x 2 mixed Analysis of Variance (ANOVA). Set Size (sizes 1, 3, 6 and 12) and Target (present or absent) were within subjects factors, and Picture Type (negative vs. neutral pictures) was a between subjects measure. Trials with incorrect responses or RTs <100 ms were removed prior to the analysis. Due to the wide variety of RTs in the larger set sizes in visual search tasks, no upper cut-off point for RTs was used. There was no main effect of Picture Type on RT, *F* (1, 30)=1.59, p> .05. As expected, there was a significant main effect of Target (present vs. absent), *F* (1, 30)= 61.42, p< .001, and an interaction of Set Size x Target, *F* (3, 30)=60.63, p< .001. No other effects or interactions approached significance. As can be seen in Figure 4.2, the search for negative targets was no more efficient than the search for neutral targets. Only two effects reached significance: As the number of distractors increased, there was a corresponding increase in RTs, and participants were roughly twice as fast to decide that a target was present than absent. Both effects are predicted by many current models of visual attention (e.g. Treisman & Gelade, 1980).

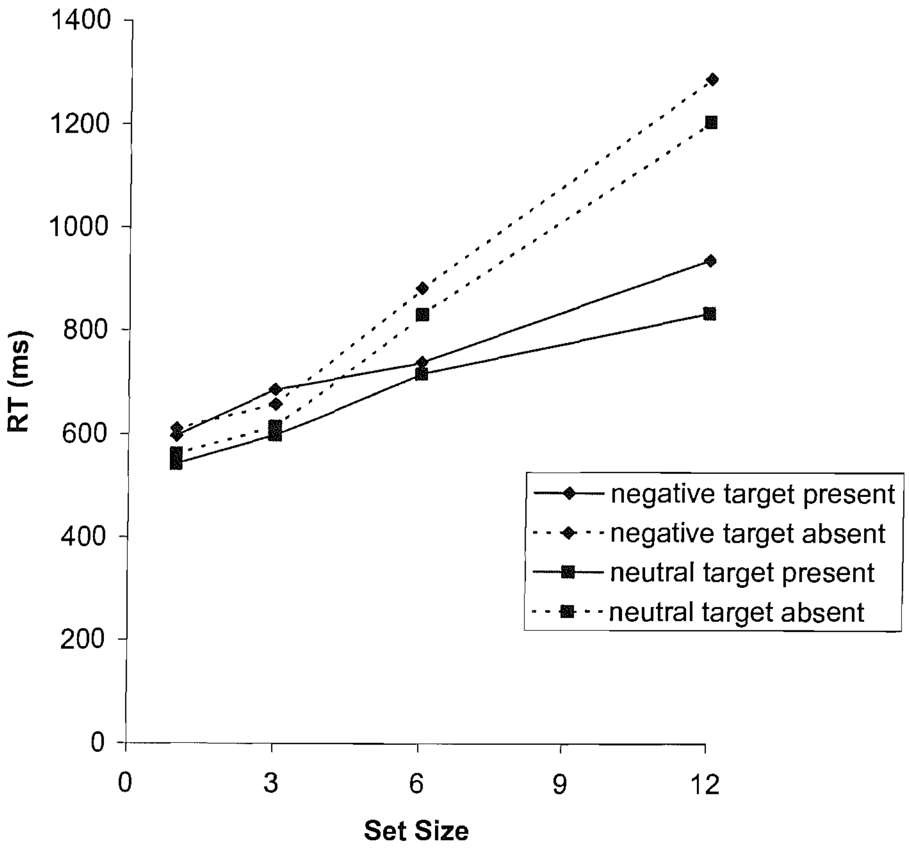

Figure 4.2. Mean visual search RTs for target present (solid lines) and target absent (dashed lines) trials, differentiated by set size and targets associated with negative (diamonds) or neutral (square) pictures.

between subjects factors and Response Pairing (T paired with good or T paired with bad words) was a within subjects factor. Response times >2000ms and <100ms and errors were excluded from data analysis. If the conditioning in the search task was effective, then the response pairing of T with bad words should produce faster responses than T with good words, but only for participants who were presented with negative images in the search task.

First, there was a significant main effect of Picture Type (Negative v Neutral Pictures), *F* (1, 28)= 6.0, p< .05, whereby participants who viewed negative pictures during the search task were significantly slower in their overall RTs during the IAT task. There was also a significant interaction between Response Pairing and Order, *F* (1, 28)= 13.96, p< .001. This interaction may be interpreted as a weakening of conditioning effects during the IAT; i.e., faster responses if congruent pairings are combined in the first data phase of the IAT than during the second data phase. Additionally, an interaction between Response Pairing and Picture Type, demonstrating the effectiveness of the conditioning, showed a trend in the correct direction. While participants viewing negative pictures associated the letter T more readily with the category bad (T-) than the category good (T+), this effect did not reach significance, *F* (1, 28)= 1.086, p>. 05. This pattern of data is illustrated in Figure 4.3.

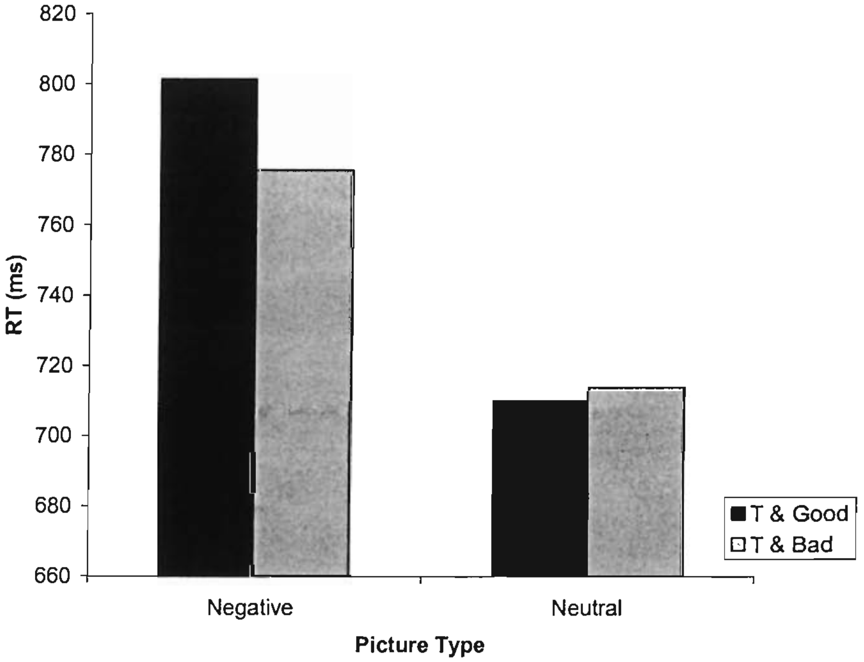

Figure 4.3. Mean IAT RTs by Picture Condition (Neutral Vs Negative Images).

4.4. Discussion

The purpose of Experiment 1was twofold. Firstly it was hypothesised that there may be differences in response times to stimuli conditioned to carry negative valence compared to stimuli conditioned with neutral pictures, in line with data from a number of studies showing more rapid search for threat (e.g. Eastwood, Smilek & Merikle, 2001; Fox et al, 2000; Öhman, Lundqvist & Esteves, 2001). Secondly, and more importantly, it was hypothesised that search slopes may be altered by the negatively conditioned stimuli compared to neutral stimuli (e.g. Öhman, 1997; Öhman, Flykt & Esteves, 2001). In particular, if the search slopes for the negative valence targets showed a flattened rate of search then this would have implications for theories of visual attention, which propose that emotion should have no effect on visual searches for the type of targets used in the present experiment.

The results from the search task in Experiment 1 are clear. Whilst there was a general trend in the visual search task for participants to respond more slowly across the board for all set sizes in both target present and target absent trials in the negative condition, this difference was not significant. More importantly, the search slopes for negative valence stimuli were more or less identical to the search slopes for the neutral stimuli.

Interpreting the search results, however, is difficult because of the results from the IAT. Although a trend was observed in the IAT in which participants responded more rapidly to the pairing of T and "bad" words than T and "good" words when T had been previously paired with negative images in the visual search task, the interaction between T and picture type did not reach significance. As can be seen in Figure 4.3, participants who had previously been given the pairing of T and neutral pictures showed almost identical RTs to T and good words and T and bad words in the IAT. Interestingly, participants' valence ratings on the 7-point scale for T fell from 4.9 before the visual search task to 4.1 after the negative search task, reflecting a modest subjective decrease in valence rating (although this difference was non-significant). Although the difference in mean RTs for the critical phases (3 and 5) in the negative group were over 25 ms compared to almost 0 ms in the neutral group, there was a large amount of error variance. It is impossible to determine whether the lack of significant differences between the groups in the visual search task are due to the fact that attention is not automatically captured at a preattentive level when the stimulus carries a threat (Öhman, Flykt & Esteves, 2001; Öhman & Mineka, 2001), or whether the stimuli were not effectively conditioned. Due to the ambiguous IAT data, a preattentive bias can be neither confirmed nor disproved at this stage.

The obvious question evoked by the data is why the IAT has not worked as expected. Is this due to the fact that the conditioning did not work in the first place, or more an artefact of the exemplar L not being positive? Put another way, as the target (T) in the negative condition was always shown in tandem with a negative image, participants should be more rapid in responding to the exemplar T and bad than the exemplar T and good. However, as the letter L was never paired with a positive stimulus (rather it was not paired with any stimuli at all) it may be much harder for participants to associate L with the concept good. Furthermore, in target present trials in the negative condition, with the exception of set size 1, the letter L was always present with the letter T to serve as a distractor. As a result, it is quite feasible that participants could associate the letter L equally with the concepts good and bad. Indeed, as no positive images were ever shown following the letter L, one could argue that the letter L is equally likely to be viewed as neutral or negative as it is to be viewed positively. Were the letter L to be more readily associated with the concept good, then this may increase the strength of the relationship between T and bad in addition to strengthening the relationship between L and good. Hence, this could increase the overall effectiveness of the IAT.

Participants were randomly assigned to each group and STAI-T scores of 39.2 and 32.9 ($F<1$, non-significant) for negative and neutral groups respectively indicate mean anxiety scores within the cut off bracket usually used to delineate low anxiety groups. Indeed, trait anxiety measurements were used solely to ensure that there were no chance differences in anxiety between the groups. Thus, the results from Experiment 1 are assumed to reflect the findings one would expect from a low anxious group in the general population (were the conditioning to prove successful) and therefore cannot be equated with studies using clinically anxious participants.

The search data show a non-significant trend for participants in the negative picture group to be slower overall for both target present and target absent trials compared to participants in the neutral group. This pattern of results corresponds with findings observed in the IAT, whereby participants who viewed negative pictures in the visual search task were significantly slower in all phases of the IAT. Due to the graphic content of the pictures in the visual search task, participants may have dwelled on the images after they disappeared, and this rumination could have been responsible for the interference with task performance. Bradley et. al (1992) found that when participants are first exposed to new pictures, RTs to probes are significantly slower for emotionally arousing pictures in comparison to neutral or calm pictures, possibly due to the allocation of more attentional resources at encoding. Furthermore, studies in which participants are shown negative and neutral images under free viewing (i.e. the participant is allowed to view each image for as long as they wish), demonstrate that participants view negative images as long as pleasant images (Bradley et. al, 1992). Thus, although participants viewed each picture set for an equal amount of time, the disturbing content of the images in the negative group is likely to have caused more interference and hence, slower RTs in the IAT. Similar effects have been evidenced in other studies using affective pictures, with negative and high arousal pictures producing more interference than pleasant and low arousal pictures (Tipples & Sharma, 2000).

Without either clear changes in slope function or evidence of effective stimuli conditioning, these experiments cannot address the question of whether attention can be captured preattentively by complex visual stimuli associated with threat. To address these issues, a second experiment using the same basic paradigm but with subtle yet important changes was run. Where Experiment 1 employed a between subjects design, in which trait anxiety was unaccounted for, Experiment 2 used a within subjects design, whereby each participant acted as his/her own control, using 3 separate groups differentiated by their trait anxiety levels (again, measured using the STAI).

Chapter 5: Experiment 2

5.1. Introduction

As previously discussed, due to the lack of significant findings in the IAT, little can be inferred about whether previously emotionally neutral stimuli (i.e. letters) can be made to carry affective valence and whether this in turn has any effect upon visual attention. Experiment 2 was designed to circumvent some of the problems inherent in Experiment 1. The most important change involved the use of a within subjects design, whereby each participant acted as his or her own control. Secondly, participants were divided into three groups (low, medium and high anxiety) on the basis of their STAI trait scores. Finally, increases were made to the number of trials in the IAT in order to give participants more practice in changing concepts and to enable any associations to be more clearly seen.

Experiment 2A

5.2. Method

Participants

46 students (20 male, 26 female) from The University of Southampton participated in the experiment in return for course credits or £5 payment. Participant's age ranged from 19 to 45 years. Mean age was 25.46 years (SD =5.78). All participants gave written informed consent and reported normal or corrected to normal vision, including normal colour perception. Participants' trait anxiety scores were measured using the Spielberger State-Trait Anxiety Inventory (STAI; Spielberger, Gorsuch & Lushene, 1970) and participants were assigned to one of 3 groups on the basis of these scores. Scores may range from 20 to 80; cut off points were as follows; 40 and below for low anxiety, 41 to 49 for medium anxiety, and 50 and above for high anxiety. 10 male and 6 female participants comprised the low anxiety group (STAI mean = 28.2 state, 33.1 trait), 7 male, 8 female (STAI mean = 37.13 state, 44.47 trait) the medium group, and 3 male, 12 female (STAI = 48.87 state, 56.73 trait) the high anxiety group. There were no significant differences between groups in age and sex distribution.

Design

A mixed design was used. The between subjects factor was the participants' trait anxiety level. Each participant was required to search for two targets, a T or an F (the within subjects factor).

For half the participants in each group the letter T was paired with negative valence images taken from the International Affective Picture System (IAPS; Lang, Bradley & Cuthbert, 1997) and the letter F with neutral valence images. The remaining half experienced the reversed pairing (i.e. T was neutral and F negative). On each trial one target was always present. All conditions were counterbalanced across each group so that 50% of participants were required to press the left-hand button of the button box if a T was present and the right-hand button if an F was present. The remaining half used the opposite mappings.

Stimuli

Stimuli used in the visual search task were arrays of target letters (T and F) and distractors, a T/F hybrid (see Shiffrin & Gardner, 1972; Gardner, 1973; see figure 5.1). As in Experiment 1, both targets and distractors were presented at one of four orientations; (0°, 90°, 180°, 270°) and their presentation was randomised. In a single trial there could only be one target present. However, the number of distractors varied across trials (either zero, two, five or eleven). Target and distractor letters were black. Each letter was 10 mm tall and 10 mm wide. Throughout all trials the background was white. Target and distractor letters were spaced around an imaginary circle with a diameter of 100 mm around fixation, as in Experiment 1. Throughout the course of the experiment, target and distractors could appear at any of twelve equally spaced locations around this imaginary circle. A black fixation cross appeared at the centre of the screen both before and throughout stimulus presentation. Neutral and negative pictures were the same as those used in Experiment 1.

Apparatus

The experiment was conducted with the same computers and response boxes used in Experiment 1.

Procedure

The procedure was the same as in Experiment 1 with a few subtle changes. As before, participants completed the State Trait Anxiety Inventory and rated the letters T and F on a 7-point scale. Participants were then given written and verbal instructions and informed that a target would always be present in each display. For 50% of participants the letter T was paired with negative images and the letter F with neutral images. The remaining 50% were given the opposite pairings (i.e. T was paired with neutral images and F was paired with negative images). As before, participants were informed of the association between target and picture type and shown sample images. Participants were then given the opportunity to withdraw from the study (2 cases) and providing they felt comfortable to continue, all participants underwent the revised visual search task as before.

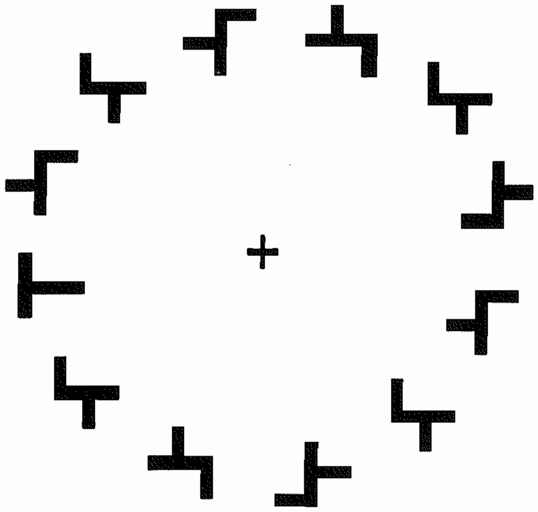

Figure 5.1. Example stimuli used in the visual search task. The target was either the letter T or the letter F. Set Size varied between 1, 3, 6 and 12 items. This example shows T as the target with 11 distractors.

To facilitate participants, a template was placed over the response box showing the appropriate letter for each response key. Following completion of the visual search task, participants rated the letters again, and then performed the IAT as before. Each participant was given one of four versions of the IAT depending on whether they had the letter T or F to the left or right of the button box. For half the participants given the pairing of T with the left hand button, the first data phase consisted of T paired with good words, followed by T paired with bad words in the second data phase. For the remainder, the order was reversed (i.e. T was paired with bad words for the first data phase and good words for the second data phase). Similarly, there were two conditions for participants given the pairing of F with the right hand response key, with half the participants given F and good words in the first data phase and F and bad words in the second data phase. For the remainder, this order was reversed. This counterbalancing procedure was used to eliminate order effects.

As in Experiment 1, to contend with any prior mapping learnt over the course of the experiment, participants who pressed the left-hand button for T and the right-hand button for F always had these exemplars to the left and right of the screen respectively in the IAT. Similarly, participants with reversed mapping had exemplars congruent with their button presses. The only other difference in the IAT to Experiment 1 was an increase in the number of versions of the IAT (from 2 to 4 versions) to contend with the counterbalancing of target side, and an increase in the number of trials in each stage. The procedure was as follows:

Phase 1 consisted of the concepts Bad and Good (or vice versa) presented in the top left and right of the monitor respectively. 32 trials were completed (16 positive and 16 negative words).

Phase 2: Associated target discrimination. Participants followed exactly the same procedure as in phase 2 of Experiment 1. Throughout all the trials in this phase, the letters T and F (or F and T) were presented in the top left and right of the screen respectively. Letters were 25 mm tall and 25 mm wide. On each trial, a single target letter (either T or F) would appear just below the centre of the monitor at any of the four orientations used in the previous visual search task. (No words appeared in this phase). Only one target letter appeared on each trial (although the letter headings remained visible at the top of the monitor throughout the phase to remind participants which letter corresponded with which side). Participants completed 32 trials (16 T and 16 F).

Phase 3: Initial combined task. In this critical phase the target concepts were T and F (or vice versa) *and* Good or Bad (or vice versa). Participants completed 128 trials (32 F, 32 T, 32 Positive Words and 32 Negative words). A full list of the words appears in Appendix IV.

Phase 4: Reversed attribute discrimination. Again this was exactly the same as Phase 1 except for a reversal of the concepts. Participants again completed 32 trials.

Phase 5: Reversed combined task. The letter responses remained as they were in Phase 3, but the switching of the responses to the words in Phase 4 was maintained in this phase, so that now the pairing of letters with word categories was different than it had been in Phase 3. Participants completed 128 trials using the same stimuli as Phase 3.

5.3. Results

Visual Search: Three-way repeated measures ANOVAs with the between subject variable Anxiety Level (low, medium or high) and the within subject variables Target Valence (negative vs. neutral) and Set Size (1, 3, 6 or 12) were run to determine whether negatively paired CS letters were detected more rapidly or with fewer errors in the visual search task than CS letters given the neutral pairing. RTs <100 ms and errors were removed prior to analysis. Due to the large variation in RTs at larger set sizes, there was no upper cut off point.

The main RT effects of the visual search task are summarised in Table 5.1. Search slopes were very inefficient for all groups, with target-trial rates of ~60ms/item for both negative and neutral associated targets. There were no significant main effects of Target Valence, or Anxiety Level. In addition, there were no significant interactions between Set Size and Anxiety Level, Target Valence and Set Size, Anxiety Level and Target Valence, or Anxiety Level x Target Valence x Set Size, (all *F*'s<1.8). There was however, a highly significant main effect of Set Size, *F* (3, 129) = 378.59, p< .001, showing an increase in RTs with set size.

	Set Size								Slope
	1	SD	3	SD	6	SD	12	SD	(ms/item)
Low anxiety level									
Negative target	612.57	(133.29)	765.32	(166.36)	939.30	(222.30)	1210.13	(338.03)	53.13
Neutral target	615.08	(165.77)	775.26	(187.02)	966.32	(272.73)	1281.93	(385.48)	59.56
Medium anxiety level									
Negative target	622.67	(161.67)	751.77	(179.23)	936.66	(184.75)	1306.59	(325.26)	62.03
Neutral target	620.49	(128.70)	806.62	(189.73)	991.68	(275.80)	1331.83	(390.85)	62.96
High anxiety level									
Negative target	662.96	(188.93)	768.82	(184.75)	950.65	(276.20)	1280.01	(444.26)	56.38
Neutral target	551.17	(92.51)	707.99	(124.94)	864.76	(186.56)	1166.03	(352.42)	54.51

Table 5.1. Mean RTs to letters associated with negative or neutral pictures by Set Size and Anxiety Level. All figures are ms. Figures in parentheses are standard deviations.

With respect to error rates, there was a main effect of Anxiety Level, *F* (2, 43) = 3.26, p< .05, with the low anxiety group evidencing the lowest error rates. Also, a significant interaction between Set Size and Anxiety Level, *F* (3, 129) = 69.2, p< .001, indicated that the increase of error rates with corresponding increases in set size was only evident in the high trait anxiety group (see figure 5.2). However, error rates were low across all groups and set sizes, with a maximum error rate of just above 7% for the high anxiety group in the largest set size (11 distractors).

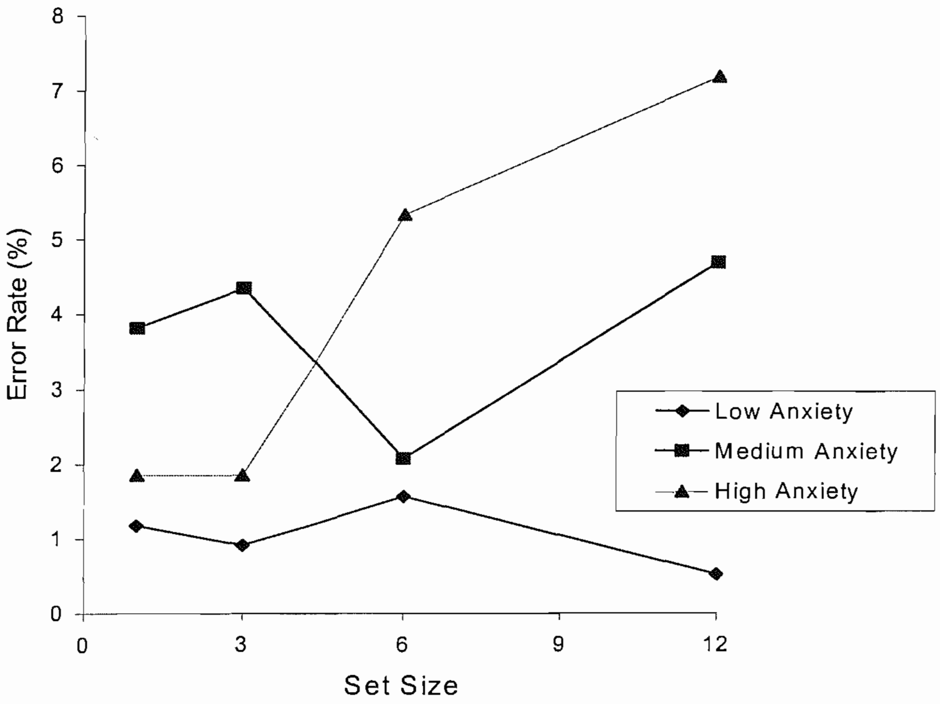

Figure 5.2. Error rates to conditioned targets by Set Size and Anxiety level.

<u>IAT</u>: As in Experiment 1, a 3 Way ANOVA was performed on the IAT data to measure the effectiveness of the conditioning. The within subjects factor, Response Congruity, consisted of two levels, Congruent and Incongruent Response Pairing (i.e. if the letter F had been paired with negative pictures in the visual search task, then the pairing of F and 'bad' words was considered to be congruent mapping. Similarly, if F was paired with 'good' words following a negative association in the visual search task, this was termed incongruent). There were two between subject's factors, each with two levels, F or T (i.e. whether the negative target was F or T) and Order (congruent first or incongruent first). In this analysis, the effectiveness of conditioning is tested by the main effect of Response Congruity. This factor was significant, *F* (1, 41) = 5.60, p< .05. In other words, regardless of whether T or F was paired with negative pictures in the visual search task, participants were faster to associate the negative target with

bad words than good words in the IAT task (734 ms (SD = 121) for negative target and bad words vs. 771 ms (SD + 113) for negative target and good words). In addition, there was an interaction of Response Congruity x Negative Target which, although non-significant, evidenced a strong trend, *F* (1, 41) = 3.78, p= .059. Participants were faster to classify the letter F with words of negative affect and the letter T with words of positive affect when they had been previously paired with negative and neutral targets respectively in the visual search task. When this pairing was reversed (i.e. when T was paired with negative pictures and F with neutral pictures in the visual search task), participants were slower to associate T with negative words and F with positive words compared to the congruent mapping. Participants apparently came into the lab associating T with positive and F with negative before undertaking any of the experimental tasks.

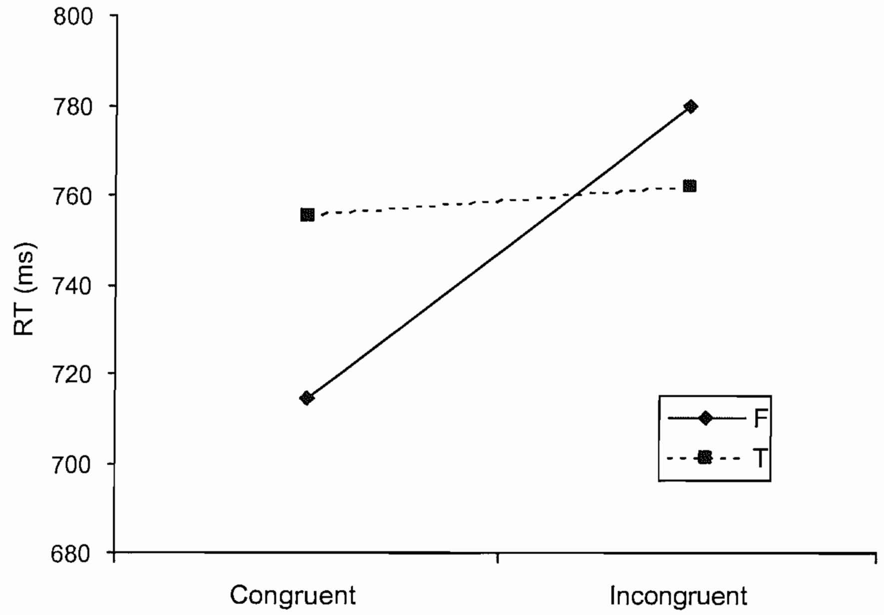

Figure5.3. Interaction plot of Congruent and Incongruent Response Pairing by Negative Target. Congruent refers to the pairing of F with bad words and T with good words. Incongruent is the pairing of F and good words and T with bad words.

Overall the IAT demonstrated effective conditioning of target letters, evidenced by significantly faster RTs to congruent pairings of target letters and word valences. As expected, participants were slower in their RTs to incongruent exemplars. A graph of the interaction of response congruity by negative target can be seen in Figure 5.3.

This effect was complicated slightly in that participants apparently found it easier to associate F with bad and T with good than to associate T with bad and F with good. However, this interaction did not achieve significance, and even if it had, it would not detract from the main result of the IAT, which shows that the conditioning was effective. Subjective ratings of T and F on the 7-point scale using change scores (taking the difference between the pre and post test scores for neutral and negative conditioned targets) were tested with a 2 x 3 mixed ANOVA.

Change consisted of 2 levels (Negative and Neutral conditioned targets), while Anxiety had 3 levels (low, medium and high). There was a main effect of Change, *F* (1, 43) = 3.82, p= .057 on the cusp of significance. However, there was no significant interaction between Change x Anxiety, *F* (2, 43) = 1.13, p> .05, i.e. all participants evidenced a subjective change in their affective valence towards letters paired with negative targets in the visual search task, with no significant difference in these ratings for participants across the different trait anxiety groups (see table 5.2).

	Negative Target	Neutral Target
Low anxiety level		
Pre-Conditioning	4.81 (1.11)	4.56 (0.96)
Post-Conditioning	4.44 (1.21)	4.63 (0.96)
Change	0.38 (0.80)	0.06 (0.25)
Medium anxiety level		
Pre-Conditioning	4.40 (0.63)	4.07 (1.10)
Post-Conditioning	4.26 (0.80)	4.20 (0.86)
Change	0.13 (0.99)	0.13 (0.83)
High anxiety level		
Pre-Conditioning	4.33 (0.82)	3.93 (0.80)
Post-Conditioning	3.63 (1.17)	4.10 (1.49)
Change	0.70 (1.00)	0.17 (1.06)

Table 5.2. Valence ratings of conditioned letters by trait anxiety level

Correlations: It is possible that more efficient visual searches for threat targets were not observed due to individual differences in conditioning (i.e. some participants evidenced little/no conditioning, whereas for others, the conditioning was very strong). To test for this possibility, correlations between the strength of conditioning (measured by the IAT) and visual search RTs were calculated. The rationale is as follows: if conditioning was successful, then subtracting the congruent RT from the incongruent RT in the IAT should give a positive value.

Similarly, if threat targets are detected more efficiently than neutral targets, then subtracting negative from neutral RTs should give a positive value. Therefore, if the visual search task is dependent upon the strength of conditioning, and the lack of significance is due to some participants being more effectively conditioned than others, the strength of conditioning should correlate with more rapid visual search RTs. As there were no significant correlations for all three levels of trait anxiety combined (r= -.024, n=46, p> .05), or by low (r=. 085, n, 16, p>.05) medium (r= -.059, n, 15, p >.05) and high anxiety (r=-.047, n=15, p>.05) this explanation can be ruled out

5.4. Discussion

Experiment 2a was designed to address two issues. Firstly, groups were differentiated on the basis of their trait anxiety scores on the STAI, in order to address specifically the impact of anxiety on visual search. Secondly, a within subjects design was employed so that each participant acted as their own control, and also in order to strengthen associations within the IAT. To this extent the experiment was successful, as in contrast to the results of Experiment 1, there was a significant main effect of Response Congruity in the IAT. However, as in Experiment 1, there were no significant effects of target valence on visual search.

In addition, there was no main effect of anxiety on visual search, although there was a slightly more efficient rate of search in the low and high anxiety groups compared to the medium group. With respect to error rates, there was a significant main effect of anxiety and an interaction effect of set size and anxiety. The maximum error rate was just above 7% in the largest set size. However, although this effect was significant, it is important to note that for all participants error rates were low and within acceptable parameters for this type of attention task. As expected, error rates increased with set size, in keeping with previous visual search experiments (e.g. Wolfe, 1994). It is possible that the lower error rates in the smaller set sizes are attributable to the lower density of items producing less bottom up activation (Wolfe, 1994).

Although error rates were low for all participants, the highest error rates were evidenced in the medium and high trait anxiety groups, possibly because participants in these groups were more sensitive to the graphic nature of the images. If so, this may explain the uniform error rates for both the negative and neutral targets, as participants in these groups could experience stronger task interference due to rumination of the negative images even when they were not present in the current trial. However, it is noteworthy that there was no significant difference in subjective ratings of the CS on the 7-point valence scales as evidenced by the non-significant interaction between change scores and anxiety.

As there was no "pop-out" effect (Treisman & Gelade, 1980), as evidenced by the search slopes, the target letters cannot be said to have grabbed attention at a preattentive level. Although one cannot unequivocally dismiss the occurrence of automatic, preattentive processing of threat stimuli, the data do not support preattentive selection of the threat stimuli in this experiment.

More important than error rates in the visual search task however, was the main effect of response congruity in the IAT. Put simply, when participants had a negative image paired with a target letter (T or F counterbalanced across participants) in the search task, they were faster to respond in the IAT when that letter was paired with negative words. Similarly, whenever a letter was paired with neutral images in the search task, participants responded quicker in the IAT when that letter was paired with positive words. Participants were significantly faster to associate congruent (conditioned) exemplars than incongruent exemplars, indicating that conditioning was successful in this experiment.

However, the IAT data may be more ambiguous than they initially appear. While it makes good experimental sense to have the target letters in the IAT presented to the same side as the button presses used in the original visual search task, this may have obscured the conditioning results. For example, if the participant pressed the left hand button whenever T was present and the right whenever F was present in the visual search task, then T and F would always appear to the left and right of the screen respectively in the IAT. While this prevents problems due to existing mapping formed during the visual search task, i.e. participants may have learnt to associate T with left and F with right (and vice versa), this raises ambiguity as to what is actually being conditioned. Put another way, does the main effect show that target letters have been successfully conditioned, or have participants simply been conditioned to associate the left button with bad and the right button with good? As a result of the prior mapping that may already exist from the first part of the experiment, one cannot separate target letters from buttons. Thus, one cannot with certainty know whether letters did actually come to elicit affective valence or whether the effect of congruence was simply due to an association with the response buttons. This is a potentially serious issue that needs to be redressed in the next experiment, as it leaves ambiguity over the interpretation of the findings.

Although the IAT in Experiment 1 evidenced a trend in the right direction, it did not reach significance. One may therefore ask why Experiment 2a was more successful (save for the issues raised above). Whereas Experiment 1 employed a between subjects design, Experiment 2a adopted a within subjects design. Therefore for each participant, one target letter was paired with neutral images and another with negative images, (counterbalanced across conditions). As a result, affective valence was transferred to each letter, with the neutrally paired target becoming more subjectively and objectively liked, whilst the negatively paired target became more disliked. As a result, the design overcame a possible problem inherent in Experiment 1, in which only one letter was paired with images, raising the possibility of "contamination". That is to say, whenever a T was displayed in the visual search task in Experiment 1, a letter L was always also present (with the exception of set size 1). Therefore, one cannot discount the possibility that the letter L also carried negative affective valence by association in the worse case, or little or no positive valence in the best. The within subjects design used in the present experiment is more likely to eliminate these problems due to the differential associations between the letters for reasons described above.

A second possibly important difference between the two experiments was in the number of trials in each phase. In Experiment 2a the number of critical trials was increased from 96 to 128 trials in each of Phases 3 and 5. Similarly, Phases 1, 2 and 4 each had the number of trials doubled, from 16 to 32 to facilitate changing concepts. It is likely that increasing the number of trials lessens order effects and allows differences in valence to be seen more clearly.

A second unexpected finding arose when analysing the IAT data, whereby participants seemed to be much faster to respond to the exemplars F and bad when F had previously been paired with negative images than T and bad when T had been previously paired with the same negative images in the visual search task (see Figure 5.3). One possibility for this may be that participants have a prior pre-existing association between the letter F and bad and the letter T and good (although this association appeared weaker for the latter category). One can only speculate as to the cause of this. However, it is feasible that participants have a strong affinity for F and bad due to the association of F and false. For example, in questionnaires, people are often required to respond T for true and F for false. This association would also explain the greater affinity for T and good, although it is unclear why this should be weaker than F and bad. However, it is possible that the stronger affinity for F and bad may to be attributable to the association with F and the word "fuck", a highly negative (and commonly used) word. As such, this would explain how F and negative is more easily associative when paired with negative images as it carries a "double whammy" effect. In addition, whilst the letter T became positively evaluated when paired with neutral images, it is important to note that the images were not in fact positive and their evaluation was only positive because of the contrast with the negative picture set. This interpretation may explain the facilitation of F and bad compared to T and good when paired with congruent pictures. Although this interaction did not reach significance, a clear trend was evidenced. To test the degree of interference as a result of the targets used, a second experiment (2b) was run. However, it is important to note that regardless of any possible prior association, participants in Experiment 2a demonstrated a significant main effect of congruence and that the negative target interaction was non-significant. Therefore, even if an in built bias towards T and good and more particularly, F and bad already exists, the strength of conditioning was strong enough to overturn this prior association.

5.5. Experiment 2b

To test any prior associations between the letters T and F, a group of participants who had not taken part in any previous studies were tested on the IAT without completing the conditioned visual search. The protocol used was the same as that used in the IAT procedure in Experiment 2a.

5.6. Method

Participants

16 Students (4 male, 12 female) from The University of Southampton participated in this experiment. Participants were recruited from the subject pool fulfilling course requirements. Participant age ranged from 18 to 47 years. Mean age was 24.12 years (SD =8.62). All participants gave written informed consent and reported normal or corrected to normal vision.

Design

An independent design was used in which each participant carried out one of the 4 IAT versions used in the previous experiment in order to balance order and hand effects across subjects.

Stimuli

Stimuli used were the same as for the IAT in Experiment 2a.

Apparatus

The apparatus used was the same as in Experiment 2a.

Procedure

The procedure used was the same as that of the second half of Experiment 2a. After completing consent forms, participants were assigned to one of the 4 versions of the IAT. There were 4 participants in each condition. No participants had completed either Experiment 1 or Experiment 2a. With the exception of the omission of STAI-T questionnaires, the procedure used was identical to the second stage of Experiment 2a.

5.7. Results

As previously, mean RTs from Phases 3 and 5 of the IAT were entered into an ANOVA to examine whether there was a main effect of Response Congruity, i.e. whether participants would be faster to associate T with good and F with bad. The term congruence is used in a slightly different way in this experiment. In Experiment 2a congruence referred to the relationship between word valence and targets paired with neutral or negative pictures in the search task, whilst in this experiment, the congruence is between word valences and the positive associations to T and the negative associations to F that participants may have had before the experiment began. The pairing of T with good words and F with bad words is considered a congruent response pairing. Similarly, the pairing of T with bad words and F with good words is referred to as incongruent. RTs <100 ms and >2000 ms and errors were excluded from the analysis. Although there was no main effect of Response Congruity, *F* (1, 14) = 2.8, p> .05, there was a significant interaction effect between Response Congruity and Order, *F* (1, 14) = 7.3, p< .01, (see Figure 5.4).

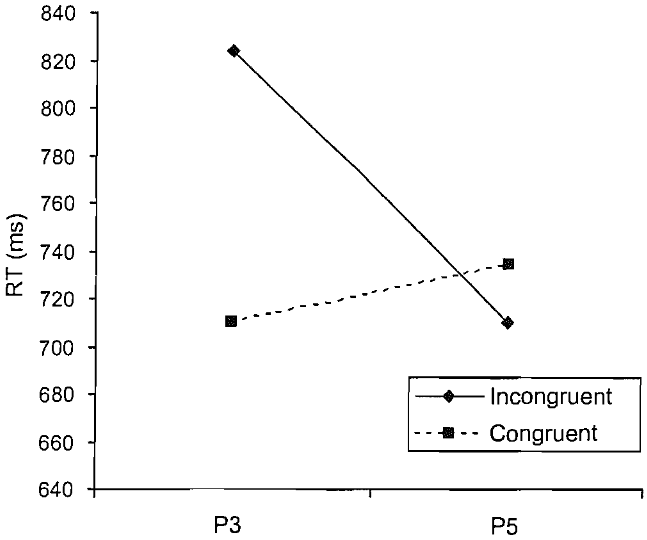

Figure 5.4. Interaction Plot of Response Congruity versus Order Effects. For participants who receive the incongruent response mapping in the first block (P3), responses are slow. However, participants who do not receive the incongruent pairing until the second block (P5) produce responses about as fast as those from congruent pairings.

5.8. Discussion

The results from Experiment 2b show that for congruent exemplars (i.e. T and Good and F and Bad) participants were quick to make associations in the first critical data phase of the IAT (phase 3). However, when the congruent pairing occurred in the fifth phase this advantage was somewhat offset, as participants had to "unlearn" the associations made previously in phase 3. For the incongruent pairings, participants were considerably slower in phase 3, whilst in the fifth phase their mean response times were roughly the same as those of the congruent exemplars in the 3rd phase. These data strongly suggest that participants did indeed have a pre-existing bias to associate T with good and more strongly still, F with bad, in the 3rd phase of the IAT. While these data intimate that participants began Experiment 2a with pre-existing associations already in firmly in place, the conditioning that associated the targets with affective images was strong enough to overturn these associations. As the neutral images came to be regarded as positive as a result of evaluative conditioning and participants were clearly able to unlearn any prior associations, it appears that the negative IAPS images are effective in eliciting strong responses of negative affect.

5.9. General Discussion

The results from Experiments 2a and 2b are important for a number of reasons. Firstly, the data offer no support for the suggestion that attention to conditioned threat stimuli can be captured at a very early stage of visual processing, independent of focused attention, as suggested by recent clinical experiments (e.g. Öhman, Flykt & Esteves, 2001). Secondly, there was no evidence at the preattentive level of visual processing to support the model of Williams et al (1988, 1997) that individuals high in trait anxiety show an enduring tendency to orient towards threat, whilst individuals low in trait anxiety show a characteristic pattern of avoidance. Indeed, although non-significant, the rate of search for the low trait anxiety group was faster for negative targets than neutral targets, whilst the high trait anxiety group evidenced an opposite pattern of results. Finally, there was no evidence at the preattentive level to support the cognitive-motivational account of anxiety advocated by Mogg & Bradley (1998), which claims that stimuli high in threat value preferentially capture attention in *all* individuals, irrespective of their level of anxiety. As IAPS images are of sufficient intensity to be classed as high in threat value (Mogg & Bradley, 2000), the lack of an effect should not be attributed to non-threatening stimuli. Note that the results presented here could still be consistent with the theories by Williams et al or by Mogg and Bradley, if we assume that the attentional biases that they describe do not arise until after stimuli have been attended enough for them to be at least partially identified. In a visual search task designed to measure preattentive processing, however, there is no evidence that these attentional biases play a role.

Although the results from Experiment 2a are contrary to recent experimental clinical findings, the data are very much in keeping with a voluminous body of research in visual search, which advocates that searches for targets sharing common features with distractor objects should be inefficient (e.g. Cave, 1999; Treisman & Gelade, 1980; Wolfe, Cave & Franzel, 1989; Wolfe, 1994; Wolfe, 1998a). Furthermore, the data are in accordance with recent cognitive paradigms that have examined visual search for threat stimuli. For example, using a similar visual search paradigm to that employed by Öhman et al (2001) Tipples, Young, Quinlan, Broks & Ellis (2002) found no differences in participants' reaction times to detect threatening animals, pleasant animals and fruit. As such, the notion of a preattentive bias towards threat (or indeed, an attentive bias) has not been supported. It is indeed possible that, as discussed in Chapter 3, the methodology adopted may play a significant contribution to the results obtained. Whilst the dot probe paradigm is particularly well suited to reveal attentional biases, it does not indicate the processing stage at which this bias occurs, as the temporal interval used (usually 500 ms) allows time for attention to shift from one location to another. Hence, any biases in attention cannot be assumed to be biases in the subject's initial engagement of attention; rather, these may be more a reflection of increased dwell times and difficulties in shifting attention (Fox, Russo, Bowles & Dutton, 2001). To this end, the visual search paradigm is particularly informative. However, from the current data, one can only conclude that, as in the Tipples et al (2002) study, there is no evidence of preattentive, or indeed, attentive biases in the general population towards threatening stimuli.

It is of course entirely possible that attentive (and preattentive) biases towards FR stimuli are a very real phenomenon, and the reason that they were not evidenced in this study was because of ineffective conditioning of the target (T or F) stimuli. However, this possibility is doubtful for two reasons. Firstly there was a change in the subjective valence ratings of the target stimuli (which although non-significant, were on the cusp) and secondly, and more importantly, because of the significant findings in the IAT. Furthermore, while the target letter selection of T and F resulted in unexpected problems due to prior mapping (as evidenced in Experiment 2b), the conditioning proved to be effective enough to overturn this.

Potentially more problematic however, is the interpretation of the conditioning, or put another way, what is actually being measured. It is possible that the significant result in the IAT is a misinterpretation of what is actually being measured by the conditioning. In other words, the significant findings of the IAT may be attributable to conditioning the response buttons rather than the visual search targets themselves. Hence, because the target letters in the IAT were always presented on the same side as the congruent targets in the visual search task, the significance of the conditioning in the IAT could be the result of participants associating the left button with good and the right button with bad and vice-versa, rather than the transfer of affective valence to the targets.

This explanation is unlikely, however, because of the problems with prior mapping (evidenced in Experiment 2b). If one is to advocate that participants were conditioned to the buttons rather than the target (letter) stimuli, then the faster responses to F and bad and T and good when paired with negative and neutral images respectively should not occur. Put another way, if the buttons had been conditioned then there should be no difference between T and F and congruence, as the button and not the letter should be carrying the conditioning. However, if the letters were the items that had been successfully conditioned, then this would explain the difficulty encountered by participants as they had to overturn the pre-existing mapping to the conditioned target, i.e. the *letter* (as evidenced in Experiment 2b). Hence, the significant findings in the conditioning measure (i.e. the IAT) can be assumed to reflect conditioning to the letters T and F rather than the buttons. In order to redress any ambiguity, however, this issue will be dealt with more fully in the next experiment.

In conclusion, the data from Experiments 1 and 2 support current models of visual attention and offer no support for the notion of preattentive detection of threat stimuli (at least in the general population). Although Experiment 1 was ambiguous in its findings due to a lack of significant results in the IAT, the present experiment offers more compelling evidence to validate contemporary models of visual attention. However, in order to further extrapolate the mechanisms of visual attention to biologically salient threat, a further experiment based on the present paradigm was run with a number of important changes. These changes and the rationale behind them are described in more detail in the following chapter.

Chapter 6: Experiment 3

6.1. Introduction

In Experiment 2, there was no evidence of enhanced detection of threat stimuli compared to neutral targets. Furthermore, the efficiency of the rate of search was virtually identical across all levels of trait anxiety. As such, the combined findings from Experiments 1 and 2 suggest that conditioned threat stimuli do not capture attention any more effectively than neutral stimuli in sub-clinical populations, irrespective of the individuals' levels of trait anxiety. Although the IAT failed to achieve significance in Experiment 1, a strong trend in the predicted direction was evidenced. Following amendments to the experimental design in the second study, the IAT reached significance, thus demonstrating effective conditioning of targets. While these data offer evidence against the notion that threat attracts preferential capture of visual attention, an alternative possibility for the findings may be that the UCS was not perceived as threatening.

This alternative interpretation, that the UCS lacked a strong enough threat value to evidence attentional biases, is not consistent with the IAT data showing that the stimuli used were effective in eliciting negative valence. Despite this evidence, there are other stimuli that might be more effective. The idea that we, as a species, are 'hardwired' to fear particular stimuli (i.e. snakes, spiders and angry faces) is used in support of phylogenetic explanations of fear learning (e.g., Mineka & Öhman, 2002; Öhman & Mineka, 2001). Certainly, the amygdala is particularly sensitive to angry faces (Morris et al, 1996). As such, if clinical models of anxiety, with their emphasis on the role of attentional biases in relation to the mediation and aetiology of anxiety, are correct (e.g. Williams et al, 1988; 1997), one would expect to see strong evidence of attentional biases towards 'inherent' FR stimuli such as snakes and spiders in individuals with high levels of snake and spider fear. In order to explore these issues, a paradigm similar to that used in Experiments 1 and 2 was employed. However, substantial changes were made to the target selection and response criteria to addresses ambiguities in the interpretation of the IAT data discussed in the previous chapter.

Experiment 3

6.2. Method

Participants

250 students from The University of Southampton were screened for snake and spider fear using validated snake and spider fear questionnaires (Klorman, Weerts, Hastings, Melamed & Lang, 1974). From the screening procedure, 63 participants were selected for inclusion in the study, separated into 4 groups: high spider fear (14 males, 2 females, mean age 25.5 years (SD=9.03), SPQ=20.63, SD=3.32), low spider fear (4 males and 12 females, mean age 23.25 years, SD=6.44; SPQ=2, SD=1.46), high snake fear (2 males and 13 females, mean age 23.53 years, SD=8.03; SNAQ mean =20.4, SD=3.5), and low snake fear (1 male and 15 females, mean age, 23.19 years, SD =7.59; SNAQ mean =1.81, SD=1.17). Participant age ranged from 18 to 50 years. All participants gave written informed consent and reported normal or corrected to normal vision, including normal colour perception. Selection criteria for the study were based upon participants scoring in the top and bottom 10% for snake and spider fears. Only participants with reliable SNAQ and SPQ scores (confirmed in a second assessment) were included in the study. As expected, the high snake fear group evidenced significantly higher scores on the SNAQ than the low snake fear group, t (14)=18.8, p<. 001. Similarly, the high fear spider group scored higher than the low fear spider group on the SPQ, t (15)=19.6, p<. 001. All participants received either course credit or £5 payment.

Design

A mixed design was used. Each participant was required to search for either of two targets, a ‘swirl’ or a ‘diamond’ (see Figure 6.1). Only one target was ever present on any given trial. For half the participants in each group the ‘swirl’ was paired with the negative UCS (a snake or a spider picture, depending on the participant’s fear) using a similar procedure to that detailed in Experiments 1 and 2, and the ‘diamond’ with neutral pictures. For the remaining half the pairing was reversed (i.e. the ‘swirl’ was paired with neutral pictures and the ‘diamond’ with snake or spider pictures). The number of diamond and swirl trials was equal across each block, as were the number of trials with each set size.

Stimuli

Stimuli used as targets in the visual search task were two complex abstract shapes, both being different arrangements of four white and four black triangles. In one target they were configured into a 'diamond' pattern, while in the other they were configured into a 'swirl.' A third configuration of the same eight triangles was used for all the distractors (see figure 6.1). Targets were always presented upright (0°). However, distractors were randomly presented at four orientations (0°, 90°, 180°, 270°). Target and distractor shapes were black and white, and each was enclosed by a black outline square. Each shape was 10 mm tall and 10 mm wide. Throughout all trials the background was white. As in the previous experiments, target and distractor shapes were evenly spaced around an imaginary circle with a diameter of 100 mm around fixation. Throughout the course of the experiment, target and distractors could appear at any of the stimulus locations around this imaginary circle. A black cross was presented in the centre of the screen 1000 ms before stimulus onset, and remained visible for the duration of each trial. The number of pictures used as the UCS was increased from 24 to 48 for each picture set. Neutral images were again taken from the IAPS. As there were not sufficient snake and spider images in the IAPS alone, pictures were selected from various Internet sites. As in Experiments 1 and 2, pictures were scaled appropriately and centred for fixation on a black background.

Apparatus

The experiment was conducted with the same computers and response boxes used in Experiments 1 and 2.

Procedure

Visual Search Task: Upon entering the laboratory, participants were asked to complete either the Snake (SNAQ) or Spider (SPQ) Fear Questionnaire (Klorman, Weerts, Hastings, Melamed & Lang, 1974) depending upon their score from the earlier screening procedure. Any participants with scores deviating +/- 5 from their initial screening score were excluded from the study, and recorded scores were taken from the second administration of the questionnaires. Two participants were eliminated by this procedure. The SNAQ and SPQ consist of 30 and 31 true or false questions respectively, and provide a subjective index of snake and spider fears (see Appendices V & VI). The minimum possible score for each questionnaire is 0, while the maximum possible score for the SNAQ is 30 and 31 for the SPQ.

Following completion of the questionnaires, participants were given written and verbal instructions regarding the experimental protocol. A sample FR image was shown, and participants were explicitly informed of the contingency between targets and pictures, and reminded that they could withdraw from the study at any time. A conditioning block ensued prior to the practice trials (rather than being incorporated as in Experiments 1 and 2), to strengthen conditioned responses (Glautier et al, 2001). This conditioning block consisted of 24 trials, 12 FR (spiders or snakes depending upon the participants SNQ/SPQ score) and 12 neutral. Each trial began with a display of one of the two target shapes (50mm x 50mm,centred on a white background) for 2 seconds, immediately followed by a congruent image (snake, spider or neutral picture), again of 2 seconds duration. This image served as the UCS, and was centred on a black background. As in Experiments 1 and 2, each neutral and aversive image occupied the entire screen. After an inter-trial interval of 2 s (during which time the monitor display was black), the next trial was automatically initiated. Target presentation was randomised and equal across the conditioning block. Therefore, depending upon their group, each participant viewed either 12 snake or 12 spider pictures, and 12 neutral pictures. During the conditioning phase participants were instructed to look at the screen and learn the contingencies between targets and picture type.

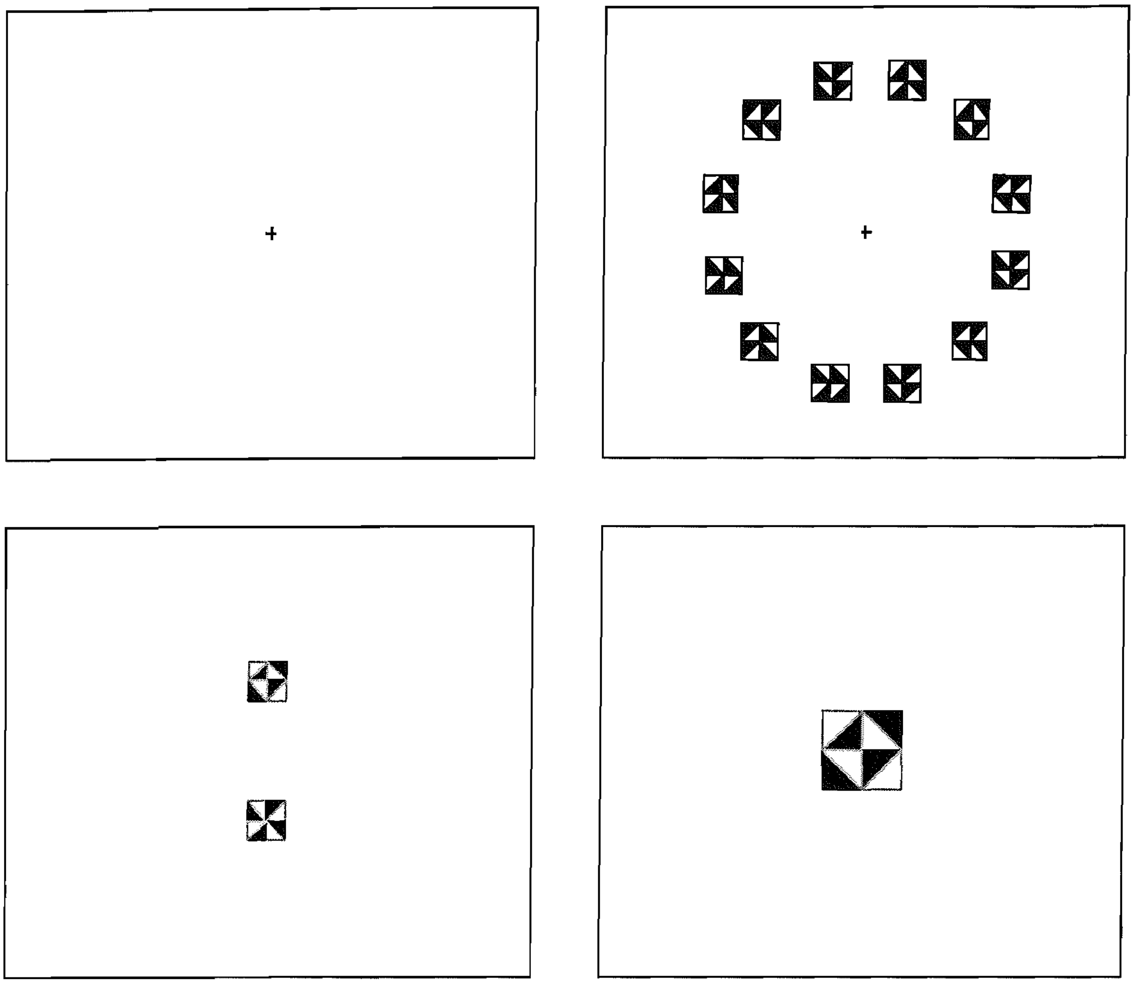

Figure 6.1. The sequence of events in a trial is shown above, beginning with the fixation cross (top left). After a SOA of 1s the visual search display appeared (top right). A negative or neutral target was always present, and the display remained visible until the participant responded. Set size was 1, 3, 6 or 12 items. Following the participants' response was a forced choice task (bottom left) for which there was no time limit. After the participant clicked over the target, an enlarged display was presented for 750 ms (bottom right) followed by a congruent picture (snake, spider or neutral image) of 2s duration.

Following the conditioning block, participants were given verbal instructions for the visual search task, and a practice block of 8 trials. During the practice block, each set size was presented twice in random order, and there were equal numbers of FR and neutral targets. The paradigm used was similar to that of Experiments 1 and 2a. As in previous experiments, participants sat approximately 60cm in front of the monitor and were instructed to make a response upon locating a target as quickly and accurately as possible. A small cross was presented in the centre of the monitor at the start of each trial and, following a SOA of 1000 ms, a search array of between 1 and 12 items was displayed until a response was initiated.

Participants completed 4 blocks of 48 trials, with set size and picture type randomised and equal across each block. Prompts for breaks were displayed after every 24 trials. However, the method of response was different from the previous experiment: participants were required to press the centre button of the button box as quickly as possible upon locating either target using their non-dominant hand (ascertained at the beginning of the experiment). This change was designed to prevent the possibility that participants learned an association to the response button location rather than the target shape. As soon as the centre button was pressed, the search display disappeared and was replaced by the two targets, presented vertically in the centre of the monitor and separated by a gap of 35 mm. Their display size was identical to those used in the visual search (i.e. 10mm x 10mm) and target location was randomised across the selection response; hence the top target was equally likely to be a swirl or diamond. Participants were required to select which target had been present in the previous display by clicking over the appropriate target with the mouse, using their dominant hand. RTs to the mouse responses were not included in data analysis and it was stressed to participants that speed was important in only the initial target response (pressing the centre button on the button box). However, error rates were included in data analysis.

Once the target shape was selected, the two-target display was briefly replaced by a larger display of just the correct target (50 x 50 mm) for 750 ms and followed by a congruent picture for a period of 2 seconds, in order to maintain the strength of the association between search stimulus and picture type. If the participant clicked over an incorrect target, the larger version of the correct target was accompanied by a 'negative' auditory tone. In addition, during the practice block instructions, participants were explicitly told that if they clicked over the incorrect target, the correct target and its appropriate picture would replace it. Hence, there was no advantage for high fear participants in deliberately clicking onto the neutral target to avoid viewing the threatening picture.

IAT Procedure: Following the visual search task, participants completed the IAT as before. Although the side of presentation of the targets was not as critical as in Experiments 1 and 2, it was nonetheless counterbalanced. More importantly, to contend with order effects, the order of presentation was counterbalanced, as in the earlier experiments. The number of trials and the procedure used in the IAT were identical to that implemented for the IAT phase in Experiments 2a and 2b (with the exception of the stimuli used; 'swirl' and 'diamond' shapes replaced the exemplars T and F). As before, the 32 negative and positive words were the same as those used in Experiments 2a and 2b. However, the word beauty replaced the word diamond in order to avoid any possible associations formed by the 'diamond' shape over the course of the experiment. It is important to note that the target shapes used in the visual search task were never given labels, in order that participants were able to form their own mental labels of these abstract shapes.

6.3. Results

Visual Search Task:

Data were analysed by a 4 x 2 x 2 x 2 x 2 mixed analysis of variance (ANOVA), with the two within-subjects variables Set Size (sizes 1, 3, 6 and 12) and Target Valence (negative versus neutral, i.e., paired with a spider or snake versus neutral picture). Between-subjects variables were Negative Target Shape (i.e., which shape was paired with the snake or spider pictures, swirl or diamond), Fear Level (high or low) and Fear Content (snake or spider). As in Experiments 1 and 2, RTs <100 ms and RTs from trials with errors were removed prior to analysis. Figures 6.2 and 6.3 and appendix VII depict the mean visual search RTs differentiated for set sizes, target valence, and fear level.

As expected, there was a main effect of Set Size, *F* (3, 165) = 245.2, p< .001. RTs increased with corresponding increases in Set Size. The most important issue is whether there is a main effect of Target Valence or a Set Size x Target Valence interaction since this would indicate that Target Valence affects visual search. There was a marginal significant effect of Target Valence, *F* (1, 55)= 3.16, p= .08, indicating that the detection of negatively valenced targets (1356 ms ± 55.8) was overall somewhat faster than the detection of neutral targets (1416 ms ± 75.7), and no Target Valence x Set Size interaction (*F*<1).

The significant 3-way interaction of Target Valence x Set Size x Negative Shape, *F* (3, 165)= 9.31, p<. 001, can be traced back to the fact that the diamond, which was associated with fear-relevant picture for half of the participants and with neutral pictures for the other half, could be located faster than the swirl in large set sizes only.

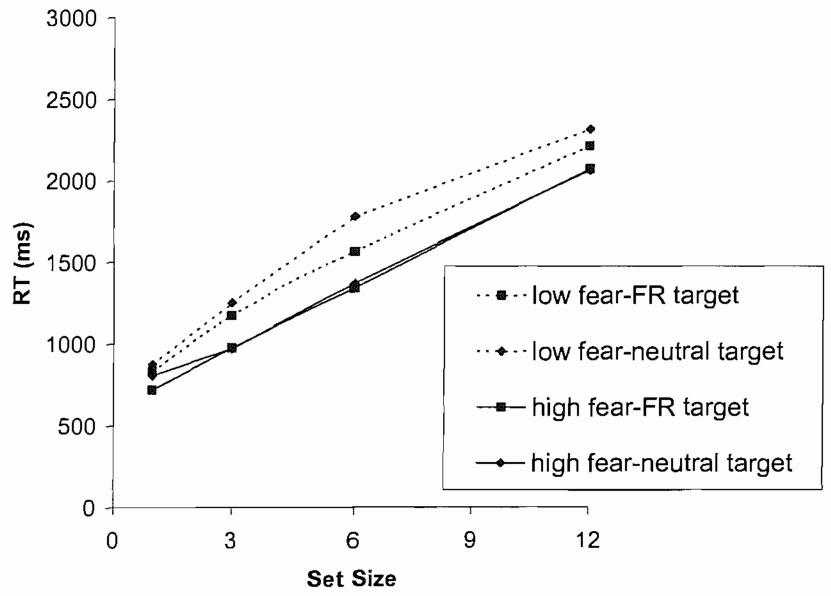

Figure 6.2. Mean visual search RTs for high (solid lines) and low spider (dashed lines) fear participants differentiated for set sizes and targets associated with fear-relevant (FR; squares) or neutral (diamonds) pictures.

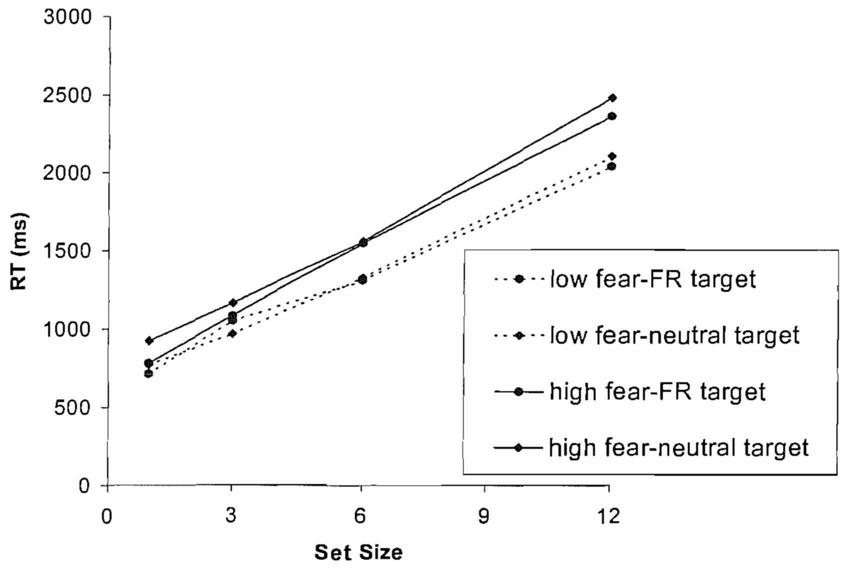

Figure 6.3. Mean visual search RTs for high (solid lines) and low snake (dashed lines) fear participants differentiated for set sizes and targets associated with fear-relevant (FR; circles) or neutral (diamonds) pictures.

Errors were entered into a 4 x 2 x 2 x 2 x 2 mixed ANOVA, using the same factors as the visual search analysis. Again, there was a main effect of Set Size, *F* (3, 165)= 5.58, p< .001 indicating that errors increased with Set Size. There was a significant interaction of Set Size x Fear Level x Fear Content, *F* (3, 165)= 3.05, p< .05 whereby high snake fear subjects made more errors in the largest set size, and a significant 3-way interaction of Target Valence x Negative Target Shape x Fear Level, *F* (1, 55)= 5.16, p<. 05 (low fear participants made more errors when the swirl was paired with snake and spider pictures than when it was paired with neutral pictures). An interaction of Negative Target Shape x Fear Content, *F* (1, 55)= 4.64, p< .05 revealed that participants made more errors when the swirl was associated with spider pictures. Importantly however, there were no significant main effects of Target Valence or Fear, or interactions of Target Valence x Fear or Set Size x Target Valence x Fear (all *F*'s<1). Overall error rates were less than 5%.

IAT: If conditioning has been successful, then RTs of high fear participants (fear of snakes or spiders) to targets associated with fear relevant (i.e., snake or spider) pictures should be faster when paired with negative words than positive words. However, participants with low levels of fear given the same pairing should evidence little or no difference. Fig. 6.4 depicts mean RTs from phases 3 and 5 of the IAT differentiated by Response Congruity for the high and low fear groups. As in Experiment 1, errors and RTs <100 and >2000 ms were excluded from the analysis.

In order to directly test the above hypothesis, the two high fear groups and the two low fear groups were combined and analysed with a 2 x 2 x 2 x 2 mixed ANOVA with one within subjects factor, Response Congruity (congruent or incongruent), and the between subjects factors Fear Level (high vs. low), Fear Content (whether the shape was associated with snake or spider pictures) and Order (congruent or incongruent first). As predicted, there was a main effect of Response Congruity, *F* (1, 55)= 4.32, p< .05. Importantly, there was a significant interaction of Response Congruity x Fear Level, *F* (1, 55)= 4.38, p<. 05. Follow-up t-tests revealed that high fear participants were significantly faster to the congruent than the incongruent response pairing, *t* (30)= 2.82, p<. 01. Low fear participants however, evidenced no significant difference between congruent and incongruent response pairings, *t* (31)= .009, p>. 05. Finally, there was a significant interaction of Response Congruity x Order, *F* (1, 55)= 10.18, p<. 01. This interaction may be interpreted as a weakening of conditioning effects during the IAT; i.e. faster responses if negative words and abstract shapes associated with fear slides are combined during the first data phase of the IAT than during the second data phase.

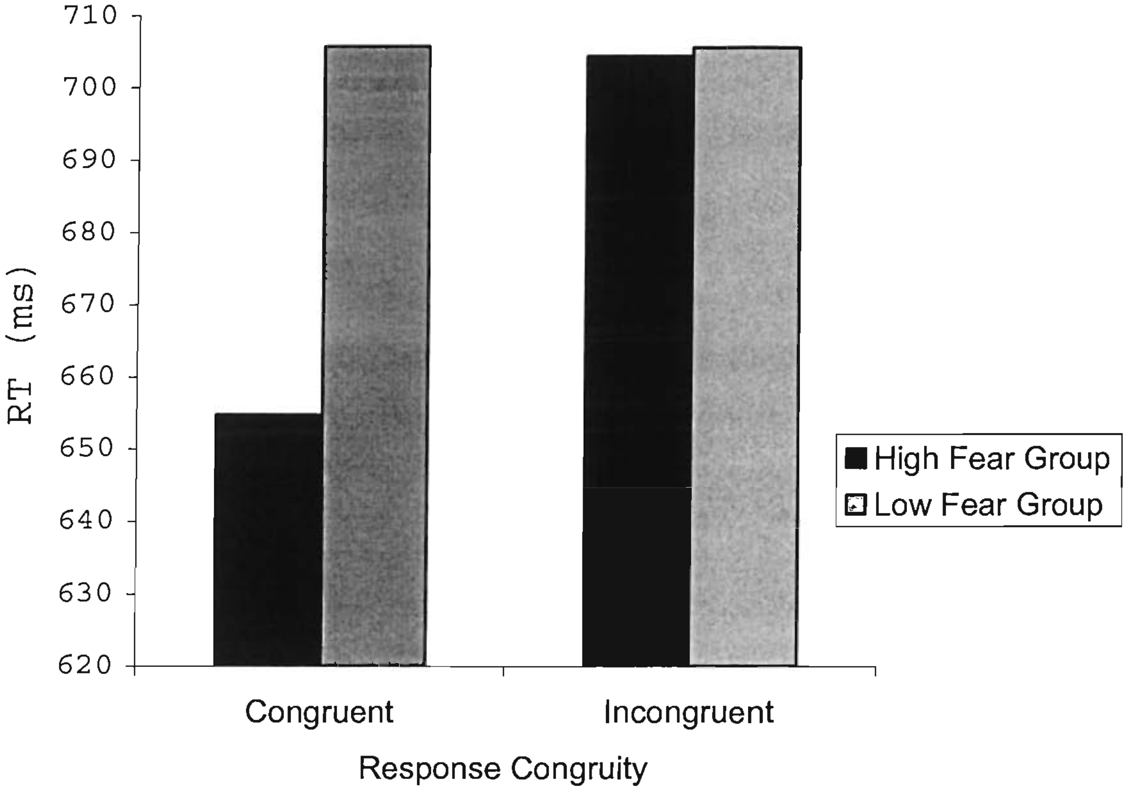

Figure 6.4. RTs in the Implicit Association Test (IAT) to conditioned targets by fear level. Response congruity refers to response pairings. Congruent refers to same responses for the shape associated with fear-relevant pictures and negative words, incongruent to same responses for the shape associated with fear-relevant pictures and positive words.

<u>Correlations:</u> Correlations of the strength of conditioning and visual search RTs were calculated as before. Again, there were no significant correlations for low or high snake fear participants (r= -.061, n, 16, p=ns; r=. 106, n, 15, p=ns respectively) or low and high spider fear participants (r= .447, n, 16, p=ns; r= -.116, n, 16, p=ns respectively).

6.4. Discussion

The results from the present experiment once again show no evidence of attentional biases towards threat stimuli. Moreover, as the pattern of results closely replicates findings from Experiments 1 and 2, the evidence implies that the data are reliable. Although, there were no significant main effects or interactions in the visual search task (other than an expected main effect of Set Size), there was a main effect of Response Congruity in the IAT and importantly, there was a significant interaction of Response Congruity x Fear. In other words, for high fear participants the shape associated with snake or spider pictures in the visual search task was negatively evaluated, whereas for low fear participants, there were no differences in valence between the neutral or FR associated shapes. This strongly suggests that as in Experiment 2, the conditioning procedure was successful. Furthermore, as there were no significant correlations between the strength of conditioning and search efficiency for the FR targets in the visual search task, the lack of a fear effect on search does not seem to be due to weak conditioning in a subset of the participants.

Although there was no main effect of Target Valence, a mild trend was evidenced, whereby participants were faster to locate FR targets than neutral targets. However, this trend did not reach significance and as in previous experiments, the rate of search for both neutral and FR targets was well above the 30ms/item threshold used to describe 'very inefficient' search (Wolfe, 1998a). Again, with the exception of a strong set size effect, no effects or interactions approached significance. As such, the results are consistent with a voluminous literature suggesting that search for conjunctions of features requires effortful, focused attention (e.g. Cave, 1999; Treisman & Gelade, 1980; Wolfe, 1995; Wolfe, 1998; Wolfe, Cave & Franzel, 1989). In light of this, the present experiment offers strong evidence against the preattentive detection of conditioned FR stimuli.

While arguing from what is essentially a null hypothesis is unquestionably difficult, the results from Experiment 3 are important for a number of reasons. The method of response in the current experiment was amended from Experiment 2 in order to redress ambiguity with respect to what was being conditioned (i.e. response button locations or targets). As the IAT result was significant and comparable to the previous study, this consolidates the validity of the IAT as a measure of conditioning. As such, the paradigm used over the course of these experiments confers a number of advantages over other studies. Previous visual search tasks have used either photographs (e.g. Hansen & Hansen, 1988; Öhman, Flykt & Esteves, 2001) or schematic images (e.g. Fox et al, 2000; Öhman, Lundqvist & Esteves, 2001) as threat stimuli with which to measure biases in visual attention. Whereas the former paradigm has the advantage of ecological validity, its disadvantage lies in a lack of strict control over the perceptual features present (see chapter 3 for a more detailed discussion). Similarly, the use of schematic stimuli overcomes the latter problem at the expense of the first (although see Purcell & Stewart, 2002). However, the paradigm employed over Experiments 1-3 capitalises on the strength of both designs without either of the inherent weaknesses. As neither of these criticisms can be levelled at the experiment, the phenomenon of preattentive detection of threat in sub-clinical populations that can guide visual search is somewhat precarious. Whilst the findings of Öhman, Flykt & Esteves (2001) are persuasive, particularly with respect to differential search rates across targets for subjects with specific fear, methodological limitations in their experiments rule out firm conclusions about the preattentive nature of threat detection. As other experiments using the visual search paradigm (e.g. Tipples et al, 2002) have failed to show differences in visual attention for threatening versus non-threatening targets, it is possible that attentional biases in anxiety may well be attributable to factors other than the initial detection of threat.

Although attentional biases have fairly reliably been shown to occur in both dot-probe studies (e.g. Asmundson, Sandler, Wilson & Walker, 1992; MacLeod, Mathews & Tata, 1986) and Emotional Stroop paradigms (Mathews & MacLeod, 1985) it is important to note that neither of these paradigms directly measure the initial allocation of attention. Rather, such biases might be attributable to problems in disengaging from threat (e.g. Fox, Russo, Bowles & Dutton, 2001; Koster, Crombez, Verschuere & De Houwer, 2003; Yiend & Mathews, 2001). If the attentional biases evidenced in these studies were indeed attributable to difficulties in disengaging attention from threat, this would explain the absence of any attentional bias in Experiments 1-3. In order to test this hypothesis, a new experimental paradigm is outlined in the following chapter.

Chapter 7: Experiment 4

7.1 Introduction

Taken together, the results from Experiments 1, 2 and 3 provide no evidence in support of the hypothesis that threat-associated stimuli can be detected preattentively. Moreover, there is no evidence to suggest that visual searches for threat stimuli are any more efficient than searches for neutral stimuli when targets and distractors are comprised of different arrangements of the same features. These results suggest a fundamental limitation on mechanisms proposed to allow threat to guide search (e.g. Fox et al, 2000; Öhman, 1997; Öhman, Flykt & Esteves, 2001; Öhman, Lundqvist & Esteves, 2001), and help to understand why a number of other studies have found no evidence in support of an attentional bias towards threat (e.g. Asmundson & Stein, 1994; Pett, 2001; Tipples et al, 2002, Yiend & Mathews (Exp 1), 2001).

Although the current results show no sign of preattentive detection of threat, data from a number of recent studies (e.g. Amir, Elias, Klumpp & Przeworski, 2003; Fox, Russo, Bowles & Dutton, 2001; Fox, Russo & Dutton, 2002; Yiend & Mathews, 2001, (Experiments 2&3)) indicate that threat stimuli may hold attention more than neutral stimuli once attention has been allocated to them (see also Chapters 3 and 6). Thus, attentional biases evidenced in paradigms such as the dot-probe may be more accurately interpreted as reflecting a difficulty in disengaging attention from threat stimuli once they have been detected, rather than a bias in the initial allocation of attention towards threat. To extrapolate at which stage (if any) attentional biases occur, a new experimental paradigm has been devised, based upon cueing paradigms used to measure visual attention (e.g. Posner, Nissen & Ogden, 1978; Posner, 1980). Whereas most cueing studies use a very limited number of potential target locations (generally two or three), the current experiment used a wide range of locations, enabling a more realistic (although still limited) measure of spatial attention. In addition, increasing the range of distances between cue and target allows comparisons with other studies in which decreased effects of attention have been evidenced with increases in eccentricity from its focus (e.g. Downing & Pinker, 1985; LaBerge, 1983; LaBerge & Brown, 1986).

Experiment 4

Following on from the work of Fox et al (2001, 2002), suggesting that attentional biases to threat reflect a difficulty in disengaging attention from threat rather than a preattentive bias (see section 3.8), the hypothesis of Experiment 4 is that participants will show a significant increase in their RTs to targets following invalid negative cues compared to invalid neutral picture cues. Based upon previous cueing studies (e.g. Posner, 1980), it is predicted that when the cue is invalid, irrespective of picture type and trait anxiety level, all participants will be slower than when the cue is valid or there is no cue (baseline trials). As the cue consists of an abrupt onset, which automatically captures attention (Jonides & Yantis, 1988; Yantis & Jonides, 1984), RTs across valid trials (in which the cue and target appear at the same location) are likely to be faster than no cue (baseline) trials, as attention has already been 'summoned' to the target location. However, when the trial is invalid (i.e. the cue and target appear at different locations in the visual field), attention must be disengaged from the cue before moving to the location of the target and finally engaging upon the stimulus. In invalid trials in which the cue is a negative picture, it is hypothesised that RTs will be slower than when the cue is a neutral picture, as participants are likely to find it harder to disengage their attention from the affective content of the pictures. Finally, previous research has shown a gradient of attention, whereby as eccentricity between an attended object and a subsequent probe is increased (LaBerge, 1983; LaBerge & Brown, 1986), or the distance between a cue and target (Downing & Pinker, 1985) increases, there is a corresponding rise in RTs. In the light of these data, it is hypothesised that RTs and error rates will rise as the distance between the cue and target is increased. It is unlikely that picture type or the participants' level of anxiety will mediate this distance effect.

7.2 Method

Participants

48 undergraduates (10 male, 38 female) from The University of Southampton took part in the experiment in return for course credits. Mean age was 19.8 years (SD=1.54). As in previous experiments, participants were screened for trait anxiety using the STAI-T (Spielberger, Gorsuch & Lushene, 1970). As there are no universally accepted criteria to delineate between different levels of trait anxiety in the literature, cut offs were changed from those used in

Experiment 2 to <30 for low trait anxiety and >40 for high trait anxiety. In addition, only 2 groups were included in the analyses that used anxiety as a factor. As the low-trait cut off score of <40 (used in Experiment 2) includes participants with scores used by some authors (e.g., Fox et al, 2001) to classify participants as being high trait anxious (>35), the cut off for inclusion into the low-trait anxious group was lowered. With this more stringent threshold for low trait anxiety, there is a better chance of detecting whether low trait anxious subjects direct attention away from threatening stimuli (Williams et al, 1988, 1997). Furthermore, using only low and high trait anxious groups allows differences between groups to be seen more clearly. 12 participants were classified as high trait anxious (STAI-T mean=50, SD=5.01) and 15 as low trait anxious (STAI-T mean=26.86, SD=3.76). The remaining 21 participants had trait anxiety scores ranging between 30-40 (mean =35.28, SD=2.91) and were classified as having normal levels of trait anxiety. Mean age of the high trait anxiety participants (2 male, 10 female) was 20.08 years (SD=1.93). For the low trait anxiety group (3 male, 12 female) mean age was 19.4 years (SD=0.9). All participants reported normal or corrected to normal vision (including normal colour perception).

Design

A mixed design was used. The between subjects factor was Anxiety, with two levels (low and high). The within subjects factors were Cue Validity (valid and invalid), Cue Valence (neutral and negative), Target Direction (arrow to the left or right) and Distance between cue and target (5 distances, from 0 (valid) to 4 (furthest away; see Figure 7.2).

Stimuli

Two types of pictures were used in this study as cues. Neutral pictures (48 in total) were taken from the IAPS (Lang, Bradley & Cuthbert, 1997) as in Experiments 1 and 2. However, as there were insufficient numbers of negative images (and due to the possibility of ‘overexposure’ within the participant pool, increasing the likelihood of desensitisation), a further set of pictures was collected from various Internet sites and rated for valence and arousal by a separate group of participants to create a total of 48 negative pictures. All images were centred on a black background and including the background measured 257 pixels x 257 pixels. Target stimuli were green closed headed arrows (the shaft dimensions being 60 x 15 pixels and the arrow head, 25 x 35 pixels), centred within a grey rectangular box 138 pixels wide x 60 pixels tall.

Apparatus

The experiment was carried out with the same equipment used in Experiments 1-3.

Procedure

Upon entering the laboratory, participants completed the trait anxiety scale of the STAI (Spielberger, Gorsuch & Lushene, 1970) and were directed to a research cubicle, where the computerised task took place. Participants were seated approximately 60 cm in front of the apparatus. Verbal instructions were given detailing the format of the experiment and a sample negative image was shown. Participants were given the opportunity to withdraw from the experiment. All participants were then given a practice block of 5 trials.

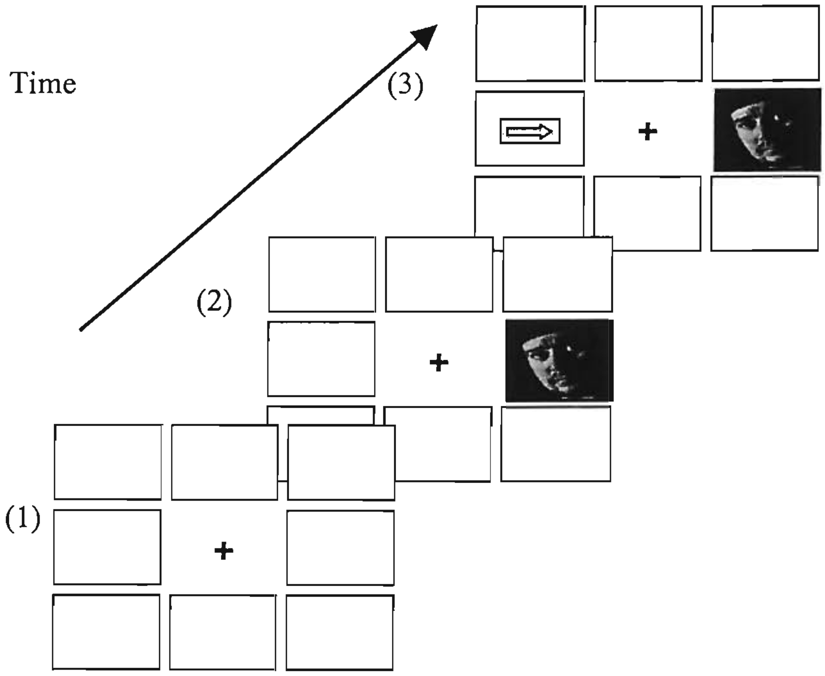

Figure 7.1. Sequence of events per trial. (1) Each trial began with a 1300 ms presentation of a fixation cross and matrix. (2) A neutral (as shown in this example) or negative cue photograph was displayed at any of the 4 middle locations for 200 ms. In baseline trials, no cue was presented and the matrix was displayed on its own for a further 200 ms. (3) A target (arrow) appeared in one of the 8 locations. The participants' task was to respond as quickly and accurately as possible by pressing a button that corresponded with the direction that the arrow

pointed. Throughout this period, the picture cue remained visible and the trial was only terminated once a response had been made. Trials in which the cue and target appeared at the same location are referred to as *valid* trials. When the spatial location of the cue and target differed, the trial was *invalid.*

Each trial began with a 1300 ms presentation of a 3 x 3 matrix of square boxes (with the central box replaced by a 15 x 15 pixel fixation cross). In cued trials this was followed by either a neutral or negative colour photograph (the cue), which randomly appeared at one of four locations within the matrix (the cue never appeared at any of the four corner locations or the centre location). After a SOA of 200 ms, an arrow (the target) was presented in the centre of one of the remaining eight boxes and remained visible (along with the picture) until the participant made a response (see Figure 7.1). In baseline trials (in which there was no cue), SOA between the initial matrix display and the onset of the target was 1500 ms. Participants were asked to depress the button on the response box that corresponded to the direction that the target arrow pointed (i.e. to press the left hand button if the arrow pointed to the left and the right hand button when the arrow pointed towards the right) as quickly and as accurately as possible. In valid trials, the target appeared at the same location as the cue (picture), with the arrow box superimposed on top of the picture. In invalid trials, the target appeared at a different location. Cues were uninformative of the location of the target (12.5% cue validity), as the target was equally likely to appear in any one of the 8 boxes. As pictures never appeared in the corner boxes or the centre of the matrix, this gave a manipulation of five possible distances between cue and target (zero distance when the cue was valid and between one and four 'boxes' away when the target was invalid; see figure 7.2).

Data were collected across 544 trials, divided into 4 blocks of 136 trials, based upon the following variables: 2 Cue valences (neutral and negative) x 4 cue positions x 8 target positions x 2 target directions (left and right) x 4 blocks (plus 32 baseline trials). Hence data were taken from 256 trials with negative pictures (32 valid, 224 invalid) and 256 trials with neutral pictures. An additional 32 no cue trials prevented participants from developing automatic response sets as a result of the fixed time period of each cue (Amir et al, 2003; Stormark, Nordby & Hugdahl, 1995). The data from these trials also provided a baseline for comparisons to be made regarding the costs and benefits for valid and invalid cues. Across each picture type, distance and cue validity were randomised and equal (224 invalid, 32 valid, 16 no cue). Participants were prompted to take a break every 48 trials.

7.3. Results

Data were initially analysed across valid and invalid trials combined, and included all participants. Trials with errors (3.76%) or RTs <100 ms were removed. Due to the large variance in RTs with increases in eccentricity, no upper cut-off point was used. A 5 x 2 repeated measures ANOVA using the within subjects factors Distance (O (valid), 1, 2, 3 and 4; see figure 7.2) and Cue Valence (neutral vs. negative) tested whether there was a main effect of Distance. Based upon previous cueing tasks, valid trials (distance=0) should be faster than invalid trials (distance=4). More important however, will be whether there is a significant interaction of Cue Valence x Distance, i.e. whether the threatening invalid cue pictures cause delayed disengagement compared to invalid neutral cues. Additional analyses examined the valid and invalid trials separately.

RTs (All Participants): Data for all participants showed a significant main effect of Cue Valence, *F* (1, 47)=23.66; p< .001, whereby all participants were slower to respond to targets when the cue was a negative picture than when the cue was a neutral picture. In addition, there was a main effect of Distance, *F* (4, 47) =132.1; p< .001. In order to determine more about this effect, a series of Bonferroni tests with a corrected α level of p< .01 were run (see also appendix VII). Each distance differed significantly from every other distance (with the exception of distances 3 and 4;see figure 7.2), whereby as eccentricity between cue and target increased, RTs rose. A Cue Valence x Distance interaction, *F* (4, 47)= 5.77; p< .001 (see Figure 7.3), was also observed. RTs were generally slower to targets the further away they were from the cue when the preceding cue was a negative picture than when the cue was a neutral picture. However, this effect was predominately attributable to slower RTs when comparing distance 0 (valid) with distance 1(see figures 7.2 and 7.3 and appendix IX). After distance 1, the difference between negative and neutral cues was roughly equal. When the data from invalid trials were analysed alone, the main effects of Cue Valence, *F* (1, 47) =22.97; p< .001, and Distance, *F* (3, 47) =121.45; p<. 001, were again apparent. However, for invalid trials alone, there was no significant interaction between Cue Valence and Distance (*F*<1).

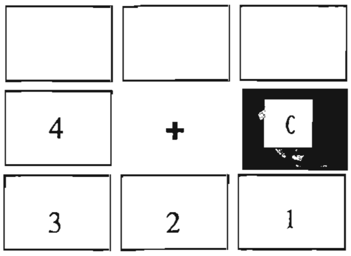

Figure 7.2. There were 5 possible distances between cue and target. Cues only ever appeared in one of the four boxes in the middle of one of the sides. Thus for valid cues, distance was 0, with a maximum possible distance of 4 boxes between cue and target.

For valid cues (distance=0), there was a significant main effect of Cue Valence, *F* (1, 47) =16.68; p< .001, whereby participants again generally evidenced slower RTs to targets when the cue was a negative picture compared to a neutral picture (see Figure 7.3). Importantly however, a planned contrast compared whether the difference between valid and invalid targets varied between negative and neutral cues. It revealed that the cost for invalid over valid cues was greater with negative cues than with neutral cues, *F* (1, 4)= 22.95; p< .001. As participants were significantly slower to negative invalid cues compared to negative valid cues, the effect observed cannot simply be attributed as an overall slowing whenever a negative cue was presented. When a cue is presented at one location and the target appears in a different location, RTs are slower than when the cue and target appear at the same spatial location. However, when the picture cue is negative, this cost is greater than when the picture cue is neutral, as would be expected if it were more difficult to disengage attention from a negative picture than from a neutral picture.

Error Rates (All Participants): There were no significant main effects of Cue Valence, or any interactions (all *F*'s<1) on error rates. However, there was a significant main effect of Distance on error rates, *F* (4, 45)= 19.36, p< .001. As before, a series of Bonferroni tests with a corrected α value of p< .01 were run, which showed a corresponding increase in errors as the distance between cue and target increased (all p values < .001; see also appendix VIII).

RTs by Anxiety: A separate analysis was carried out on participants with scores at either end of the STAI-T (12 high, 15 low). The same factors (Cue Validity, Cue Valence and Distance) were used as in the previous analyses, with the addition of Anxiety as a between-subjects factor. There was no main effect of Anxiety (*F*<1). However, as before, there were main effects of Cue Valence, *F* (1, 25)= 12.37; p< .001, and Distance, *F* (4, 25)= 64.41; p< .001, and an interaction of Cue Valence x Distance, *F* (1, 25)= 6.03; p< .001. (RTs were generally slower to targets the further away they were from the cue when the preceding cue was a negative picture than when the cue was a neutral picture). No other effects or interactions approached significance (all *F*'s<1).

Error Rates by Anxiety: There was a significant main effect of Distance, *F* (4, 104)=14.85, p<. 001, whereby the error rate increased as the distance between the cue and target was increased. No interactions with Anxiety (*F*<1) or any other significant main effects or interactions were evidenced.

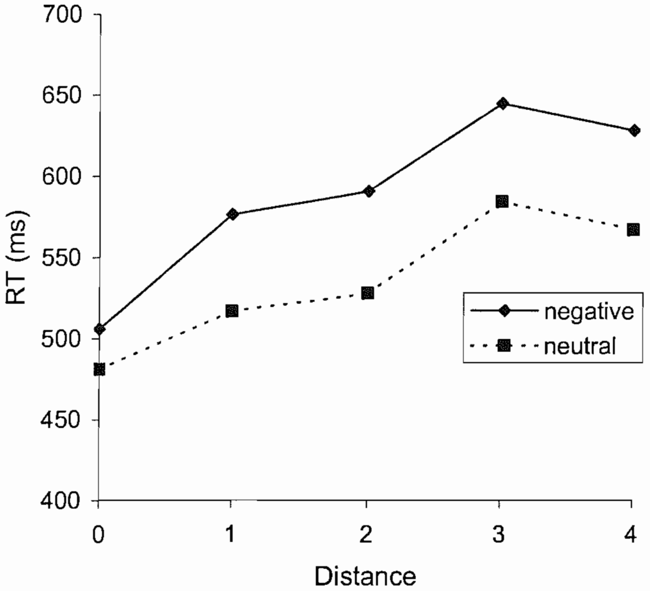

Figure 7.3. RTs by Cue Valence and Distance.

7.4. Discussion

The results from the present experiment are important for a number of reasons. First, the cue produced the same type of attentional gradient found in previous studies (e.g. Downing & Pinker, 1985; LaBerge, 1983; LaBerge & Brown, 1986). However, it is worth noting that there closer physical proximity of the cue to distance 4 than distance 3. Secondly, there was a main effect of Cue Valence, which was larger for invalid than valid cues. Participants found it harder to disengage from negative picture cues compared to neutral pictures, as evidenced by the increased costs for negative cues. It is also worth noting that when the cue was a negative photo, it seemed to slow responses generally. As such, the data provide firm evidence that pictures high in threat value interfere with the disengagement of attention in subjects generally, irrespective of individual levels of trait anxiety. Such findings, whilst not predicted by the clinical model of anxiety proposed by Williams et al, (1988, 1997) are in keeping with the cognitive-motivational account of Mogg & Bradley (1998), which proposes that severe threat will capture attention in all individuals, irrespective of their level of anxiety.

An unexpected finding, already briefly discussed, was the main effect of Cue Valence for valid cues, whereby participants were slower to respond to targets when the cue was a negative picture than when the cue was a neutral picture. As previous work (e.g. Posner, 1980) has demonstrated substantial benefits for valid cues, whereby cues are assumed to reflect automatic (or exogenous) attentional capture, differences in cue type were a surprising finding. As cues summon spatial attention to the location of the cue, RTs to valid cues are rapid, as attention can be assumed to be at the location of the to be presented target. Therefore, one would expect to see similar RTs to valid cues regardless of cue valence as such effects are assumed to be automatic. Furthermore, if the thalamo-amygdala pathway enables rapid processing of threat stimuli (LeDoux, 1996) as an evolutionarily adaptive survival mechanism, one would expect to see a converse pattern of results, i.e. faster responses to valid targets for threatening vs. neutral cues. In keeping with the data from studies one, two and three, it appears that once again threat confers no advantage. Such findings are nevertheless consistent with data from other laboratories (e.g. Fox et al, 2001, Experiment 3), in which participants evidenced slower RTs to valid threatening schematic faces compared to valid schematic happy faces. Perhaps threatening pictures cause a general concern or other negative emotions that would interfere with any task (see also Koster, Crombez, Verschuere & De Houwer, 2003). However, it is also possible that the threat stimuli are identified at a very early stage of visual processing and initiate automatic defence mechanisms, e.g., freezing. An alternative and more plausible explanation, based on the findings from the previous studies however, is that threat appears to generally disrupt the cognitive system (as evidenced by slower RTs in the IAT for participants in the negative picture condition in Experiment 1). When attention has to be moved from threatening pictures (as in the case of invalid trials) this compounds to slow responses still further, whereby the generally slowing in cognitive processing interacts with a difficulty in withdrawing attention away from the more salient stimulus.

A potential problem with this (as in all studies using picture stimuli) however, is that one cue may have more salience than another cue due to differences in low-level perceptual properties. Thus, any differences between cues may be an artefact of these properties rather than their affective content. To redress this possibility, conditioned cues could be used, thus allowing tight control of perceptual properties. Nonetheless, these data appear to offer strong support for the hypothesis that picture threat stimuli interfere with the disengage component of attention.

Chapter 8: Experiment 5

8.1. Introduction

In order to examine further the role of low-level perceptual features and determine whether affective stimuli are more salient than neutral stimuli, Experiment 5 used an amended version of the cueing paradigm from Experiment 4. Instead of picture cues, however, conditioned abstract shapes were used in order to ensure that each cue was matched for perceptual properties. The most important question will be whether participants once again show significant costs for invalid negative cues compared to invalid neutral cues as they did in Experiment 4. As the Unconditioned Stimulus (UCS) will be the same picture set used in the previous experiment, the threat value is likely to be perceived as being high. Thus, it is hypothesised that should the conditioned cues hold attention in a similar manner to the previous study, *all* participants will show greater costs for invalid negative cues than for invalid neutral cues due to the content of the UCS. However, if the effects that were observed in Experiment 4 were attributable to stimulus salience rather than threat associated with it, then RTs are likely to be similar across both types of cue.

Experiment 5

8.2 Method

Participants

48 undergraduates (14 male, 34 female) from The University of Southampton took part in the experiment in return for course credits. Mean age was 20.67 years (SD= 4.33). As in previous experiments, participants were screened for trait anxiety using the STAI-T (Spielberger, Gorsuch & Lushene, 1970) and separated into high and low trait anxiety groups. Of these, 17 participants were classified as being high in trait anxiety (STAI-T mean = 46.5, SD= 8.3), and 9 participants as low trait anxious (STAI-T mean = 26.56, SD= 2.65) based upon cut off scores of >40 and <30 respectively. A further 22 participants with scores between these cut-off points (STAI-T mean = 35.36, SD= 3.06) provided data used in the overall conditions that did not include anxiety as a factor. Mean age of the high trait anxiety participants (4 male, 13 female) was 19.88 years (SD= 0.86) and for the low trait anxiety group (5 male, 4 female) mean age was 23.1 years (SD = 9.41). All participants reported normal or corrected to normal vision (including normal colour perception).

Design

A mixed design was used. The between subjects factor was Anxiety, with two levels (low and high). The within subjects factors were Cue Validity (valid and invalid), Cue Valence (neutral and negative), Cue Shape ('diamond' or 'swirl') and Distance between cue and target (5 distances).

Apparatus

The experiment was carried out with the same equipment used in Experiments 1-4.

Stimuli

The experiment used the picture set from Experiment 4 for the conditioning block that preceded the experimental trials. Pictures (the UCS) were aligned onto a black background and centred for fixation. The display size of each image, along with the black background, was 1152 X 870 pixels, and occupied the entire screen. Cues were the same shapes used in Experiment 3, a 'swirl' and a 'diamond'. In order to facilitate conditioning, shapes were coloured and their colour counterbalanced across participants, i.e. for half the participants, the swirl was blue and the diamond was brown, and for the remainder the pairing was reversed (i.e. the swirl was brown and the 'diamond' blue). The RGB colour mixture used for the blue was (r=100, g=170, b=210) and for the brown (r=200, g=160, b=100). The two colours were matched for luminance (50 c/m^2) using a Minolta GS-100 colour meter.

Each shape was 257 pixels tall by 257 pixels wide in both the conditioning block and the cueing paradigm. Target stimuli were green closed headed arrows (the shaft dimensions being 60 x 15 pixels and the arrow head, 25 x 35 pixels), centred within grey rectangular box 138 pixels wide x 60 pixels tall.

Procedure

Cueing Task: Upon entering the laboratory, participants completed the trait anxiety scale of the STAI (Spielberger, Gorsuch & Lushene, 1970) and were directed to a research cubicle, where the computerised task took place. Verbal instructions were given detailing the format of the experiment and a sample negative image was shown. Participants were given the opportunity to withdraw from the experiment without prejudice if they felt uncomfortable with the content of the pictures.

Before any experimental trials took place, all participants underwent a conditioning block. During this period, participants were seated approximately 60cm in front of the monitor in a dimly lit room. Participants were explicitly informed of the relationship between the shape and the picture content. Both shape and colour were counterbalanced across participants. The conditioning block consisted of the presentation of 24 pairings of the diamond and swirl shapes with negative and neutral pictures (or vice-versa). Each trial consisted of a 2 second presentation of either the swirl or the diamond, followed by a 2 second presentation of the UCS (a negative or neutral picture depending upon which shape was paired with which image). The presentation of the pictures was randomised across the conditioning block and the timing and display controlled by the computer program. Thus, each participant viewed 12 pairings of the diamond with neutral pictures and 12 pairings of the swirl with negative pictures or vice-versa.

Once the conditioning period had been completed, participants were given written and verbal instructions detailing the experimental protocol, and a practice block of 5 trials. The procedure was identical to that of Experiment 4, with the exception that the cues used were the diamond and swirl shapes viewed in the conditioning block rather then the picture cues used in the previous experiment. All timing details and the number of trials were the same as Experiment 4. However, in the practice block, all trials were reinforced; i.e. once a response to the direction of the target had been made, a picture congruent with the target was displayed for 2 seconds. Similarly, in experimental trials there was a consequence approximately every three trials out of eight, whereby a picture congruent with the cued shape was displayed for 2 seconds. These reinforcement trials were randomised and equal across blocks, so that participants had 48 reinforcement trials per block. As before, data were collected across 4 blocks of 136 trials. The number and organisation of trials was identical to Experiment 4.

IAT: Upon completing the cueing task, participants remained in the experimental cubicle and undertook an IAT. As in previous experiments, the purpose of this was to measure whether conditioning had been successful. With successful conditioning, participants would be faster to associate the shape previously paired with negative images in the cueing task with negative words than positive words. The procedure used and the number of trials in each phase was the same as that used in Experiment 3. However, the diamond and swirl shapes were coloured blue and brown respectively (or vice-versa) using the colours used in the cueing task. As before, both shapes were matched for equal luminance (50 c/m^2). To contend with order effects, four versions of the IAT were used; (V1 & V2: blue swirl and brown diamond, with order of presentation of target and target concept switched between them, and V3 & V4, blue diamond and brown swirl). In all other respects, the procedure was identical to that of Experiment 4.

8.3. Results

Cue Task RTs (All Participants): Data were analysed by a series of ANOVAs. Trials with errors (2.04%) and RTs <100 ms were removed. As in Experiment 4, due to the large variance in RTs with increases in eccentricity, no upper cut-off point was used. To facilitate the interpretation of the data, results were initially analysed for all participants across valid and invalid trials, prior to a separate analysis that included Anxiety level. Across valid and invalid trials combined, there was no main effect of Cue Valence, F (1, 47)= 1.99, p> .05, although there was a significant main effect of Distance, F (4, 188)= 214.63, p<. 001 (see figure 8.1). Bonferroni adjusted comparisons (see appendix X) revealed a general pattern whereby as the distance between the cue and target increased, RTs increased correspondingly. These differences were significant across all distances, except between distance 1 and 2 (mean RTs 518.1 vs. 521.5 ms respectively). As in Experiment 4, RTs to targets that appeared 3 steps away from the cue were slower than RTs for targets that appeared 4 steps away (580.9 vs. 557.8 ms respectively), which is predicted by the actual distances between cues and targets. Most important was whether there was an interaction of Cue Valence x Distance (i.e., whether conditioned negative invalid cues held attention more effectively than conditioned neutral invalid cues). This interaction did not approach significance (F <1). No other effects or interactions approached significance.

Cue Error Rates (All Participants): Error rates were low (2%), with no main effect of Cue Valence ($F < 1$) or interaction of Cue Valence x Distance, F (4, 188)= 1.97, p>. 05. However, there was a main effect of Distance, F (4, 188)= 27.79, p<. 001, whereby error rates generally increased with corresponding increases in distance between the cue and target (see Appendix XI).

Cue Task RTs by Anxiety: When the data were separated by trait anxiety, there were no main effects of Cue Valence or any significant interactions (all F's <1). As before, there was a main effect of Distance, F (4, 96)= 85.62, p<. 001, whereby RTs rose as the distance between cue and target was increased. An analysis of invalid trials only showed a significant main effect of Distance, F (3, 135)= 148.06, p< .001, but no main effects of Cue Valence, F (1, 45)= 1.76, p>. 05 or Anxiety, F (2, 45)= 1.1, p>. 05. Finally, there were no significant main effects or interactions (all F's <1) when the data were analysed by valid trials only.

Cue Task Error Rates by Anxiety: For valid and invalid trials combined, there was a main effect of Distance, F (4, 96)= 15.85, p<. 001. Errors increased with corresponding increases in distance between the cue and target. A significant interaction of Distance x Anxiety, F (4, 96)= 2.64, p<. 05 revealed that participants with high levels of trait anxiety made more errors as the distance between the cue and target increased compared to participants with low levels of trait anxiety. However, there were no significant main effects of Cue Valence, F (1, 24)= 2.67, p> .05 or Anxiety ($F < 1$). Importantly, there was no significant interaction of Distance x Cue Valence, F (4, 96)= 2.32, p>. 05.

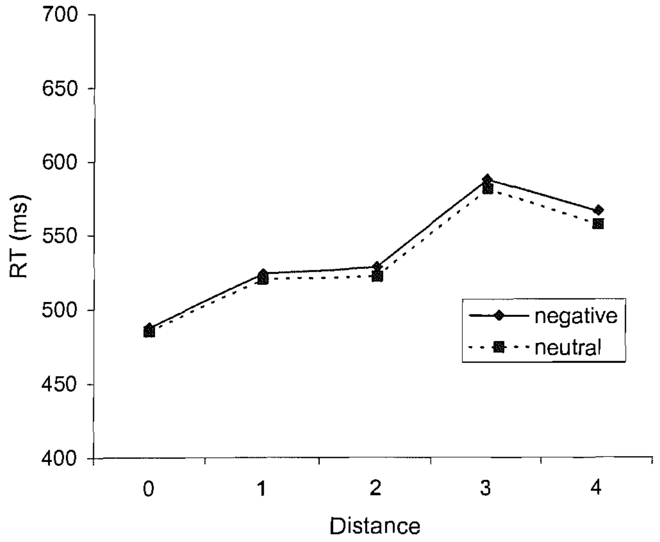

Figure 8.1. RTs to negative (solid line) and neutral valenced (dashed line) conditioned cues by distance.

<u>IAT:</u> If conditioning has been successful, then RTs to negatively valenced cues (from the cueing task) should be faster for 'bad' words than 'good' words in the IAT. Similarly, faster RTs to good words should be evidenced for neutral valenced cues. Mean RTs from phases 3 and 5 of the IAT were entered into a 2 x 2 x 2 repeated measures ANOVA. There was one within subject's factor, Response Congruity (congruent or incongruent; i.e. if a 'swirl' had been paired with negative pictures in the visual search task, then the pairing of the swirl and 'bad' words was considered to be a congruent response pairing. Similarly, if the swirl was paired with 'good' words following a negative association in the visual search task, this was considered an incongruent response pairing). The two between subject's factors were Cue Shape ('swirl' or 'diamond' associated with negative pictures) and Order (congruent or incongruent first). As in Experiments 1-3, errors and RTs <100 and >2000 ms were excluded from the analysis.

There was a significant main effect of Response Congruity, *F* (1, 40)= 9.05, p<. 01, whereby participants were faster to associate negative cues with bad words than good words (mean RTs 617.4 vs. 652.1ms respectively). In addition, there were significant interactions between

8.4. Discussion

The main effect of Response Congruity in the IAT indicates that cues were effectively conditioned; that is to say, there was a transfer of affective valence onto cues associated with negative and neutral pictures. Indeed, the level of significance reached was higher than in previous experiments, probably due to the fact that participants could associate each cue with two properties, shape and colour. However, despite the effectiveness of the conditioning, there was no main effect of Cue Valence and more importantly, no interaction of Cue Valence x Distance. In contrast to Experiment 4, threat associated cues did not increase attentional dwell time, nor was there any slowing in RTs to valid threat associated cues compared to valid neutral cues. There was however, a main effect of distance in keeping with previous studies showing a gradient of attention (e.g., Downing & Pinker, 1985; LaBerge, 1983; LaBerge & Brown, 1986). As in Experiment 4, there was a general pattern whereby RTs rose as the number of steps between cue and target increased, with the exception of distances 3 and 4 (presumably since distance 3 was physically further away than distance 4). There was no significant difference between distances 1 and 2 (although the mean RTs were slightly larger). As this pattern of results was not observed in other studies, it appears more likely to be a chance variation.

Although the positive result in the IAT is consistent with previous studies, paradoxically the very reason that might have led to the ease of conditioning may be the reason why no difficulty in disengaging attention from the threat stimulus emerged. In the current experiment, cues were comprised of relatively simple arrangements of geometric shapes. Detecting either of two basic properties, i.e., colour or shape, could define the valence of each cue. As such, participants only needed to remember and detect either of these features in order to form the correct associative valence. Although participants would be able to form associations of valence quickly upon the presentation of the cue, as no further processing was required there would be little to hold the subjects' attention. When picture cues were used, however, each picture would require more extensive computation and probably deeper cognitive processing. Due to the content of the negative pictures, it is quite feasible that their increased salience compared to neutral pictures would result in greater distraction, rumination and a general slowing (as observed in the slower IAT responses for the group given the negative targets in Experiment 1). For conditioned cues however, these effects were not observed due to ease of processing each cue type.

With these factors in mind, it is important to consider which mechanisms and properties may be responsible for the increased dwell times to threat associated stimuli exhibited in Experiment 4 and previous studies showing similar effects (e.g., Amir, Elias, Klumpp & Przeworski, 2003; Fox, Russo, Bowles & Dutton, 2001; Fox, Russo & Dutton, 2002; Tipples & Sharma, 2000; Yiend & Mathews, 2001, (Experiments 2&3)). Assuming that the conditioning in Experiment 5 was sufficient, then the different results between Experiments 4 and 5 suggest that attention will only be held at a location if further visual processing there will continue to yield relevant information. If early processing of a stimulus indicates that it could be particularly threatening, then it will continue to hold attention in all individuals (e.g. Mogg & Bradley, 1998). However, if the cue is neutral or low in threat, attention can be readily moved in individuals low in anxiety. For individuals with high levels of anxiety, however, even mild threat can disrupt the disengagement of attention (e.g., Fox et al, 2001). Thus, it appears likely that whilst the strongest predictor of pathological anxiety is elevated trait anxiety (Eysenck, 1992), a major cause in the maintenance of anxiety is a difficulty in disengaging attention from threat related stimuli. However, in order for threatening stimuli to elicit these effects, the stimuli must be particularly salient. Hence, conditioned threat cues lack either sufficient impact (unlikely when one considers that negative valence words and schematic faces have been sufficient to hold attention in previous studies), or they do not require the extensive cognitive processing that would hold attention. Clearly, further work is needed in this area.

Chapter 9:General Discussion

9.1 Introduction

The rationale for Experiments 1-3 was to examine whether threat associated stimuli could be detected preattentively in the absence of basic features. Whereas many cognitive models of visual attention suggest that search can only be efficient when basic features such as orientation, size and colour can be used to guide search (e.g. Cave, 1999; Treisman & Gelade, 1980, Wolfe, Cave & Franzel, 1989), clinical theory (e.g. Mogg & Bradley, 1998; Williams et al, 1988, 1997) and neurobiological evidence (e.g., LeDoux, 1996) suggests that threat automatically captures attention. If threat can be detected automatically, then threat would have to be added to the list of basic features (see Wolfe, 1998a) known to support preattentive processing. In light of this, the set of studies outlined in Chapters 4-6 are important to further validate either clinical theory or cognitive models of visual attention. Should the search slopes for threat-associated targets have a flat function, this would require a revision of models of visual cognition, which suggest that the efficiency of a search is determined only by features shared by targets and distractors.

9.2 Attentional Processes: Experiments 1-3

Experiment 1 used a between subjects design in which half the participants searched for a threat associated target, with the remainder searching for a neutral target. There were no differences in search slopes between groups. However, although there was a trend in the predicted direction in the conditioning measure (the IAT), this did not reach significance. Consequently, a within subjects design was used in Experiment 2 in which participants searched for threat-associated and neutral targets. Participants were separated into 3 groups (low, medium and high trait anxiety) and performed exactly the same search task. There were no differences in search rate across both target type and anxiety level. However, as there was a significant main effect in the IAT, targets were successfully conditioned. Finally, in Experiment 3, 63 participants (selected from a large student pool screened for snake and spider fear) were placed into one of 4 groups (low or high snake or spider fear). Again there were no differences between groups (low vs. high fear) in their rate of search for either target (snake or spider vs. neutral targets), with all participants evidencing slow, inefficient search. However, there were significant main effects in the IAT indicating successful conditioning of the targets.

Taken together, the data from Experiments 1-3 show no evidence that threat associated stimuli can be detected preattentively. In fact, there was no advantage for targets paired with negative stimuli compared to targets paired with neutral stimuli. Instead, threat-associated pairings merely affected overall performance in subsequent tasks that required participants to decide on stimuli characteristics (e.g., slower RTs in the IAT in Experiment 1 for participants given negative pairings in the visual search task relative to participants given neutral pairings) without affecting search rates. These findings suggest that as long as there are no simple features (e.g., colour, size, orientation, angles, simple line configurations, etc) that differentiate threatening targets from non-threatening distractors, any differences in search for threatening targets occur only *after* stimuli have been identified by focused attention.

This conclusion (i.e., slower performance following negative valence) is supported by comparable findings from other studies using different paradigms. For example, Eastwood, Smilek & Merikle (2003), asked participants to count features belonging to schematic faces that varied in their affective expression. Comparable with our IAT findings, performance was affected more by faces displaying negative affect compared to faces expressing positive or neutral emotion. Similarly, in a study in which participants were presented with arrays of faces containing either all the same or a discrepant affective expression (Fox et al, 2000), participants were slower to detect the absence of a discrepant face when the display consisted of angry faces relative to displays consisting of happy faces. Although in the Fox et al study, participants did evidence faster RTs to threatening faces compared to happy and neutral faces, there was no evidence of preattentive detection.

Certainly, faces appear to possess unique attributes that are particularly effective in capturing visual attention. However, while visual searches for faces generally show an "anger superiority effect", that is, a search asymmetry whereby angry faces are detected more quickly than neutral or happy faces (e.g. Eastwood, Smilek & Merikle, 2001; Fox et al, 2000; Öhman, Lundqvist & Esteves, 2001), few (if indeed, any) studies have provided convincing evidence that angry faces can be identified preattentively unless basic featural differences are present. Of the experiments in which search slopes for angry faces have yielded a flat function (e.g., Hansen & Hansen, 1988; White, 1995) or within the 5-6 ms/item generally taken to indicate preattentive processing (Treisman & Souther, 1985), these effects have been confounded by low-level featural properties (e.g. Hansen & Hansen, 1988) and hence, open to alternative interpretations (see Purcell et al, 1996). For example, in the White (1995) study, RTs were independent of set size for both happy *and* angry faces. More importantly however, when the faces were inverted, flat search functions were still evidenced. As inversion is thought to result in the loss of configural information (De Gelder, 2000) and disrupt holistic facial processing (Yin, 1969), these data indicate that basic features were used to guide search rather than affective properties.

In a recent replication of Öhman et al's (2001) study, Purcell & Stewart (2002) obtained almost identical data using a control face that consisted of a rearrangement of the features present in the threatening face but portraying no affective expression. Purcell & Stewart concluded that the advantage for angry faces is due to a configuration of lines pointing outwards towards the edge of the face. Certainly, the V shaped eyebrow configuration appears particularly salient and could constitute a basic feature that may account for the more rapid detection of angry faces. However, as a recent study by Tipples, Atkinson & Young (2003) demonstrated an advantage for V shaped eyebrows only when other internal facial features were present (there were no differences in RTs for schematic faces containing V and inverted V shaped eyebrows alone), it appears that some form of holistic face processing may combine a number of basic features. Interestingly, imaging studies suggest that fearful eyes by themselves may be enough to activate amygdala responses (e.g., Morris, deBonis & Dolan, 2002) and that affective expression is much harder to recognise if conveyed in only the lower half of the face (i.e., with 'neutral' eyes). This effect is noticeable more so for angry and fearful than happy faces (De Gelder, 2000). However, as these paradigms do not typically include the use of distractor stimuli, further investigation including distractors would be necessary to determine whether any aspect of these configurations can be detected preattentively (cf. Pessoa et al, 2002).

Irrespective of which component or components are responsible for the rapid detection of angry faces, data from the study of Eastwood et al (2001) show that, relative to the baseline, search is slowed more by a happy target than it is speeded by a threatening one (cf. Fox et al, 2000). Certainly, schematic faces are problematic in terms of ecological validity, as imaging studies have yielded results inconsistent with the general findings obtained from experiments using more naturalistic facial photographs. While enhanced amygdala responses to emotional relative to neutral schematic faces have been evidenced (e.g. Wright et al, 2002), in the Wright et al study, activity in the left amygdala was no larger for angry relative to happy faces, in contrast to the common finding of increased left amygdala activation to angry and fearful faces (e.g. Morris et al, 1996). Schematic faces also pose further problems with respect to familiarity. As smiley faces are more frequently encountered than angry faces (Öhman, Flykt & Lundqvist, 2000), their familiarity may account for the differential search rate in the study of Fox et al (2000) & Öhman et al (2001), as familiarity in itself is known to affect search efficiency (Greene & Rayner, 2001; Wang, Cavanagh & Green, 1994). In particular, participants are faster to detect an unfamiliar target embedded amongst familiar distractors than vice-versa (Greene & Rayner, 2001). Wang et al (1994) propose that whereas familiar items can be processed rapidly, unfamiliar items require extra processing, which attracts attention to them. As such, an angry face embedded amongst a homogeneous background of happy faces might be detected more rapidly than vice-versa, as it is unfamiliar. However, this interpretation resides upon the notion that angry schematic faces are indeed unfamiliar, a fact that requires further validation. As the distractors in the Eastwood et al (2001) study were always neutral faces, their familiarity was equivalent in each condition. In light of this, any effects of familiarity are difficult to interpret. However, irrespective of this, data from the above studies remains open to a number of viable alternate explanations other than the threatening faces preferentially capturing visual attention.

Leaving aside issues concerning the use of schematic faces, other studies have used photographic stimuli in order to achieve more ecological validity. Öhman, Flykt & Esteves (2001) did find evidence of preattentive detection of Fear Relevant (FR) targets (snakes and spiders) amongst Fear Irrelevant (flowers and mushrooms) distractors, with the strength of this effect mediated by the subjects' particular fear; i.e., high spider fear participants were faster to find snakes amongst flowers than vice-versa, and faster still to find spiders. These findings, although persuasive, are equivocal. As photographs were used, controlling for the absence of basic features is extremely difficult. Although many different pictures were used to reduce the chances of low-level perceptual confounds, one cannot exclude this possibility as an explanation. It is possible that snakes (and spiders) possess certain unique features, which lead to their more rapid detection. In fact, Öhman, Flykt & Esteves themselves suggest that images of snakes and spiders might contain a "signature feature that, because of evolutionary history, is as readily picked up by the visual system as a vertical line among slanted ones" (Öhman et al, 2001, p. 476), thus leading to their automatic detection. However, as Tipples et al (2002) obtained similar results for the detection of 'cuddly animals' and fruit, one could argue that nature has also equipped us with fruit and kitten detectors. It is more likely that numerous objects can be detected against a background of flowers and mushrooms due to basic featural differences. Furthermore, as both electrophysiological (e.g., Holmes, Vuilleumier & Eimer, 2003) and imaging studies (e.g., Pessoa et al, 2002, although see Vuilleumier et al, 2001) have illustrated that spatial attention is necessary for the detection of threat, it appears that threatening stimuli are not detected without attention.

9.3 Attentional Biases Reflect Delayed Disengagement from Threat: Experiments 4-5

If spatial attention is necessary to identify the threatening nature of a stimulus, then any attentional biases must occur only *after* the target has been attended in the course of a serial search. The results of Experiment 4, in which participants responded to the direction of arrows following threatening or neutral pictures in an adaptation of an exogenous cueing task, showed increased RTs to valid and invalid threat cue trials relative to neutral cues, consistent with the idea that threat effects occur after selection rather than before. As contrast analyses revealed that this effect was greater for invalid cues, these data can be assumed to reflect difficulty in disengaging attention from threatening stimuli rather than threat stimuli simply slowing cognitive processing. These findings parallel data obtained by other researchers (e.g., Amir et al, 2003; Fox et al, 2001, 2002; Yiend & Mathews, 2002), in which greater costs were evidenced for negative than neutral invalid cues. However, when conditioned cues were used in a further experiment (Experiment 5), there were no differences in dwell times between cue types, despite strong conditioning effects (evidenced in an IAT). It is proposed that as cues in

Experiment 5 could be identified as threatening or neutral using colour or shape alone (or by a combination of both properties), relatively little computational operations were required. Consequently, as less visual information needed to be extracted, the conditioned abstract stimuli required no further processing. However, when complex pictures were used as cues (as in Experiment 4), more extensive computations were required by the visual system, perhaps to learn more about the nature of the threat. As such, the increased salience and explicit nature of the negative cues held attention more effectively (see section 8.4 for a more detailed account). An important point to note is that this does not imply that conditioned cues are ineffective as threat stimuli (as exemplified by the IAT data). Rather, if no further visual analysis is required, then there is no reason for attention to remain static when other items compete for attention.

9.4 Limitations: Is the IAT an Appropriate Test to Measure Conditioning Effects?

The conclusion that association with threat has no effect on visual search requires arguing from the null hypothesis, and thus it becomes important to explore every other possible explanation for the lack of an effect. One possibility is that the conditioning was ineffective, but the significant results from the IAT argue against this possibility. Another possibility is that the IAT is not an appropriate test of conditioning. A test like the IAT, based on behavioural measures, seems far preferable to self-report measures, which are generally unreliable (de Jong et al, 2001) and open to demand characteristics (Herman et al, 2002). The rationale for the use of the IAT is described in Chapter 4 and will not be reiterated here. However, it is important that any potential criticisms of the IAT as a tool for the measurement of conditioning are addressed in order to validate the findings that threatening stimuli hold no attentional advantage when targets and distractors are matched for shared features.

A potential problem with the IAT is that as it measures one target relative to the other, it cannot be considered an absolute indication of threat (De Houwer, 2002). For example, faster RTs to the pairing of flowers with positively valenced words and insects with negatively valenced words than vice-versa might lead to the conclusion that insects are implicitly disliked whilst flowers are implicitly liked. However, if insects are associated more with negative than positive valence, while flowers are associated with positive and negative attributes to the same extent, then the same conclusion is likely to be drawn. In other words, the IAT measures the relative valence of the two categories being tested rather than the absolute valence of each. To circumvent this, studies using neutral and positive, or neutral and negative targets have been used and demonstrated similar findings (De Houwer, 2002). In the case of Experiments 2 and 3, one target was associated with negative attributes, whilst the other was associated with neutral attributes. Although the evaluative conditioning procedure is likely to have increased the valence of targets paired with neutral stimuli, the valence of the negatively paired stimuli will primarily be responsible for the main effect of response congruity seen in the IAT. In accordance with this hypothesis, Hermans et al (2002) had participants undergo a classical conditioning (CC) and an evaluative conditioning (EC) procedure. In the CC paradigm, neutral faces (CS+) were paired with an electric shock (UCS) or no consequence (CS-). In the EC task, an identical procedure was followed, with the exception that the faces were followed by neutral (CS $_{\text{pos}}$) or negative (CS $_{\text{pos}}$) adjectives (the UCS). Using valence ratings and an affective priming task, significant conditioning to the CS+ was evidenced across both paradigms, with the magnitude of this effect roughly equivalent across the two measures. More importantly, the data showed that in each case this effect was due to the CS+ being perceived as more negative rather than the CS- becoming more positive.

Similarly, the different effects shown by participants with low and high snake and spider fear in Experiment 3 support the idea that the change in valence of the CS is attributable to a change in valence in the negative CS rather than increased valence of the neutral CS. As participants with low fear evidenced no significant main effect in the IAT, whereas there was a significant main effect of congruent response pairing for participants with high fear, the obvious interpretation is that snake and spider associated targets became threatening for participants with high snake or spider fear but not for participants with low snake or spider fear. In addition, as other studies have successfully used the IAT to differentiate between participants with high and low snake and spider fears (e.g., Teachman, Gregg & Woody, 2001), the IAT appears well suited to differentiate between groups with low and high animal fear.

A variety of other methodologies have been developed to measure attitudes and associations between items with positive and negative valence in recent years such as the go/no-go association task (GNAT; Nosek & Banaji, 2001) and the extrinsic affective Simon task (EAST; De Houwer, 2003), with the suitability of each test dependant upon the task demands. However, recent studies have shown that the IAT correlates with neurophysiological measures. For example, in a study by Phelps et al (2001), white American subjects were shown photographs of unfamiliar black and white males during fMRI scans. It was found that the amount of amygdala activation to black faces correlated with the degree of racial prejudice evidenced in an implicit test (the IAT) and potentiated startle. However, when the faces were of familiar, positively regarded males, no correlations were evidenced. As the role of the amygdala in the activation and expression of fear is well documented (see Chapter 3), and as there is evidence that the IAT correlates with this degree of activation, the use of the IAT as a measurement of conditioning to threat associated targets may yield similar results to physiological measures of conditioning.

9.5 Is the Main Effect of Target in the IAT a Measurement of Disgust Rather than Fear?

While the IAT reliably measures implicit valence, it is possible that the IAT results evidenced in Experiments 2 and 3 may reflect the measurement of disgust rather than fear. Therefore, the lack of any attentional bias in Experiments 1-3 might simply reflect the fact that the stimuli used here were not threatening enough to trigger the threat detectors that would bias attention. This explanation is rejected on two grounds. Firstly, attentional biases have been evidenced to disgust stimuli as well as threat stimuli (e.g., Charash & McKay, 2002), so the logic that predicts preattentive detection of threat also predicts preattentive detection of disgust. Either way, these experiments show no evidence that the emotional content of these stimuli can be detected preattentively. Secondly, disgust and fear are intimately related, particularly in spider phobia (de Jong & Merckelbach, 1998; Sawchuk et al, 2002; Smits et al, 2002; Tolin et al, 1997). Due to this close relationship, it is often difficult to extrapolate whether disgust or fear is the primary emotion for a given fear stimulus. Whereas blood-injury-injection (BII) phobia is primarily characterised by increased parasympathetic activity (characteristic of disgust reactions), spider fear subjects typically evidence increased sympathetic and parasympathetic activity more closely resembling fear states (Sarlo et al, 2002; Tolin et al, 1997). Therefore, whilst BII phobia is primarily associated with disgust, fear is the principal emotive state in spider fear (Sawchuk et al, 2002; Tolin et al, 1997). As IAPS images are a validated set of affective picture stimuli (Bradley & Lang, 2000; Lang et al, 1997), which have been used in numerous studies to measure attentional biases to threat (e.g., Yiend & Mathews, 2002), their ability to evoke threat related responses is well documented. In light of this, the most plausible interpretation of the IAT data is that negative targets were successfully associated with threat and that the IAT reliably measured this transfer of affective valence.

9.6 Future directions

An important limitation of any study used to measure attentional biases is the nature of the threat stimuli used. As has been discussed in Chapter 3, conclusions from many previous search studies remain equivocal, as basic featural differences are seldom eliminated. Using schematic faces is less problematic in this respect (although see Purcell & Stewart, 2002) but comes at a cost of lower ecological validity. In light of these shortcomings, the paradigm developed over the course of Experiments 1-3 aimed to circumvent these problems by using carefully controlled abstract stimuli (CS) paired with negative or neutral pictures (UCS). While this methodology allows greater control of the perceptual features present, combined with a higher threat value, it is possible that although the CS were strong enough to evoke effects in the IAT, they do not pose a strong enough threat to allow any attentional biases to surface. While phylogenetic stimuli such as angry faces produce strong conditioning effects (e.g., Armony & Dolan, 2002), it remains possible that abstract stimuli do not present a strong enough threat. This interpretation is unlikely however, as attentional biases have been shown in both emotional Stroop (e.g., Williams et al, 1996) and dot-probe paradigms (e.g., MacLeod et al, 1986) using affective words. However, Koster, Crombez, Van Damme, Verschuere & De Houwer (2003) recently used a variation of the Posner cueing paradigm, in which cues (coloured squares) were conditioned to carry either negative (CS+ followed by a 100dB burst of white noise) or neutral valence (CS- followed by a 71dB neutral tone). Participants evidenced significantly faster RTs to valid CS+ than valid CS- targets, indicating that attention was preferentially drawn towards targets associated with threat. Additionally, RTs were slower to invalid CS+ targets relative to invalid CS- targets, indicating that participants found it difficult to disengage their attention from the threat associated target. In keeping with the idea that strong threat will capture attention in all individuals (Mogg & Bradley, 1998), these effects were not mediated by state or trait anxiety.

While these data show preferential attentional capture for threat-associated colour patches, a separate experiment (Koster, Crombez, Verschuere & De Houwer, 2003) used IAPS pictures in a dot-probe paradigm. However, unlike previous studies, a baseline condition (in which 2 neutral pictures were presented simultaneously) was included, so that RTs to probes following neutral pictures could be compared with RTs to probes following threat pictures. Three types of IAPS images were used; neutral (N) pictures (e.g. an umbrella) medium threat (MT) pictures (e.g. a man with a knife) and high threat (HT) pictures (e.g. mutilations). From these, three types of picture pairs were created; N-N, MT-N and HT-N. If attention is captured by threat (i.e. increased vigilance), then RTs to probes appearing at the location previously occupied by a threat picture should be faster than RTs to probes in the baseline (N-N) condition. On the other hand, if the attentional bias is due to problems disengaging from threat, then RTs to probes occurring at the location previously occupied by neutral pictures will be slower than RTs to probes appearing at the location of threatening pictures in the MT-N and HT-N trials. There were no significant differences between participants' RTs to mild-threat versus neutral pictures (i.e. RTs to probes occurring at the location of the threat picture in the MT-N pairing were roughly the same as probes following N-N trails). However, RTs to high threat trials were significantly *slower* than neutral trials (i.e. RTs to probes occurring at the location of threat pictures in HT-N trials were slower than RTs to probes following baseline N-N trials). In keeping with previous dot-probe experiments (e.g. MacLeod, Mathews & Tata, 1986), RTs to probes occurring at neutral locations on threat-neutral trials were slower than probes presented at the location of threat pictures on HT-N trails, suggesting the attentional bias was due to difficulty in disengaging attention from threat pictures rather than a bias in the initial allocation of attention. Consequently, while negative valenced pictures demonstrate that attentional biases may occur *post* identification, it is feasible that they lack a strong enough threat value to enable preattentive biases to materialise. Further research using IAPS pictures combined with electric shock or aversive white noise would be a profitable avenue for future research, particularly in Experiment 3, as studies (e.g. Öhman, 1993) have shown that participants are especially sensitive to tactile pairings (i.e. electric shocks) as a UCS to animal stimuli.

It is of course possible that the lack of any preattentive biases in Experiments 1-3 might be attributable to the non-clinical samples used over the course of the experiments. However, attentional biases have been found in numerous studies using college populations, separated by state (e.g. Bradley, Mogg & Millar, 2000) and trait (e.g., Fox et al, 2002) anxiety. Furthermore, although one cannot directly compare student and clinical populations, trait anxiety in the high anxious student group in Experiment 2 (STAI-T=56.7) was similar to levels measured in a clinical sample with GAD (STAI-T=58.8) tested by Mogg, Millar & Bradley (2000). Further experiments using clinical groups would be informative to see whether these effects are only evidenced in extreme populations, although this option is unfeasible for ethical reasons. Furthermore, as Mogg & Bradley (1988) have argued, while mild threat stimuli can effectively capture attention in highly anxious individuals, more potent stimuli such as threatening pictures, should readily capture attention in individuals with anxiety levels within the normal population range. However, the present experiments have shown that if there are no basic features to guide attention (as in Experiments 1-3), then any biases will only become apparent in later stages (see Experiment 4) and will not occur if no further processing is required (Experiment 5).

Future studies could profit from the use of psychophysiological measures such as GSRs and Heart Rate responses in order to provide a useful alternative index of conditioning efficacy to that of the behavioural measures used in the current paradigms. In addition, ERPs could enable the time course between preattentive and attentive mechanisms to be measured and form a useful link between behavioural measures and autonomic activity. Finally, as there were no RT differences between invalid negative and invalid neutral conditioned cues in Experiment 5, despite strong conditioning effects, it is possible that the difficulty subjects experienced in disengaging attention from threatening pictures in Experiment 4 might be attributable to something other than their affective content. To test this hypothesis, a further study could examine the degree of correlation between affective content and the time taken to disengage attention. As the picture cues used in Experiment 4 have normative valence ratings, this could be achieved using a similar paradigm to that of Experiments 4 and 5, whereby a subject would see a negative cue and then respond to a probe at a different location. The RT in reacting to this probe would be measured and the affective rating for each picture paired with the average RT to respond to a probe appearing after this cue. If the strength of the emotional content determines how hard it is to disengage attention, there should be a correlation between RT and affective rating. As in the Koster et al (2003) study, it is hypothesised that as the valence of the picture decreases (i.e., the picture becomes more negative), the time taken to disengage attention from the negative picture cue is likely to decrease. This might also interact with trait anxiety in accordance with the model of Williams et al (1988; 1997).

9.7 Conclusion

Contemporary models of visual cognition (e.g. Cave, 1999; Wolfe et al, 1989; Wolfe, 1995) purport that basic features (e.g., colour and orientation) can be identified preattentively and used to guide visual search. More complex properties that require the conjoining of features typically require more extensive computation than is traditionally associated with preattentive vision. Thus search latencies for items other than basic features are dependent upon the number of distractors present. Neurobiological (e.g., LeDoux, 1996) and behavioural evidence (e.g. Öhman et al, 2001) however, suggests that complex visual stimuli can be detected preattentively if strongly associated with threat. In support of current models of visual attention, there was no evidence that anything as general as 'threat' can be detected preattentively when targets and distractors are matched for shared features. While practice and increased top-down activation can make searches for threatening stimuli more efficient, attentional biases are unlikely to be attributable to preattentive mechanisms unless basic features are associated with threat. However, once a threatening stimulus is located, it is harder to disengage attention from it relative to a neutral stimulus, providing the stimulus is complex enough to warrant further processing.

References

Agras, S., Sylvester, D., & Oliveau, D. (1969). The epidemiology of common fears and phobias. *Comprehensive Psychiatry, 10,* 151-156.

American Psychiatric Association. (1994). Diagnostic and Statistical Manual of Mental Disorders (Fourth Edition). Washington, DC.

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 and Therapy* (In Press).

Asmundson, G.J.G., Sandler, L.S., Wilson, K.G., & Walker, J.R. (1992). Selective Attention towards physical threat in patients with Panic Disorder. *Journal of Anxiety Disorders. 6 (4),* 295-302.

Asmundson,G.J.G., & Stein, M.B. (1994). Dot-Probe evaluation of cognitive biases in patients with Panic Disorder-A failure to replicate and extend. *Anxiety, 1 (3),* 123-128.

Beck, A.T. (1976). Cognitive Therapy and the Emotional Disorders. New York: International Universities Press.

Beck, A.T., Emery, G & Greenberg, R.L. (1985). Anxiety Disorders and Phobias: A Cognitive Perspective. New York: Basic Books.

Bieling, P.J., Antony, M.M., & Swinson, R.P. (1998). The State-Trait Anxiety Inventory, Trait Version: structure and content re-examined. *Behaviour Research and Therapy. 36,* 777-788.

Bower, G.H. (1981). Mood and Memory. *Amer. Psychol., 36,* 129-148.

Bouton, M.E. (1994). Conditioning, remembering and forgetting. *Journal of Experimental Psychology: Animal Behaviour Processes, 20,* 219-231.

Bradley, B.P., Mogg, K., & Millar. (2000). Covert and overt orienting of attention to emotional faces in anxiety. *Cognition & Emotion, 14,* 789-808.

Bradley, M.M., Greenwald, M.K., Petry, M., & Lang, P.J. (1992). Remembering pictures: Pleasure and arousal in memory. *Journal of Experimental Psychology: Learning, Memory & Cognition, 18,* 379-390.

Bradley, M.M., & Lang, P.J. (2000). Measuring emotion: Behaviour, feeling and physiology. In R. Lane, & L. Nadel (Eds). The Cognitive Neuroscience of Emotion (pp. 242-276). New York: Oxford University Press.

Broadbent, D.E. (1958). Perception and Communication. New York: Pergamon Press.

Broadbent, D.E. (1982). Task combination and selective intake of information. *Acta Psychologica,* 50, 253-290.

Bruce, V., Green, P.R., & Georgeson, M.A. (1996). Visual Perception: Physiology, Psychology And Ecology (3RD Edition). Bath: Psychology Press.

Cave, K.R. (1999). The FeatureGate model of visual selection. *Psychological Research, 62,* 182-194.

Cave, K.R. & Bichot, N.P. (1999). Visuospatial attention: Beyond a spotlight model. *Psychonomic Bulletin & Review, 6 (2),* 204-223.

Cave, K.R. & Wolfe, J.M. (1990). Modelling the Role of Parallel Processing in Visual Search. *Cognitive Psychology, 22,* 225-271.

Cepeda, N.J., Cave, K.R., Bichot, N.P., & Kim, M.S. (1998). Spatial selection via feature-driven inhibition of distractor locations. *Perception and Psychophysics, 60, (5)* 727-746.

Charney, D.S., Nagly, L.M., Bremner, D., Goddard, A.W., Yehuda, R. & Southwick, A.W (1996) in B.S Fogel, Schiffer & S.M Rao. (Eds). Neuropsychiatry (pp. 257-281). Baltimore: Williams & Wilkins.

Cherry, E.C. (1953). Some experiments on the recognition of speech with one and two ears. *Journal of the Acoustical Society of America, 25,* 975-979.

Clark, L.A., Watson, D., & Mineka, S. (1994). Temperament, Personality, and the Mood and Anxiety Disorders. *Journal of Abnormal Psychology, 103 (1)* 103-116.

Copland, J.D. & Lydiard, R.B. (1998). Brain Circuits in Panic Disorder. *Biological Psychiatry, 44,* 1264-1276.

De Gelder, B. (2000). Recognising emotions by ear and by eye. In R. Lane, & L. Nadel (Eds). The Cognitive Neuroscience of Emotion. (pp. 84-105). New York: Oxford University Press.

De Houwer, J. (2002). The Implicit Association Test as a tool for studying dysfunctional associations in psychopathology: strengths and limitations. *Journal of Behaviour Therapy and Experimental Psychiatry, 33,* 115-133.

De Houwer, J. (2003). The Extrinsic Affective Simon Task. *Experimental Psychology, 50 (2),* 77-85.

De Houwer, J., & Thomas, S. (2000). Associative learning of likes and dislikes: A review of 25 years of research on human evaluative conditioning. Manuscript submitted for publication.

De Jong, P., & Merckelbach, H. (1998). Blood-Injection-Injury Phobia and Fear of Spiders: Domain Specific Individual Differences in Disgust Sensitivity. *Person. Individ. Diff, 24 (2),* 153-158.

De Jong, P., Pasman, W., Kindt, M., & van der Hout, M.A. (2001). A reaction time paradigm to assess (implicit) complaint-specific dysfunctional beliefs. *Behaviour Research & Therapy, 39,* 101-113.

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

Downing, C.J. & Pinker, S. (1985). The spatial structure of visual attention. In M. Posner & O. Martin. (Eds.). *Attention and Performance XI* (pp. 171-187). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

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

Duncan, J. (1984). Selective Attention and the Organisation of Visual Information. *Journal of Experimental Psychology: General, 113 (4)* 501-517.

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

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

Eastwood, J.D., Smilek, D., & Merikle, P.M. (2001). Differential attentional guidance by unattended faces expressing positive and negative emotion. *Perception & Psychophysics, 63, (6),* 1004-1013.

Eastwood, J.D., Smilek, D., & Merikle, P.M. (2003). Negative facial expression captures attention and disrupts performance. *Perception & Psychophysics, 65, (3),* 325-358.

Egly, R., & Homa, D. (1984). Sensitisation of the visual field. *Journal of Experimental Psychology, Human Perception & performance, 10,* 778-793.

Ekman, P. (1984). Expression and nature of emotion in K. Scherer, & P. Ekman. (Eds).Approaches to Emotion Hillsdale, NJ: Erlbaum.

Enns J.T., & Rensink, R.A. (1990). The Influence of Scene-Based Properties on Visual Search. *Science, 247,* 721-723.

Eriksen, C.W., & Yeh, Y. (1985). Allocation of attention in the visual field. *Journal of Experimental Psychology, Human Perception & performance, 11,* 583-597.

Esteves, F., & Öhman, A. (1993). Masking the face: Recognition of emotional facial expressions as a function of the parameters of backward masking. *Scandinavian Journal of Psychology, 34,* 1-18.

Eysenck, M.E. (1992). Anxiety: The Cognitive Perspective. Hove: Lawrence Erlbaum Associates.

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 (1),* 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, (4)* 681-700.

Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition & Emotion, 16* (*3*), 355-379.

Gardner, G.T. (1973). Evidence for Independent Parallel Channels in Tachistocopic Perception. *Cognitive Psychology, 4,* 130-155.

Glautier, S., De Houwer, J., & Redhead, E. (2001). Evaluative Conditioning and Affective Priming: Awareness, Contingency, and Cue-Competition. Unpublished Manuscript.

Gratacòs, M., Nadal, M., Martín-Santos, Pujana, M.A., Gago, J., Peral, B., Armengol, L., Ponsa, I., Miró, R., Bulbena, A. & Estivill, X. (2001) A Polymorphic Genomic Duplication on Human Chromosome 15 Is a Susceptibility Factor for Panic and Phobic Disorders. *Cell, 106,* 367-379.

Greenwald, A.G., McGee, D.E. & Schwartz, J.L.K. (1998). Measuring Individual Differences in Implicit Cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74 (6)* 1464-1480.

Greenfield, S. (2001). The Private Life Of The Brain. London: Penguin Books.

Gray, J.A. (1982). The Neuropsychology of Anxiety. Oxford: Clarendon Press.

Gray, J.A., & McNaughton, N. (2000). The Neuropsychology of Anxiety: Second Edition. New York: Oxford University Press.

Hampton, C., Purcell, D.G., Bersine, L., Hansen, C.H., & Hansen, R.D. (1989). Probing 'pop-out': Another look at the face-in-the-crowd effect. *Bulletin of the Psychonomic Society, 27,* 563-566.

Hansen, C.H., & Hansen, R.D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology, 54,* 917-924.

He, Z.J., & Nakayama, K. (1992). Surface versus features in visual search. *Nature, 359,* 231-233.

Heller, W., Nitschke, J.B., Etienne, M.A. & Miller, G.A. (1997). Patterns of Regional Brain Activity Differentiate Types of Anxiety. *Journal of Abnormal Psychology, 106, (3)* 376-385.

Hermans, D., Vansteenwegan, D., Crombez, G., Baeyens, F., & Eelen, P. (2002). Expectancy-learning and evaluative learning in human classical conditioning: affective priming as an indirect and unobtrusive measure of conditioned stimulus valence. *Behaviour Research and Therapy, 40,* 217-234.

Hoffman, J.E. (1998). Visual attention and eye movements. In. H. Pashler, (Ed.). Attention (pp. 119-153). Hove: Psychology Press.

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., & Müller, H. (1993). Search via recursive rejection (SERR): A connectionist model of visual search. *Cognitive Psychology*, *25,* 43-110.

Jonides, J. (1981). Voluntary vs. automatic control over the mind's eye's movement. In J.B. Long & A.D. Baddeley (Eds.). *Attention and performance IX* (pp. 187-203).

Joseph, J.S, Chun, M.M. & Nakayama, K. (1997). Attentional requirements in a 'preattentive' feature search task. *Nature*, *387,* 805-807.

Kahneman, D. (1973). Attention and Effort. NJ: Prentice-Hall.

Keogh, E., Dillon, C., Georgiou, G., & Hunt, C. (2001). Selective attentional biases for physical threat in physical anxiety sensitivity. *Journal of Anxiety Disorders, 15*, 299-315.

Kim, M.S., & Cave, K.R. (1995). Spatial Attention in Visual Search for Features and Feature Conjunctions. *Psychological Science, 6 (6)* 376-380.

Klauer, K.C. (1998). Affective Priming. *European Review of Social Psychology, 8,* 67-103.

Klorman, R., Weerts, T.C., Hastings, J.E., Melamed, B.G., & Lang, P.J. (1974). Psychometric Description of Some Specific-Fear Questionnaires. *Behaviour Therapy, 5,* 401-409.

Koster, E.H.W., Crombez, G., Verschuere, B., & De Houwer, J. (2003). Selective attention to threat in the dot probe paradigm: differentiating vigilance and difficulty to disengage. *Behaviour Research and Therapy.* In Press.

Koster, E.H.W., Crombez, G., Van Damme, S.V., Verschuere, B., & De Houwer, J. (2003). Is attention drawn to signals of imminent threat? *Emotion.* In Press.

Kristjánsson, Á., & Tse, P.U. (2001). Curvature discontinuities are cues for rapid shape analysis. *Perception & Psychophysics*, *63,* 390-403.

LaBerge, D. (1993). Spatial Extent of Attention to Letters and Words. *Journal of Experimental Psychology: Human Perception and Performance, 9 (3)* 371-379.

LaBerge, D., & Brown, V. (1986). Variations in size of the visual field in which targets are presented: An attentional range effect. *Perception & Psychophysics, 40 (3)* 188-200.

Lang, P.L., Bradley, M.M., & Cuthbert, B.N. (1997). International Affective Picture System (IAPS): Technical Manual and Affective Ratings.

Lang, P.L., Bradley, M.M., & Cuthbert, B.N. (1998). Emotion, Motivation and Anxiety: Brain Mechanisms and Psychopathology. *Biological Psychiatry*, *44,* 1248-1263.

Lavie, N. (1995). Perceptual Load as a Necessary Condition for Selective Attention. *Journal of Experimental Psychology: Human Perception and Performance, 21 (3),* 451-468.

LeDoux, J.E. (1996). The Emotional Brain. London:Phoenix.

LeDoux, J.E. (1998). Fear and the Brain: Where Have We Been, and Where Are We Going? *Biological Psychiatry*, *44,* 1229-1238.

LeDoux, J.E. (2000). Emotion Circuits in the Brain. *Annu. Rev. Neuroscience, 23,* 155-184.

Lesch, PK., Bengel, D., Heils, A., Sabol, SZ., Greenberg, B.D., Petri, S., Benjamin, J., Müller, CR., Hamer, DH., & Murphy, DL. (1996). Association of Anxiety-Related Traits with a Polymorphism in the Serotonin Transporter Gene Regulatory Region. *Science*, *274,* 1527-1531.

Luck, S. (1998) in H. Pashler. (Ed) Attention (pp. 257-295). Hove: Psychology Press.

M. Maguire, R. Morgan & R. Reiner (Eds.). (1996). The Oxford Handbook Of Criminology. Oxford: Clarendon Press.

MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional Bias in Emotional Disorders. *Journal of Abnormal Psychology, 95 (1),* 15-20.

Marks, I.M., & Nesse, R.M. (1994). Fear and fitness: An evolutionary analysis of anxiety disorders. *Ethology and Sociobiology, 15,* 247-261.

Mathews, A. (1990). Why Worry? The Cognitive Function of Anxiety. *Behaviour Research and Therapy, 28 (6)* 455-468.

Mathews, A., May, J., Mogg, K., & Eysenck, M. (1990). Attentional Bias in Anxiety. Selective Search or Defective Filtering? *Journal of Abnormal Psychology, 99 (2)* 166-173.

McNally, R.J. (1998). Information-processing Abnormalities in Anxiety Disorders: Implications for Cognitive Neuroscience. *Cognition and Emotion, 12 (3)* 479-495.

Medin, D.L., & Ross, B.H. (1996). Cognitive Psychology. Second Edition. Fort Worth: Harcourt Brace College Publishers.

Mineka, S., & Öhman, A. (2002). Phobias and Preparedness: The Selective, Automatic, and Encapsulated Nature of Fear. *Biological Psychiatry, 52,* 927-937.

Mogg, K., & Bradley, B.P. (1998). A cognitive-motivational analysis of anxiety. *Behaviour Research and Therapy, 36,* 809-848.

Mogg K., & Bradley, B.P. (1999). Some methodological issues in assessing attentional biases for threatening faces in anxiety: a replication study using a modified version of the probe detection task. *Behaviour Research and Therapy, 37,* 595-604.

Mogg, K, McNamara, J., Powys, M., Rawlinson, H., Seiffer, A., & Bradley, B.P. (2000). Selective attention to threat: A test of two cognitive models of anxiety. *Cognition & Emotion, 14* (3) 375-399.

Mogg, K, Millar, N. & Bradley, B.P. (2000). Biases in Eye Movements to Threatening Facial Expressions in Generalized Anxiety Disorder and Depressive Disorder. *Journal of Abnormal Psychology, 109 (4)* 695-704.

Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology, 11,* 56-60.

Morris, J.S., deBonis, M., & Dolan, R.J. (2002). Human Amygdala Responses To Fearful Eyes. *NeuroImage, 17,* 214-222.

Morris, J.S., Frith, C.D., Perrett, D.I., Rowland, D., Young, A.W., Calder, A.J., & Dolan, R.J. (1996). A different neural response in the human amygdala to fearful and happy expressions. *Nature, 383,* 812-815.

Morris, J.S., Öhman, A., & Dolan, R.J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature, 393*, 467-473.

Morris, W.N. (1992). A functional analysis of the role of mood in affect systems. In M.S. Clark. (Ed). Emotion. Newbury Park: Sage.

Mozer, M.C. (1991). The Perception of Multiple Objects: A Connectionist Approach. Cambridge: The MIT Press.

Neisser, U. (1967). Cognitive Psychology. New York: Appleton-Century-Crofts.

Norman, D.A. (1986) Towards a theory of memory and attention. *Psychological Review*, *75*, 522-536.

Nosek, B., & Banaji, M.R. (2001). The go/no-go association task. *Social Cognition, 19,* 625-666.

Öhman, A. (1992). Fear and Anxiety as emotional phenomena: Clinical, phenomenological, evolutionary perspectives, and information processing mechanisms. In M. Lewis & J.M. Haviland (Eds). Handbook of the emotions (pp. 511-536). New York: Guilford.

Öhman, A. (1993). Stimulus prepotency and fear learning: Data and theory. In N. Birbaumer & A. Öhman (Eds). The structure of emotion: psychophysical, cognitive and clinical aspects (pp. 218-239). Seattle:WA. Hogrefe and Huber.

Öhman, A. (1996). Preferential Preattentive Processing of Threat in Anxiety: Preparedness and Attentional Biases. In R.M. Rapee. (Ed). Current controversies in the Anxiety Disorders (pp. 253-290). New York: Guilford Press.

Öhman A (1997). As Fast as the Blink of an Eye: Evolutionary Prepardness for Preattentive Processing of Threat. In P.J. Lang, R.F. Simons, & M. Balabin (Eds.). Attention And Orienting: Sensory and Motivational Processes (pp. 165-184). New Jersey: Lawrence Erlbaum Associates Publishers.

Öhman, A., Esteves, F. & Soares, J. F. (1995). Preparedness and preattentive associative learning: Electrodermal conditioning to masked stimuli. *Journal of Psychophysiology*, *9,* 99-108.

Öhman, A., Flykt, A. & Esteves, F. (2001). Emotion Drives Attention: Detecting the Snake in the Grass. *Journal of Experimental Psychology: General*, *130 (3)* 466-478.

Öhman, A., Flykt, A. & Lundqvist, D. (2000). Unconscious emotion: Evolutionary perspectives, psychophysical data, and neurophysiological mechanisms. In R. Lane, & L. Nadel (Eds). The Cognitive Neuroscience of Emotion (pp. 296-327). New York: Oxford University Press.

Öhman, A., Lundqvist, D. & Esteves, F. (2001). The face in the crowd revisited; A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology, 80,* 381-396.

Öhman, A., & Mineka, S. (2001). Fears, Phobias, and Prepardness: Toward an Evolved Module of Fear and Fear Learning. *Psychological Review, 108 (3)* 483-522.

Öhman, A., & Mineka, S. (2003). The Malicious Serpent: Snakes as a Prototypical Stimulus for an Evolved Module of Fear. *Current Directions In Psychological Science, 12 (1)* 5-9.

Palmer, S.E. (1999). Vision Science: Photons to Phenomenology. Cambridge: MIT Press.

Pashler, H. (1998). Attention. (Ed). Hove: Psychology Press.

Pashler, H., Johnston, J.C. & Ruthruff, E. (2001). Attention and Performance. *Annual Review of Psychology, 52,* 629-51.

Pett, C. (2001). A multi-method study of anxiety. Unpublished Ph. D thesis.

Pessoa, L., Kastner, S., & Ungerleider, L.G. (2002). Attentional control of the processing of neutral and emotional stimuli. *Cognitive Brain Research, 15,* 31-45.

Phelps, E.A., O'Connor, K.J., Cunningham W.A., Funayama, E.S., Gatenby, J.C., Gore, J.C., Banaji, M.R. (2001). Performance on indirect measures of race evaluation predicts amygdala activation. *Journal of Cognitive Neuroscience, 12 (5),* 729-738.

Pinker, S. (1997). How the mind works. London: Penguin Books.

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

Posner, M.I., & Cohen, Y. (1984). Components of Visual Orienting. In H. Bouma & D. Bowhuis (Eds). Attention & Performance X (pp. 531-556). Hillsdale, NJ: Erlbaum.

Posner, M.I., Nissen, M.J., & Ogden, W.C. Attended and Unattended Processing Modes: The Role of Set for Spatial Location. In H.L. Pick & I.J. Saltman (Eds). (1978). Modes of Perceiving and Processing Information (pp.137-157). Hillsdale, N.J: Lawrence Erlbaum Associates.

Posner, M.I. & Peterson. S.E. (1990). The Attention System of the Human Brain. *Ann. Rev. Neurosci, 13.* 25-42.

Posner, M.I., Snyder, C.R.R., & Davidson, B.J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General, 109*, 160-174.

Power, M., & Dalgleish, T. (1997). Cognition and Emotion: From Order to Disorder. Hove:Psychology Press.

Purcell, D.J., Stewart, A.L., & Skov, R.B. (1996). It takes a confounded face to pop out of a crowd. *Perception*, *25*, 1091-1108.

Purcell, D.J., & Stewart, A. (2002). The Face in the Crowd: Yet Another Confound. Poster presented at the 43rd Annual Meeting of the Psychonomic Society, Kansas City, Missouri.

Remington, R.W., & Pierce, L. (1984). Moving attention: Evidence for time-invariant shifts of visual selective attention. *Perception & Psychophysics, 35, 393-399.*

Rolls, E.T. (1999). The Brain and Emotion. Oxford: Oxford University Press.

Rosen, J.B., & Schulkin, J. (1998). From Normal Fear to Pathological Anxiety. *Psychological Review, 105 (2)* 325-350.

Rizzolatti, G., Riggio, L., Dascola, I., & Umiltá, C. (1987). Reorienting attention across the horizontal and vertical meridians: evidence in favour of a premotor theory of attention. *Neuropsychologica, 25 (1A)* 31-40.

Sawchuk, C.N., Lohr, J.M., Westendorf, D.H., Meunier, S.A., & Tolin, D.F. (2002). Emotional responding to fearful and disgusting stimuli in specific phobics. *Behaviour Research and Therapy, 40,* 1031-1046.

Sarlo, M., Palomba, D., Angilli, A., & Stegagno, L. (2002). Blood phobia and spider phobia: two specific phobias with different autonomic cardiac modulations. *Biological Psychology, 60,* (2), 91-108.

Seligman, M.E.P. (1970). On the generality of the laws of learning. *Psychological Review, 77*, 406-418.

Shriffin, R.M., & Gardner, G.T. (1972). Visual processing capacity and attentional control. *Journal of Experimental Psychology*, *93*, 78-82.

Smits, J.A.J., Telch, M.J., & Randall, P.K. (2002). An examination of the decline in fear and disgust during exposure-based treatment. *Behaviour Research and Therapy, 40,* 1243-1253.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74,* 1-29.

Spielberger, C.D., Gorsuch, R.C., & Lushene, R.E. (1970). Manual for the State-Trait Anxiety Inventory. Palo Alto, CA: Consulting Psychologists Press.

Stanley, M.A., & Beck, J.G. (2000). Anxiety Disorders. *Clinical Psychology Review, 20 (6)* 731-754.

Stormark, K.M., Nordby, H., & Hugdahl, K. (1995). Attentional shifts to emotionally charged cues: Behaviourl and ERP data. *Cognition and Emotion, 9 (5)*, 507-623.

Stroop, J.R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643-662.

Taylor, J.A. (1953). A personality scale of manifest anxiety. *Journal of Abnormal and Social Psychology*, *48*, 285-290.

Teachman, B.A., Gregg, A.P., & Woody, S.R. (2001). Implicit associations of fear-relevant stimuli among individuals with snake and spider fears. *Journal of Abnormal Psychology, 110,* 226-235.

Tipples, J., Atkinson, A.P., & Young, A.W. (2002). The Eyebrow Frown: A Salient Social Signal. *Emotion, 2 (3),* 288-296.

Tipples, J., & Sharma, D. (2000). Orienting to exogenous cues and attentional bias to affective pictures reflect separate processes. *British Journal of Psychology, 91,* 87-97.

Tipples, J., Young, A.W., Quinlan, P., Broks, P. & Ellis, A. (2002). Searching for threat. *The Quarterly Journal of Experimental Psychology*, *55 (A) 3*, 1007-1026.

Tolin, D.F., Lohr, J.M., Sawchuk, C.N., & Lee, T.C. (1997). Disgust and disgust sensitivity in blood-injection-injury and spider phobia. *Behaviour Research & Therapy, 35 (10),* 949-953.

Treisman, A. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology. 12,* 242-248.

Treisman, A. (1969). Strategies and models of selective attention. *Psychological Review*, *76,* 282-299.

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.

Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, *114*, 285-310.

Treisman, A. (1993). The perception of features and objects. In A. Baddeley & L. Weiskrantz (Eds.). Attention: Selection, awareness and control (pp. 5-35). Oxford: Clarendon Press.

Tsal, Y. (1983). Movements of attention across the visual field. *Journal of Experimental Psychology: Human Perception and Performance, 9*, 523-530.

Vasey, M.W., ElHag, N., & Daleiden, E.L. (1996). Anxiety and the processing of emotionally threatening stimuli: Distinctive patterns of selective attention among high and low-test anxious children. *Child Development*, *67, (3)* 1173-1185.

Vuilleumer, 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.

Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and popout in visual search. *Perception & Psychophysics, 56*, 495-500.

Whalen, P.J., Rauch, S.L., Etcoff, N.L., McInerney, S.C., Lee, M.B., & Jenikie, M.A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *J. Neurosci, 18,* 411-418.

White, M. (1995). Preattentive Analysis of Facial Expressions of Emotion. *Cognition & Emotion, 9 (5),* 439-460.

Williams, J.M., Mathews, A., & MacLeod, C. (1996). The Emotional Stroop task and psychopathology. *Psychological Bulletin, 120 (1)* 3-24.

Williams, J.M.G., Watts, F.N., MacLeod, C., & Mathews, A. (1988). Cognitive Psychology and Emotional Disorders. Wiley and Sons. Chichester.

Williams, J.M.G., Watts, F.N., MacLeod, C., & Mathews, A. (1997). Cognitive Psychology and Emotional Disorders. (Second Edition). Wiley and Sons. Chichester.

Wolfe, J.M. (1994). Guided Search 2.0. A revised model of visual search. *Psychonomic Bulletin & Review, 2,* 202-238.

Wolfe, J.M. (1998a). Visual Search. In. H. Pashler, (Ed.). Attention (pp. 13-73). Hove: Psychology Press.

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

Wolfe, J.M. (2001). Asymmetries in visual search: An introduction. *Perception & Psychophysics*, *63 (3)* 381-389.

Wolfe, J.M., Cave, K.R., & Franzel, S.L. (1989) Guided Search: An Alternative to the Feature Integration Model of Visual Search. *Journal of Experimental Psychology: Human Perception and Performance, 15 (3)* 419-433.

Wright, C.I., Martis, B., Shin, L.M., Fischer H., & Rauch S.L. (2002) Enhanced amygdala responses to emotional versus neutral schematic facial expressions. *NeuroReport, 13*, 785-790.

Yantis, S., & Hillstrom, A.P. (1994). Stimulus-driven attentional capture and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 95-107.

Yantis, S., & Jonides, J. (1996). Attentional capture by abrupt visual onsets: New perceptual objects or visual masking? *Journal of Experimental Psychology: Human Perception and Performance, 22 (6)* 1505-1513.

Yiend, J., & Mathews, A. (2002). Anxiety and attention to threatening pictures. *The Quarterly Journal of Experimental Psychology, 54 (3)* 665-681.

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

Zajonc, R.B. (1984). Feeling and thinking: On the primacy of affect. *American Psychologist, 39*, 117-124.

Zald, D.H. (2003). The human amygdala and the emotional evaluation of sensory stimuli. *Brain Research Reviews, 41,* 88-123.

Zigmond, A.S., & Snaith, R.P. (1983). The hospital anxiety and depression scale. *Acta Psychiatrica Scandinavia*, *67 (6)* 361-370.

Appendix I

Table of IAPS images used in Experiments 1 and 2a

Negative Valence

Description	Slide Number	Mean Valence (SD)	Mean Arousal (SD)
Mutilation	3000	1.45 (1.20)	7.26 (2.10)
Mutilation	3010	1.71 (1.19)	7.16 (2.24)
Mutilation	3030	1.91 (1.56)	6.76 (2.10)
Mutilation	3062	1.87 (1.31)	5.78 (2.57)
Mutilation	3063	1.49 (0.96)	6.35 (2.60)
Mutilation	3064	1.45 (0.97)	6.41 (2.62)
Mafia Hit	3010	1.79 (1.28)	7.26 (1.86)
Burn Victim	3053	1.31 (0.97)	6.91 (2.57)
Burn Victim	3102	1.40 (1.14)	6.58 (2.69)
Burn Victim	3110	1.79 (1.30)	6.70 (2.16)
Knife	6350	1.90 (1.29)	7.29 (1.87)
Sliced Hand	9405	1.83 (1.17)	6.08 (2.40)
Body	3120	1.56 (1.09)	6.84 (2.36)
Body	3140	1.83 (1.17)	6.36 (1.97)
Aimed Gun	6260	2.44 (1.54)	6.93 (1.93)
Gun in mouth	3530	1.80 (1.32)	6.82 (2.09)
Dead Body	9252	1.98 (1.59)	6.64 (2.33)
Mangled Face	3060	1.79 (1.56)	7.12 (2.09)
Soldier	9410	1.51 (1.15)	7.07 (2.06)
Dog	9570	1.68 (1.23)	6.14 (2.31)
Cat	9571	1.96 (1.50)	5.64 (2.50)
Finger	3150	2.26 (1.57)	6.55 (2.20)
Crying Boy	2800	1.78 (1.14)	5.49 (2.11)
Throat Slash	3071	1.88 (1.39)	6.86 (2.05)

Neutral Valence

Description	Slide Number	Mean Valence (SD)	Mean Arousal (SD)
Man	2190	4.83 (1.28)	2.41 (1.80)
Neutral Face	2200	4.79 (1.38)	3.18 (2.17)
Mushroom	5500	5.42 (1.58)	3.00 (2.42)
Rolling Pin	7000	5.00 (0.84)	2.42 (1.79)
Basket	7010	4.94 (1.07)	1.76 (1.48)
Fan	7020	4.97 (1.04)	2.17 (1.71)
Hair Dryer	7050	4.93 (0.81)	2.75 (1.80)
Fork	7080	5.27 (1.09)	2.32 (1.84)
Book	7090	5.19 (1.46)	2.61 (2.03)
Fire Hydrant	7100	5.24 (1.20)	2.89 (1.70)
Truck	7130	4.77 (1.03)	3.35 (1.90)
Umbrella	7150	4.72 (1.00)	2.61 (1.76)
Painted Fabric	7160	5.02 (1.10)	3.07 (2.07)
Light Bulb	7170	5.14 (1.28)	3.21 (2.05)
Armchair	7180	4.73 (1.31)	3.43 (1.95)
Building	7500	5.33 (1.44)	3.26 (2.18)
Office	7550	5.27 (1.40)	3.95 (1.91)
Office	7700	4.25 (1.45)	2.95 (2.17)
3 Oval Blocks	7187	5.07 (1.02)	2.30 (1.75)

Lamp	7175	4.87 (1.00)	1.72 (1.26)
Sad Face	2230	4.53 (1.22)	4.13 (1.68)
Boy	2410	4.62 (1.72)	4.13 (2.29)
Neutral Girl	2440	4.49 (1.03)	2.63 (1.70)
Elderly Man	2480	4.77 (1.64)	2.66 (1.78)

Appendix II

Please read each statement below and then circle the number that best describes how you *generally* feel. There are no right or wrong answers so do not spend too long on any one statement.

	ALMOST NEVER	SOMETIMES	OFTEN	ALMOST ALWAYS
1. I feel pleasant	1	2	3	4
2. I feel nervous and restless	1	2	3	4
3. I feel satisfied with myself	1	2	3	4
4. I wish I could be as happy as others seem to be	1	2	3	4
5. I feel like a failure	1	2	3	4
6. I feel rested	1	2	3	4
7. I am "calm, cool and collected"	1	2	3	4
8. I feel that difficulties are piling up so that I cannot overcome them	1	2	3	4
9. I worry too much over something that doesn't really matter	1	2	3	4
10. I am happy	1	2	3	4
11. I have disturbing thoughts	1	2	3	4
12. I lack self-confidence	1	2	3	4
13. I feel secure	1	2	3	4
14. I make decisions easily	1	2	3	4
15. I feel inadequate	1	2	3	4
16. I am content	1	2	3	4
17. Some unimportant thought runs through my mind and bothers me	1	2	3	4
18. I take disappointments so keenly that I can't put them out of my mind	1	2	3	4
19. I am a steady person	1	2	3	4
20. I get in a state of tension or turmoil as I think over my recent concerns and interests	1	2	3	4

Appendix III

Word List used for the IAT, Experiment 1.

Positive words: caress, freedom, health, love, peace, cheer, friend, heaven, loyal, pleasure, gentle, honest, lucky, rainbow, gift, honour, miracle, sunrise, laughter, paradise.

Negative words: abuse, crash, filth, murder, sickness, accident, death, grief, poison, stink, assault, disaster, hatred, pollute, tragedy, poverty, ugly, cancer, rotten, vomit.

Appendix IV

Word List used for the IAT, Experiment 2a, 2b, 3 & 5

Positive Words. Caress, Freedom, Health, Love, Peace, Cheer, Friend, Heaven, Loyal, Pleasure, Diamond, Gentle, Honest, Lucky, Rainbow, Diploma, Gift, Honour, Miracle, Sunrise, Family, Happy, Laughter, Paradise, Holiday, Warmth, Sunset, Humour, Surprise, Lush, Joy, Ecstasy.

Negative Words. Abuse, Crash, Filth, Murder, Sickness, Accident, Death, Grief, Poison, Stink, Assault, Disaster, Hatred, Pollute, Tragedy, Bomb, Divorce, Jail, Poverty, Ugly, Cancer, Evil, Rotten, Vomit, Agony, Prison, Tumour, Malice, Scar, Pain, Injure, Crash.

Appendix V

SNAQ

The following questionnaire is comprised of 30 items regarding your feelings about snakes. Please answer *every* question by ticking either the true or false box. Work *quickly* and don't spend too much time on any one question. Remember to give your *first* impression. All information is confidential and will not be seen by anyone other than the experimenter.

Name....................................... e mail....................................

	True	False
1. I avoid going to parks or on camping trips because there may be snakes about.	☐	☐
2. I would feel some anxiety holding a toy snake in my hand.	☐	☐
3. If a picture of a snake appears on screen during a film I turn my head away.	☐	☐
4. I dislike looking at pictures of snakes in a magazine.	☐	☐
5. Although it may not be so, I think of snakes as slimy.	☐	☐
6. I enjoy watching snakes at the zoo.	☐	☐
7. I am terrified by the thought of watching a harmless snake.	☐	☐
8. If somebody says there are snakes about I become alert and on edge.	☐	☐
9. I would not go camping in the forest if I thought there were snakes about.	☐	☐
10. I would feel uncomfortable wearing a snakeskin belt	☐	☐
11. When I see a snake, I feel tense and restless.	☐	☐
12. I enjoy reading articles about snakes and other reptiles.	☐	☐
13. I feel sick when I see a snake.	☐	☐
14. Snakes are sometimes useful.	☐	☐
15. I shudder when I think of snakes.	☐	☐
16. I don't mind being near a non-poisonous snake if there is somebody with me in whom I have confidence.	☐	☐
17. Some snakes are very attractive to look at.	☐	☐
18. I don't believe anyone could hold a snake without some fear.	☐	☐
19. The way snakes move is repulsive.	☐	☐

	True	False
20. It wouldn't bother me to touch a dead snake with a long stick.	☐	☐
21. If I came upon a snake in the woods I would probably run.	☐	☐
22. I am more afraid of snakes than any other animal.	☐	☐
23. I would not want to travel in tropical countries because of the greater prevalence of snakes.	☐	☐
24. I have no fear of non-poisonous snakes.	☐	☐
25. I wouldn't take a course in biology if I thought you had to dissect snakes.	☐	☐
26. Snakes are very graceful animals.	☐	☐
27. I think that I am no more afraid of snakes than the average person.	☐	☐
28. I would prefer not to finish a story if something about snakes was introduced into the plot.	☐	☐
29. Even if I was late for an important appointment, the thought of snakes would stop me from taking a shortcut through an open field.	☐	☐
30. Not only am I afraid of snakes, but worms and most reptiles make me feel anxious.	☐	☐

Appendix VI

SPQ

The following questionnaire is comprised of 31 items regarding your feelings about spiders. Please answer *every* question by ticking either the true or false box. Work *quickly* and don't spend too much time on any one question. Remember to give your *first* impression. All information is confidential and will not be seen by anyone other than the experimenter.

Name.......................................   e mail......................................

	True	False
1. I avoid going to parks or on camping trips because there may be spiders about.	☐	☐
2. I would feel some anxiety holding a toy spider in my hand.	☐	☐
3. If a picture of a spider crawling on a person appears on screen during a film I turn my head away.	☐	☐
4. I dislike looking at pictures of spiders in a magazine.	☐	☐
5. If there is a spider on the ceiling above my bed, I can't go to sleep until somebody kills it for me.	☐	☐
6. I enjoy watching spiders build webs.	☐	☐
7. I am terrified by the thought of watching a harmless spider.	☐	☐
8. If somebody says there are spiders about I become alert and on edge.	☐	☐
9. I would not go down into the basement to get something if I thought there were spiders about.	☐	☐
10. I would feel uncomfortable if a spider crawled out of my shoe as I took it out of the cupboard to put it on	☐	☐
11. When I see a spider, I feel tense and restless.	☐	☐
12. I enjoy reading articles about spiders.	☐	☐
13. I feel sick when I see a spider.	☐	☐
14. Spiders are sometimes useful.	☐	☐
15. I shudder when I think of spiders.	☐	☐
16. I don't mind being near a harmless spider if there is somebody with me in whom I have confidence.	☐	☐
17. Some spiders are very attractive to look at.	☐	☐
18. I don't believe anyone could hold a spider without some fear.	☐	☐
19. The way spiders move is repulsive.	☐	☐

	True	False
20. It wouldn't bother me to touch a dead spider with a long stick.	☐	☐
21. If I came upon a spider whilst cleaning I would probably run.	☐	☐
22. I am more afraid of spiders than any other animal.	☐	☐
23. I would not want to travel to Mexico or Central America because of the greater prevalence of tarantulas.	☐	☐
24. I am cautious when buying fruit because bananas may attract spiders.	☐	☐
25. I wouldn't take a course in biology if it involved handling live spiders.	☐	☐
26. Spider webs are very artistic.	☐	☐
27. I think that I am no more afraid of spiders than the average person.	☐	☐
28. I would prefer not to finish a story if something about spiders was introduced into the plot.	☐	☐
29. Even if I was late for an important appointment, the thought of spiders would stop me from taking a shortcut through an underpass.	☐	☐
30. Not only am I afraid of spiders, but millipedes and caterpillars make me feel anxious.	☐	☐
31. I have no fear of non-poisonous spiders.	☐	☐

Appendix VII

Bonferroni comparisons for distance (Experiment 4).

(I) DISTANCE	(J) DISTANCE	Mean Difference (I-J)	Std.Error	Sig.[a]	95% Confidence Interval for Difference[a] Lower Bound	Upper Bound
0	1	-52.896*	8.556	.000	-78.154	-27.638
	2	-66.662*	9.243	.000	-93.948	-39.375
	3	-121.398*	8.734	.000	-147.181	-95.614
	4	-108.838*	8.932	.000	-135.207	-82.469
1	0	52.896*	8.556	.000	27.638	78.154
	2	-13.766*	3.374	.002	-23.725	-3.806
	3	-68.502*	4.352	.000	-81.350	-55.654
	4	-55.942*	5.887	.000	-73.320	-38.565
2	0	66.662*	9.243	.000	39.375	93.948
	1	13.766*	3.374	.002	3.806	23.725
	3	-54.736*	3.618	.000	-65.416	-44.056
	4	-42.177*	4.762	.000	-56.235	-28.118
3	0	121.398*	8.734	.000	95.614	147.181
	1	68.502*	4.352	.000	55.654	81.350
	2	54.736*	3.618	.000	44.056	65.416
	4	12.559	4.725	.108	-1.389	26.508
4	0	108.838*	8.932	.000	82.469	135.207
	1	55.942*	5.887	.000	38.565	73.320
	2	42.177*	4.762	.000	28.118	56.235
	3	-12.559	4.725	.108	-26.508	1.389

Based on estimated marginal means
*. The mean difference is significant at the .05 level.
a. Adjustment for multiple comparisons: Bonferroni.

Appendix VIII

Bonferroni comparisons for distance for error rates (Experiment 4).

		Mean Difference			95% Confidence Interval for Difference[a]	
(I) DISTANCE	(J) DISTANCE	(I-J)	Std.Error	Sig.[a]	Lower Bound	Upper Bound
0	1	-1.658*	.425	.003	-2.912	-.404
	2	-2.394*	.483	.000	-3.820	-.967
	3	-4.790*	.696	.000	-6.845	-2.735
	4	-3.556*	.685	.000	-5.577	-1.534
1	0	1.658*	.425	.003	.404	2.912
	2	-7.35	.319	.257	-1.676	.205
	3	-3.131*	.589	.000	-4.870	-1.393
	4	-1.897*	.625	.040	-3.742	-5.257E-02
2	0	2.394*	.483	.000	.967	3.820
	1	.735	.319	.257	-.205	1.676
	3	-2.396*	.577	.001	-4.100	-.692
	4	-1.162	.694	1.000	-3.212	.888
3	0	4.790*	.696	.000	2.735	6.845
	1	3.131*	.589	.000	1.393	4.870
	2	2.396*	.577	.001	.692	4.100
	4	1.234	.648	.632	-6.78	3.147
4	0	3.556*	.685	.000	1.534	5.577
	1	1.897*	.625	.040	5.257E-02	3.742
	2	1.162	.694	1.000	-.888	3.212
	3	-1.234	.648	.632	-3.147	.678

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Bonferroni.

Appendix IX

	Trait Anxiety	Mean	SD	Number
NegDist0	Low Anxiety	516.52	61.75	15
	High Anxiety	489.84	60.05	12
	All Subjects	505.86	69.99	48
NegDist1	Low Anxiety	580.08	112.98	15
	High Anxiety	560.16	96.59	12
	All Subjects	575.98	130.64	48
NegDist2	Low Anxiety	598.16	115.85	15
	High Anxiety	574.67	95.75	12
	All Subjects	590.29	131.02	48
NegDist3	Low Anxiety	655.64	112.88	15
	High Anxiety	634.04	106.43	12
	All Subjects	644.42	117.99	48
NegDist4	Low Anxiety	663.38	137.83	15
	High Anxiety	621.64	111.14	12
	All Subjects	627.56	118.23	48
NeutDist0	Low Anxiety	478.57	44.27	15
	High Anxiety	469.00	43.97	12
	All Subjects	481.04	49.12	48
NeutDist1	Low Anxiety	520.65	53.05	15
	High Anxiety	503.58	67.04	12
	All Subjects	516.80	60.66	48
NeutDist2	Low Anxiety	536.88	47.15	15
	High Anxiety	517.69	73.92	12
	All Subjects	527.84	62.25	48
NeutDist3	Low Anxiety	587.26	57.20	15
	High Anxiety	563.30	52.24	12
	All Subjects	583.97	64.77	48
NeutDist4	Low Anxiety	581.99	63.40	15
	High Anxiety	555.39	82.54	12
	All Subjects	566.79	71.37	48

Appendix X

Bonferroni comparisons for distance (Experiment 5).

		Mean Difference (I-J)	Std.Error	Sig.[a]	95% Confidence Interval for Difference[a]	
(I) DISTANCE	(J) DISTANCE				Lower Bound	Upper Bound
0	1	-33.785*	4.631	.000	-47.457	-20.113
	2	-37.154*	3.744	.000	-48.205	-26.103
	3	-96.591*	5.026	.000	-111.429	-81.753
	4	-73.446*	4.828	.000	-87.698	-59.193
1	0	33.785*	4.631	.000	20.113	47.457
	2	-3.369	2.997	1.000	-12.216	5.477
	3	-62.806*	3.537	.000	-73.249	-52.363
	4	-39.661*	4.920	.000	-54.185	-25.137
2	0	37.154*	3.744	.000	26.103	48.205
	1	3.369	2.997	1.000	-5.477	12.216
	3	-59.437*	2.516	.000	-66.865	-52.009
	4	-36.292*	2.944	.000	-44.981	-27.602
3	0	96.591*	5.026	.000	81.753	111.429
	1	62.806*	3.537	.000	52.363	73.249
	2	59.437*	2.516	.000	52.009	66.865
	4	23.145*	3.609	.000	12.492	33.798
4	0	73.446*	4.828	.000	59.193	87.698
	1	39.661*	4.920	.000	25.137	54.185
	2	36.292*	2.944	.000	27.602	44.981
	3	-23.145*	3.609	.000	-33.798	-12.492

Based on estimated marginal means
*. The mean difference is significant at the .05 level.
a. Adjustment for multiple comparisons: Bonferroni.

Appendix XI

Bonferroni comparisons for distance for errors (Experiment 5).

(I) DISTANCE	(J) DISTANCE	Mean Difference (I-J)	Std.Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
0	1	-1.033*	.290	.009	-1.889	-.176
	2	-1.414*	.313	.000	-2.338	-.490
	3	-1.355*	.270	.000	-2.152	-.558
	4	-.841*	.242	.011	-1.554	-.128
1	0	1.033*	.290	.009	.176	1.889
	2	-.381	.322	1.000	-1.332	.569
	3	-.323	.330	1.000	-1.296	.651
	4	.192	.272	1.000	-.610	.994
2	0	1.414*	.313	.000	.490	2.338
	1	.381	.322	1.000	-.569	1.332
	3	5.872E-02	.338	1.000	-.940	1.057
	4	.573	.337	.959	-.421	1.567
3	0	1.355*	.270	.000	.558	2.152
	1	.323	.330	1.000	-.651	1.296
	2	-5.872E-02	.338	1.000	-1.057	.940
	4	-.514	.300	.935	-.372	1.400
4	0	.841*	.242	.011	.128	1.554
	1	-.192	.272	1.000	-.994	.610
	2	-.573	.337	.959	-1.567	.421
	3	-.514	.300	.935	-1.400	.372

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Bonferroni.

Appendix XII

	Trait Anxiety	Mean	SD	Number
NegDist0	Low Anxiety	465.63	46.91	9
	High Anxiety	504.97	93.66	17
	All Subjects	491.35	81.75	48
NegDist1	Low Anxiety	497.75	43.41	9
	High Anxiety	539.04	98.91	17
	Total	524.75	85.24	48
NegDist2	Low Anxiety	501.88	38.80	9
	High Anxiety	545.56	97.56	17
	All Subjects	530.44	83.80	48
NegDist3	Low Anxiety	559.98	39.56	9
	High Anxiety	603.17	108.13	17
	All Subjects	588.22	91.78	48
NegDist4	Low Anxiety	542.33	39.51	9
	High Anxiety	582.69	100.63	17
	All Subjects	568.71	85.82	48
NeutDist0	Low Anxiety	474.62	42.68	9
	High Anxiety	501.37	85.34	17
	All Subjects	492.11	73.57	48
NeutDist1	Low Anxiety	493.15	33.92	9
	High Anxiety	545.27	119.58	17
	All Subjects	527.22	100.80	48
NeutDist2	Low Anxiety	499.54	32.34	9
	High Anxiety	542.03	92.71	17
	All Subjects	527.32	79.13	48
NeutDist3	Low Anxiety	565.59	52.68	9
	High Anxiety	596.11	91.44	17
	All Subjects	585.54	80.36	48
NeutDist4	Low Anxiety	532.72	42.74	9
	High Anxiety	574.83	86.68	17
	All Subjects	560.25	76.23	48